

Strategic Interaction and Interstate Crises: A Bayesian Quantal Response Estimator for Incomplete Information Games

Justin Esarey

*Department of Political Science, Florida State University,
556 Bellamy Building, Tallahassee, FL 32306
e-mail: jee03c@fsu.edu (corresponding author)*

Bumba Mukherjee

*Department of Political Science and Department of Economics and Econometrics,
407 Decio Hall, University of Notre Dame, South Bend, IN 46556
e-mail: bumba_mukherjee@nd.edu*

Will H. Moore

*Department of Political Science, Florida State University,
561 Bellamy Building, Tallahassee, FL 32306
e-mail: will.moore@fsu.edu*

Private information characteristics like *resolve* and *audience costs* are powerful influences over strategic international behavior, especially crisis bargaining. As a consequence, states face asymmetric information when interacting with one another and will presumably try to learn about each others' private characteristics by observing each others' behavior. A satisfying statistical treatment would account for the existence of asymmetric information and model the learning process. This study develops a formal and statistical framework for incomplete information games that we term the *Bayesian Quantal Response Equilibrium Model* (BQRE model). Our BQRE model offers three advantages over existing work: it directly incorporates asymmetric information into the statistical model's structure, estimates the influence of private information characteristics on behavior, and mimics the temporal learning process that we believe takes place in international politics.

1 Introduction

In 1999, Curtis Signorino and Alastair Smith separately argued that the conventional application of logit and probit models produces biased estimates when used to test hypotheses derived from theories of strategic interaction (usually game theoretic models).

Authors' note: For their assistance with the *Phoenix* cluster at Florida State University (FSU) we wish to thank Jim Wilgenbusch and Jeff MacDonald of the Computational and Information Science Technology Department. The Department of Political Science at FSU graciously provided us with access to *Quantify*, its shared parallel computing resource. We also wish to thank Ken Schultz, David Siegel, Curt Signorino, Jeff Staton, our anonymous reviewers, and the participants and attendees of our 2007 APSA conference panel and 2007 Political Methodology conference poster session for their helpful comments and suggestions.

© The Author 2008. Published by Oxford University Press on behalf of the Society for Political Methodology. All rights reserved. For Permissions, please email: journals.permissions@oxfordjournals.org

These authors' papers demonstrated that it is not enough to operationalize variables that match theoretical concepts from the game theoretic model, then estimate an additive logit or probit model with these variables. Three currents in the discipline made this work especially important. First, the use of maximum likelihood techniques, such as logit and probit, to estimate categorical variable models had become commonplace during the 1990s. Second, the rising prominence of models and theories that explicitly assume that political outcomes are produced by a strategic interactive process made the critique relevant to broad areas of work. Third, the National Science Foundation Political Science Program's Empirical Implication of Theoretical Models (EITM) project called for greater attention to the linkage between theories and their empirical tests, thus providing a foundation in which to place this work.¹

Signorino and Smith's critique was especially relevant to the literature on interstate crisis bargaining. If states consider alternative options and the expected response of their adversaries, then select the option they believe will make them best off,² such a process must be reflected in the structure of an empirical test. A variety of techniques for appropriately specifying an empirical test for a strategic theoretical model have been developed, using quantal response equilibrium (QRE) models (Signorino 1999; Signorino and Yilmaz 2003), censoring models (Smith 1999), Perfect Bayesian Equilibrium (PBE) (Lewis and Schultz 2003), and standard discrete choice models with appropriate interaction terms derived via comparative statics (Bas, Signorino, and Walker 2007; Carrubba, Yuen, and Zorn 2007).

Yet none of these efforts are able to effectively capture the signaling and updating dynamics created by hidden-state characteristics (Signorino 2003: 343; Wand 2006), such as a state's *resolve* to stand firm in a crisis (Schelling 1966; Snyder and Diesing 1977; Reed 2003) or the *audience costs* it faces from its public for going to war (Fearon 1994a). Theoretical models that include private information, generally called incomplete information models, are numerous in the literature on interstate conflict.³ This literature emphasizes that, when one state is unaware of the value of other states' audience costs, capabilities, and/or resolve, complex dynamics of signaling and learning emerge that can greatly impact the probability of conflict outbreak.

We posit a formal model of crisis bargaining that (i) captures important aspects of incomplete information games (i.e., updating and signaling) and (ii) implies a statistical model that can be fitted with Bayesian econometric methods. In principle, the stage game we employ could be matched to a standard statistical model using the strategy of comparative statics, but we opt to follow the trail blazed by Signorino (1999) and Smith (1999) and develop a novel statistical model directly derived from action probabilities predicted from our game. Though useful for interstate crisis bargaining scenarios, our model is also generalizable to a wide variety of settings.⁴

We call our integrated formal and statistical framework a *Bayesian Quantal Response Equilibrium Model* (BQRE model). Our model integrates the influence of asymmetric information and learning on behavior directly into the statistical estimation process, creating a statistical model that matches the game model a researcher theorizes as the

¹See Aldrich and Alt (2003) and Granato and Scioli (2004) for overviews of this approach.

²Virtually all stripes of realist theory fit this description, as do neoliberal theories such as those developed, for example, by Keohane (1984). All rational choice-based theories of international politics fit this description.

³For examples of incomplete information models in the literature on interstate conflict, see Morrow (1989, 1999); Powell (2004); Fearon (1994a, 1994b, 1995); Reiter and Stam (1998); Reed (2003); Smith and Stam (2004); Fey and Ramsay (2006); Sartori and Meirowitz (2006); Schultz (1999); and Slantchev (2003).

⁴Our model is applicable to any finite, sequential-move extensive form game of incomplete information, as long as (i) prior beliefs are fixed and deterministic, (ii) players update their beliefs according to Bayes's rule, and (iii) payoffs at every terminal node are affine. See Lemma 1 and the proof in the appendix for details.

data-generating process. Leaders and diplomats presumably learn about one another's unobservable characteristics by repeatedly interacting, observing one another's behavior, and refining their beliefs to be consistent with this observed behavior. This process occurs in our model as well.

We begin with a two-player crisis bargaining stage game in which one state has a choice to challenge another state (or not). If challenged, a defending state can choose to fight the challenge (or acquiesce). Finally, a challenger state that has been fought can follow through with its challenge and start a war (or back down). Because more aggressive states are more likely to challenge, a defending state can learn about a challenging state's type if a challenge takes place. In this way, states learn about each other through dyadic interaction. In addition, leaders learn by observing interactions that they are not involved in. For example, the Cuban Missile Crisis informed all states about United States and Soviet resolve once the crisis had played out. The model we develop accounts for both these kinds of learning: states update their beliefs within the dyadic stage games they play and additionally update their beliefs after observing the outcome of other stage games in which they are not involved.

We are also interested in estimating the normally unobservable type characteristics of states, such as their *resolve* or *audience costs*. In a crisis bargaining game one might assume that, *ceteris paribus*, some states are more aggressive than others. We can empirically estimate this aggressiveness (due to differences in resolve, audience costs, etc.) using our model.

The paper proceeds as follows. Section 2 reviews previous EITM-inspired research in this area, noting where this research places the frontier for our contribution. Section 3 presents the BQRE model in the context of crisis bargaining. Section 4 discusses how BQRE can be applied to an empirical data set and compares the performance of BQRE to PBE and QRE alternatives using Monte Carlo simulations. Section 5 concludes with a discussion of possibilities for future work.

2 Statistical Models of Strategic Interaction

As noted above, the Signorino (1999) and Smith (1999) studies drew critical attention to the common practice of testing hypotheses from a strategic theory with a loosely specified ordinary least squares, probit, or logit model that simply contains an additive combination of operationalized variables.⁵ Although these models are appealing because of the ease with which they can be estimated using widely available software, as conventionally used these models do not account for the role of strategic interaction between actors in a game theoretic model of politics. Hence, these conventional approaches can lead to biased estimates.

Two broad approaches exist to properly estimate the parameters implied by a strategic choice theory. In both cases, one solves a game using a solution concept and then determines an appropriate statistical model given this solution. The first approach, one that has been used for decades, requires that the researcher first produce comparative static predictions, then properly specify and then estimate the implied product or interaction terms for a standard statistical model, for example, probit or logit (Carrubba, Yuen, and Zorn 2007; though see Signorino 2007). A variation on this approach involves estimating multiple probits or logits for each stage of the game, substituting fitted probabilities for nodes lower on the game tree into the probit/logit for a node higher on the tree (Bas,

⁵International Relations scholars often use standard discrete choice models or bivariate probit models to test implications of game theoretic models that predict (i) likelihood of conflict, (ii) occurrence of crises, and (iii) failure of deterrence. See, for example, Partell and Palmer (1999), Schultz (1999), Gelpi and Greisdorf (2001), Reiter and Stam (2002), and Reed (2003).

Signorino, and Walker 2007; Carrubba, Yuen, and Zorn 2007). Although this approach is straightforward for games with unique solutions and can in principle be applied to all games, it can become difficult to apply in complex games. In particular, games of incomplete information often present multiple equilibria, each of which can be equally likely under the solution concept, and hence the comparative static predictions or action probabilities to use are unclear.

A second approach derives the likelihood function directly from the action probabilities predicted by the model. Certain solution concepts (such as QRE) provide refinements that predict (under certain conditions) a unique mixed strategy equilibrium, wherein each agent plays each of her possible actions at a game node with some probability. This feature translates to a likelihood function directly: the probability of an outcome's occurrence in the likelihood is the product of action probabilities necessary to reach that outcome. Hence, a novel likelihood function is created whose free parameters can be estimated via maximum likelihood methods. Signorino, Smith, and Lewis and Schultz have all adopted this approach, and we follow their lead. Different authors have developed different twists on this approach, such as Smith's (1999) Strategically Censored Discrete Choice (SCDC) model, but because our approach is based on QRE, we focus on reviewing this particular approach.⁶ Though still challenging, we believe this approach is simpler to apply to an incomplete information game than the comparative static method because the problem of multiple equilibria is eliminated. McKelvey and Palfrey (1998, Theorem 3) show that QREs are unique in extensive form games of perfect information; in Lemma 1 and the Appendix, we show that uniqueness extends to a class of sequential-move extensive form incomplete information games.

2.1 QRE-based Estimators

How do QRE-based estimators work? Signorino's (1999) approach is based on the Logit QRE solution concept of McKelvey and Palfrey (1995, 1998). Logit QRE works by adding an error component to the players' moves (Palfrey 2007). More specifically, QRE introduces a parameter λ that describes the relationship between payoffs and actions. When this term is equal to zero, the player makes choices at random (i.e., with no relationship to her payoffs). As the parameter approaches infinity, the actor plays the game with less and less error (e.g., she comes closer to responding perfectly to the incentives provided by her payoffs). Whereas McKelvey and Palfrey were interested in isolating the degree of play error in the laboratory, where payoffs are fixed and known, Signorino chose to focus on the relationship between observable behavior and covariates outside the laboratory, where payoffs are unknown but are related to covariates that can be used to fit a statistical model. His approach fixes the error parameter to one by assumption to identify the model, causing the coefficients he estimates using this model to be reduced form parameters that incorporate this error propensity.

Logit QRE has been used to directly test game models of crisis bargaining by students of international relations.⁷ Although useful and instructive, the QRE-based model proposed by Signorino in 1999 does not allow players in a game model to learn about each others' private information (in a crisis bargaining model, their *resolve* or *audience costs*) as they interact over time. Indeed, Signorino (2003: 343) states that the assumptions of his model impose "a price: players cannot update in the model." Because signaling and

⁶Although Smith's SCDC model is an important advance, it assumes complete information and thus—without appropriate revision—cannot be used to test hypotheses from incomplete information models.

⁷See, for example, Bas, Signorino, and Walker (2007).

updating are central to players' behavior in incomplete information games, there is room for improvement upon Signorino's models.

2.2 *The PBE Estimator*

Lewis and Schultz (2003) set out to make this improvement. They proposed a PBE estimator to statistically model games of incomplete information. Their statistical model is built around a two-player crisis bargaining game that accounts for signaling and updating by assuming that a state's payoffs are partially random, with the random component of the payoff being private information. That is, each state knows the distribution from which the other state's payoffs are drawn, but not the specific value (Lewis and Schultz 2003: 347–8). A state knows its own payoffs exactly, including the random component. Hence, from the perspective of an outsider, a state's choices are probabilistic; from the perspective of the state itself, choices are determined by its private information parameter. The game is solved using PBE.

Given the structure of this game, one would expect the estimator to account for signaling and updating dynamics. Yet Wand (2006: 108–9, 117–9) charges that the PBE estimator does not actually account for signaling or updating, he shows that there is no real difference between the defending state's prior and posterior distribution of the initiating state's type after the initiator moves in the PBE statistical model, thus suggesting that the PBE estimator of Lewis and Schultz (2003) does not fully capture updating or signaling dynamics. Furthermore, Wand (2006: 117) claimed that the different results obtained from a PBE estimator compared to a QRE-based statistical model were not due to "intrinsic consequences of signaling and updating dynamics" as suggested by Lewis and Schultz (2003), but instead resulted from these models' different assumptions about the placement of stochastic terms and the size of their variance.

2.3 *Our Model's Role in the EITM Research Program*

The primary objective of our project is to construct and evaluate an original statistical model for incomplete information games that explicitly accounts for updating and signaling dynamics in the estimation process, avoiding the pitfalls raised by Wand. The model that we construct and present here—the *Bayesian Quantal Response Model*—extends existing work in two important ways. First, the BQRE model allows us to estimate the influence of private Bayesian types in incomplete information games, including signaling behavior and consequent Bayesian learning/updating. We do not know of other work that combines the QRE solution concept with Bayesian statistics to both construct and directly estimate statistical models for incomplete information games. Second, our statistical model uses panel or pooled data sets to estimate heterogenous hidden unit-level characteristics (for instance, variation in states' propensity to engage in war) when testing hypotheses implied by incomplete information game theoretic models.

3 **A Dynamic Model of Crisis Bargaining: Quantal Response with Bayesian Updating**

We begin with a verbal sketch of our theoretical model in the context of crisis bargaining, then turn to its formal presentation and the derivation of the statistical model. Like Signorino (1999), Smith (1999), and Lewis and Schultz (2003), we develop a theoretical model of interstate crisis bargaining. In doing so, we seek to produce both a formal and a statistical model that can best capture the dynamics that many scholars of interstate conflict would agree are involved in the process. In particular, we assume incomplete

information that produces both signaling and Bayesian updating as states change their beliefs about other states' types.

3.1 *A Verbal Sketch of the BQRE Model*

How will states behave as they jostle for resources on the world stage? We model crisis bargaining as a series of simultaneously played, directed-dyadic bargaining games. Each state plays every other state as both challenger and defender in a two-player bargaining stage game at a particular point in time t , and all games happen simultaneously (i.e., they begin and end simultaneously) during this time period. In these games, each state has both publicly known characteristics (corresponding, for instance, to its industrial capacity and military strength) and private characteristics (such as the audience cost it pays to back down from a threat or start a war.) Within a dyadic stage game at time t , each state tries to learn about its opponent's characteristics by observing its behavior as the game unfolds. States also learn about this opponent by observing its interactions with other states, but only at the end of the time period t once all games have ended and their outcomes have been revealed. States carry what they have learned about their opponents' private information into the next round of crisis bargaining at time $t + 1$. As such, future crises will be influenced by what has been learned about states' taste for conflict from past crises.

Note that there are two distinct mechanisms by which states learn about one another's types. First, in any given dyadic interaction at time t , the defender can learn about the challenger (i.e., the challenger's behavior reveals information about its type) and change its behavior accordingly at later decision nodes in the same dyadic stage game at time t . Second, once a round of bargaining is completed at time t , all states can observe the play of all the other states and revise their beliefs accordingly before the start of stage games at time $t + 1$. As explained in detail below, we take advantage of Bayesian statistical methods to model this second learning mechanism.

Several features of the model deserve note up front, though we describe these features in greater detail in the formal explication below. First, rather than (implicitly) assume that states choose without error among actions at each choice node when engaged in crisis bargaining, we assume that choice has a random component corresponding to mistakes and/or unsystematic influences on utility. This randomness can occur for a variety of reasons. For example, states do not always have the command and control needed to get followers to carry out a leader's orders. Orders are sometimes misunderstood, even when lower-level military commanders and diplomatic officers attempt to follow them. Further, advisors to a decision maker may misrepresent information in an effort to influence the decision. Intelligence information possessed by the state may be inaccurate, distorting a leader's view of how actions relate to outcomes and payoffs. For all these reasons, we believe that our error assumption is justified in the crisis bargaining scenario. Second, though our model allows signaling within a dyadic bargaining scenario, it does not allow states to strategically bluff about their type with one state at time t in order to deliberately influence outcomes in future crises with other states at later times. This is an unusual feature for a game theoretic model, and we describe it in more detail in Section 3.6.

3.2 *Formal Setup*

Our model can be fleshed out more formally. Our stage game is structurally similar to the crisis bargaining model discussed in Lewis and Schultz (2003: 349) to facilitate comparison, though private information is handled differently. We present the extensive form of this game in Fig. 1. This game has two players, state A (challenger) and state B (defender),

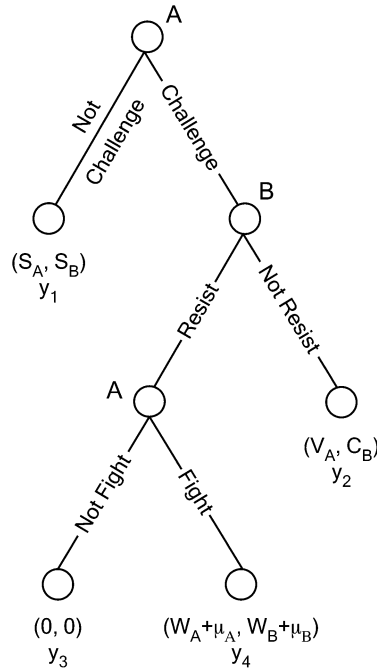


Fig. 1 A crisis bargaining game.

where state *A* hopes to possess a good currently held by state *B*. State *A* moves first, choosing whether to challenge state *B* for the good. If state *A* chooses not to make a challenge, the status quo (labeled as y_1) prevails. If *A* does make a challenge, then state *B* decides whether or not to resist *A*'s demand for the good. If state *B* does not resist, the crisis ends peacefully with *B* conceding the good to *A* (outcome y_2). If state *B* resists, then *A* must decide whether or not to fight. If *A* chooses not to fight, then it backs down from its threat (outcome y_3) and state *B* retains its control over the good. If *A* decides to fight and goes ahead with its demand for the good possessed by *B*, then the game ends with a war (outcome y_4).

State *A*'s payoff from the status quo is S_A , whereas state *A*'s payoff if *B* concedes is V_A . State *B*'s payoff from the status quo is S_B , whereas state *B*'s payoff from conceding to *A* is C_B . Outcome y_3 (*A* challenges, *B* resists, *A* backs down) produces payoffs of zero to both states, an assumption we make for analytic convenience that, as described below, makes it possible to identify the model's parameters. Finally, state *A*'s payoff from war can thus be written as $W_A + \mu_A$, whereas state *B*'s payoff from war is $W_B + \mu_B$. We describe the components of these payoffs in turn.

S_A , V_A , W_A , and W_B correspond to a state's publicly observable components of utility. For example, a state with a larger military than its opponent may expect positive payoffs from going to war due to its high probability of success in such a conflict. Likewise, its opponent may perceive negative payoffs from war. In this case, the relative military capabilities of a state will impact the values of W_A and W_B . A state with a large economy may expect greater payoffs to the status quo than states with less to lose, and hence state *A*'s economic output may partially determine S_A . These capabilities are assumed to be publicly known before the start of the game.

The μ_A and μ_B parameters correspond to state *A*'s and *B*'s private components of utility, respectively, and capture each state's aversion against (or propensity for) war outside of

the payoffs already modeled within the game structure. This propensity has been called *resolve* in some contexts (e.g., Schelling 1966; Snyder and Diesing 1977; Reed 2003) and *audience costs* in others (e.g., Fearon 1994). A state's taste for war is constant (not random)⁸ and unknown to the outside world (including researchers), but other states will form beliefs about it and update those beliefs based on observable actions taken by the state. We assume that other states only have beliefs about what these private parameters might be, beliefs that they update via Bayes's rule as they observe a state's behavior. A state that frequently follows through on its military threats, for example, will become known as a state with a high resolve, an aggressive state, and so on.⁹

Because we believe that a state's aggressiveness is not a random variable but is related to unobservable characteristics of the state, it is appropriate to assume that μ_A and μ_B are nonstochastic and are not necessarily drawn from a common distribution for all states. In addition to accurately reflecting our intuition, this specification of the μ parameters has an important statistical implication. Many statistical analyses implicitly assume that all states are, *ceteris paribus*, equally likely to engage in aggressive expansionary behavior, stand firm when challenged, and so forth. This is problematic since ignoring unit-level heterogeneity will bias statistical estimation by introducing correlation between the residuals and the regressors if types are correlated with observable variables included in the estimated payoff terms (Greene 2003; Hsiao 2003). These biases are particularly pronounced in a strategically interdependent setting, like that of a game theory model. For example, in an international relations scenario, not only do states behave according to their own types but they may also modify their behavior based on the *other* state's type.

Because interstate interaction unfolds over time, a single iteration stage game is less than ideal. Nonetheless, International Relations scholars have been able to tease a number of important implications from such games. Unfortunately, when one moves to data to test the implications of those models one is immediately confronted with ad hoc choices that are frequently treated as statistical nuisances (e.g., the panel properties of pooled cross-sectional time-series data). Although this approach makes data analysis tractable, it does so at the cost of producing a gap between the theory and the empirical test of its hypotheses and thus leaves room for improvement. In an effort to reduce that gap and create a tight link between theory and statistical analysis, we take a novel approach to the temporal dimension of our game. We justify this theoretical move on the grounds that it makes possible a tight connection between a game theoretic model and a statistical model.

In our model, states play the crisis bargaining game against every other state in directed-dyadic fashion; that is, at every time t every state plays as both challenger (state A) and defender (state B) with each of its partners, giving every state the opportunity to be the aggressor. Because a state's history of play is ex post publicly observable information (i.e., states know how all other states have played with all their partners in past interactions), states have identical information at the start of every time period and thus should share identical prior beliefs about each others' private information at the beginning of each time period. Beliefs may diverge during time t as each dyadic game progresses; for example, if state A challenges state B , then state B will update his beliefs about μ_A in a way that other states interacting with A in other dyads will not. But at the end of time t , when all players

⁸To keep the presentation clear and the model tractable, this paper consistently invokes an assumption of unchanging taste for war. This is a simplification that can be relaxed to, for example, allow for trends over time.

⁹We assume that the audience costs or resolve characteristics are specific to a state, not to a dyad; that is, the United States displays the same audience costs in fighting Iran that it might in fighting Iraq, Afghanistan, or Grenada. This assumption can be relaxed, though it greatly increases the ratio of parameters to data points in the model.

observe outcomes in other dyads and hence all have the same information to use for updating their beliefs, these beliefs will again converge before the beginning of the next round of play at $t + 1$.

In principle, one could set up this model as a dynamic game of incomplete information, solve it analytically, derive comparative statics, and then specify an appropriate functional form for the independent variables of a standard maximum likelihood estimate discrete choice model (per Carrubba, Yuen, and Zorn 2007). Unfortunately, the analytical solution and comparative static derivations are far from trivial, and as complex as the technique presented here may appear to be, we found it far more analytically tractable than this alternative. In particular, we extend the work of McKelvey and Palfrey (1998, Theorem 3) to show in the Appendix that the QRE solution concept uniquely selects a single equilibrium under certain conditions, obviating the need to select among equilibria when defining a likelihood function based on the theoretical model. The appeal of this approach is in its analytic convenience, and it is not without its limitations (as will be detailed below.)

3.3 Solving the Static Stage Game

How can predictions of state behavior be derived from our theoretical model? We start by considering the static interaction between two states in a dyad at time t inside a stage game, then build from this solution to incorporate temporal dynamics. Our approach to solving the static portion of the model is based on the agent QRE solution concept articulated by McKelvey and Palfrey (1998) and further extended by Signorino (1999, 2003), hence the QRE portion of the BQRE model. This solution concept implies that states take actions probabilistically,¹⁰ with a state's likelihood of taking an action proportional to the relative utility it expects to derive from that action compared to alternatives. As noted above, we assume that states will not always make the optimum choice given their objective payoffs. The QRE provides an explicit mechanism to model this process as random deviation from expected choices given payoffs. More explicitly, we assume that states are more likely to take actions that yield more utility to them than alternatives, and become ever more likely as this relative utility advantage grows.

Let $U_i(a_j)$ be the utility of action j to state i . We assume that $U_i(a_j) = V_i(a_j) + u_j$, where $V_i(a_j)$ is the expected value to state i of action j and u_j is the error in utility. The expected value of taking action a_j , $V_i(a_j) = \sum_k O_i(h_k) \times \Pr(h_k)$, where $O_i(h_k)$ is the payoff to player i of reaching terminal history h_k and $\Pr(h_k)$ is the probability of reaching terminal history h_k given the equilibrium strategies of both players. Note that in the agent QRE model, $\Pr(h_k) = \prod_r \Pr(a_r)$, $a_r \in h_k$. In this framework, state i 's probability of choosing action a_1 over a_2 at any node with two alternatives is

$$\begin{aligned} \Pr(A_1) &= \Pr(U_i(a_1) > U_i(a_2)) \\ &= \Pr(V_i(a_1) + u_1 > V_i(a_2) + u_2) \\ &= \Pr(u_2 - u_1 < V_i(a_1) - V_i(a_2)). \end{aligned} \quad (1)$$

¹⁰The agent QRE assumes that each node in a game is operated by an independent agent, and as noted above, prevented Bayesian updating which in Signorino's (2003) model. Though our model shares the agent error structure of Signorino's model, we do not interpret this error as private information. Rather, the "agent error" in utility at each node is purely random error, representing transient and unsystematic influences on behavior about which updating is not possible.

Rather than treating each error component of utility individually, we assume instead that the difference in utility between the two options at each choice node, $\varepsilon_{(1-2)} = u_1 - u_2$, takes a logistic distribution:¹¹

$$\varepsilon_{(2-1)} \sim \frac{\exp(-x)}{1 + \exp(-x)}. \quad (2)$$

With this structure in hand, we can write down the probability that state i will take action a_1 as

$$\begin{aligned} \Pr(a_1) &= \Pr(\varepsilon_{(2-1)} < V_i(a_1) - V_i(a_2)) \\ &= \frac{1}{1 + \exp[-(V_i(a_1) - V_i(a_2))]} \\ &= \frac{\exp[V_i(a_1)]}{\exp[V_i(a_1)] + \exp[V_i(a_2)]}. \end{aligned} \quad (3)$$

Since there are only two actions available at each information set in the crisis bargaining game of Fig. 1, $\Pr(A_2) = 1 - \Pr(A_1)$.

The BQRE solution of the crisis bargaining game in Fig. 1 involves calculating QRE strategies where (i) each state i 's type is captured by the parameter μ_i and (ii) states try to infer via Bayesian updating their opponent's type from its actions. Recall from the introduction to this section that μ_i is a constant, nonrandom parameter on state i 's payoff to war that is private information specific to state i ; other states form and update beliefs about μ_i through interaction with state i . This constant parameter μ_i allows our model to account for unit heterogeneity. We show below that the QRE solution explicitly accounts for updating by each player about their opponent's type via Bayes's rule after observing their opponent's action.

Let the parameter μ_i be fixed for each state i . That is, an individual state i 's bias toward a strategy is a fixed state-specific parameter. Suppose that state i does not know μ_j , $j \neq i$, but does know its own μ_i . Based on equation (3), the probability of action at every node of the crisis bargaining game in Fig. 1 is defined as follows:

$$\Pr(\text{Challenge}) = \frac{1}{1 + \exp[V_1(\text{Not Challenge}) - V_1(\text{Challenge})]}. \quad (4)$$

$$\Pr(\text{Resist}) = \frac{1}{1 + \exp[V_2(\text{Not Resist}) - V_2(\text{Resist})]}. \quad (5)$$

$$\Pr(\text{Fight}) = \frac{1}{1 + \exp[V_1(\text{Not Fight}) - V_1(\text{Fight})]}. \quad (6)$$

¹¹The assumption of treating the difference in utility as a random variable, rather than each utility error component as a separate random variable, does two things for our analysis. First, it simplifies the mathematics involved in our calculations and empirical estimations without loss of generality. Second, it makes at least as much substantive sense to conceive of individuals making random comparisons of two alternatives as a unit rather than making random assessments of each object individually and then combining these random assessments to make a decision. Identical results can be derived by assuming that u_1 and u_2 take independent and identical Type I extreme value distributions; this assumption also allows the framework to be extended to games with more than two actions at a decision node (see McKelvey and Palfrey 1995, 1998).

For notational convenience, we label $\Pr(\text{Fight})$ in equation (6) as $\Pr(F)$, $\Pr(\text{Resist})$ in equation (5) as $\Pr(R)$, and $\Pr(\text{Challenge})$ in equation (4) as $\Pr(C)$. Having defined the basic form of the solution for the probability of challenging by state A (equation (4)), the probability of state B resisting state A 's challenge (equation (5)), and the probability of state A fighting when state B resists (equation (6)), we derive below the full functional form of the QRE solutions for $\Pr(C)$, $\Pr(R)$, and $\Pr(F)$ taking into account that the two states update beliefs about their opponent's type by using Bayes's rule.

We begin with $\Pr(F)$. Note that for the last node in the game, at which state A decides to fight a war or back down from the challenge it previously issued, state B 's payoffs do not enter into the expression and there is no uncertainty about the expected payoffs V_1 that will be received as a result of either action. Hence, the payoffs for state A are $V_1(\text{Not Fight}) = 0$ and $V_1(\text{Fight}) = W_A + \mu_A$; recall that μ_A is not random and its value is known to A . Given the aforementioned payoffs and the basic form of the solution of $\Pr(F)$ in equation (6), the QRE solution for state A 's probability of fighting is

$$\Pr(F) = \frac{1}{1 + \exp[-(W_A + \mu_A)]}. \quad (7)$$

Moving up the game tree to solve for the QRE solution of $\Pr(R)$ for state B , we find that the solution becomes less clear because the payoffs $V_2(\text{Not Resist})$ and $V_2(\text{Resist})$ are somewhat complicated to write down. In the case of resistance, state B 's payoff is dependent on state A 's later probability of fighting a war. This probability, as we have just seen, is contingent on the values of W_A as well as state A 's private information about its payoff from war μ_A . If state B 's prior belief about μ_A can be updated contingent on the fact that state A challenged the status quo, then state B should be able to make use of this information in deciding whether to resist the challenge or capitulate to A . Note that, though state B has a belief about μ_A that takes the form of a probability distribution, we assume that B forms a constant expectation about the value from resisting a challenge when making its decision.¹² We capture state B 's updating of its prior belief about μ_A —after observing state A 's challenge to the status quo—in two steps. First, we define the expected values of “not resisting” and “resisting” for state B after A challenges the status quo. The expected value of not resisting for state B is $V_2(\text{Not Resist}) = C_B$. The expected value of resisting is $V_2(\text{Resist}) = E[\Pr(F|\text{Challenge}) \times (W_B + \mu_B)]$. After observing state A challenge the status quo, state B can update its prior belief about μ_A :

$$\begin{aligned} V_2(\text{Resist}) &= E \left[\frac{1}{1 + \exp(-(W_A + \mu_A))} \right] (W_B + \mu_B) \\ &= \left[\int \Pr(F) \times f(\mu_A|C) d\mu_A \right] (W_B + \mu_B), \end{aligned} \quad (8)$$

where $f(\mu_A|C)$ is the updated probability distribution function for state B 's beliefs about μ_A .

¹²Note how $\varepsilon_{(2-1)}$ (from equations 2 and 3) and μ_A enter differently into B 's decision process. The $\varepsilon_{(2-1)}$ parameter represents a random error component whose specific realization will influence a player's judgment of which option yields greater utility. The realization of $\varepsilon_{(2-1)}$ changes from observation to observation due to randomness, causing B 's behavior to be random in turn. By contrast, μ_A is not random, but its value is unknown to B and hence B has a probabilistic belief about the possible values of μ_A . There are not multiple realizations of μ_A that will influence B 's judgment depending on which one is drawn. Rather, B considers the distribution of its beliefs about μ_A and makes a judgment about the expected utility from resisting a challenge. The influence of μ_A on B 's decision making is therefore not stochastic, and no additional randomness is introduced into B 's decision making by the presence of μ_A . Hence, it is appropriate to assume (as we have) that B will take expectations over its beliefs about μ_A when assessing the relative utility of resisting and capitulating.

Second, since state B updates its prior belief about μ_A via Bayes's rule, we can write an expression for $f(\mu_A|C)$ in equation (8) by using Bayes's rule itself:

$$f(\mu_A | C) = \frac{f(C | \mu_A) \times f(\mu_A)}{\int f(C | \mu_A) \times f(\mu_A) d\mu_A}. \quad (9)$$

Let $\Pr(C|\mu_A)$ be the conditional probability that state A challenges state B given its type μ_A . From equation (8), equation (9), the fact that $V_2(\text{Not Resist}) = C_B$, and the basic form of $\Pr(R)$ in equation (5), we can write the QRE solution of $\Pr(R)$ as

$$\Pr(R) = \frac{1}{1 + \exp\left(C_B - \left[\int \Pr(F) \times \frac{\Pr(C|\mu_A)}{\int \Pr(C|\mu_A) \times f(\mu_A) d\mu_A} d\mu_A \right] \times (W_B + \mu_B)\right)}. \quad (10)$$

We now move further up the game tree to solve for the QRE solution of $\Pr(C)$. To solve for $\Pr(C)$, we need to write down $V_1(\text{Not Challenge})$ and $V_1(\text{Challenge})$ for state A . State A 's expected value of not challenging is $V_1(\text{Not Challenge}) = S_A$, whereas the expected value of challenging state B is $V_1(\text{challenge}) = E[\Pr(R)] \times (\Pr(F) \times (W_A + \mu_A) + (1 - E[\Pr(R)]) \times V_A)$. Note that the expectation $E[\Pr(R)]$ is over state A 's prior beliefs about μ_B . From the expected values of $V_1(\text{Challenge})$ and $V_1(\text{Not Challenge})$ and the basic form of the solution of $\Pr(C)$ in equation (4), we obtain the following QRE solution of the probability that state A challenges state B :

$$\Pr(C) = \frac{1}{1 + \exp(S_A - (E[\Pr(R)] \times \Pr(F) \times (W_A + \mu_A) + (1 - E[\Pr(R)]) \times V_A))}. \quad (11)$$

From equations (10) and (11), one can observe that deriving the full functional form of the QRE solution is analytically difficult. To solve for $\Pr(C)$, we would need to substitute the solution of $\Pr(R)$ from equation (10) in $\Pr(C)$ in equation (11) to obtain the full functional form of $\Pr(C)$. Fortunately, we do not need to find an explicit and analytically tractable closed form solution for $\Pr(C)$ because we can prove more generally the existence and uniqueness of a QRE solution for finite extensive form games of incomplete information under certain conditions (including the game in Fig. 1).

Lemma 1. *There exists a unique BQRE solution for each choice probability for finite extensive form games of incomplete information when the players (i) have prior beliefs about their opponent's private information that are fixed and deterministic, (ii) update their posterior beliefs by using Bayes's rule and (iii) have affine payoffs at every terminal node of the game.*

Proof: See Appendix.

In short, even though their inherent complexity prohibits us from analytically solving for closed form expressions of $\Pr(C)$ and $\Pr(R)$, we know that there exists a unique QRE solution in the crisis bargaining game with types and Bayesian updating, and we can therefore compute this solution using numerical methods.

3.4 Deriving a Likelihood Function from the Static Stage Game

The solutions in equations (9–11) can be used to construct a statistical model that explicitly accounts for strategic interaction between states as well as the role of player types and

Bayesian updating in the estimation process. It is straightforward to construct a likelihood function using the solutions that we derived earlier for the three choice probabilities in the crisis game: $\Pr(F)$, $\Pr(R)$, and $\Pr(C)$. Because our solutions, which are included in the likelihood function, explicitly account for Bayesian updating and signaling, our model can account for the impact of types and updating of beliefs by states when we estimate our likelihood function using a data set.

Let $y_q^{(p)}$ be the p th outcome observation (i.e., a terminal history with an associated outcome, like war or capitulation), where q indexes the four possible outcomes (as indicated in Fig. 1). A comprehensive likelihood function of the crisis bargaining game given the solutions for $\Pr(F)$, $\Pr(R)$, and $\Pr(C)$ is

$$L = \prod_p [1 - \Pr(C)]^{y_1^{(p)}} \times [\Pr(C) \times (1 - \Pr(R))]^{y_2^{(p)}} \times [\Pr(C) \times \Pr(R) \times (1 - \Pr(F))]^{y_3^{(p)}} \times [\Pr(C) \times \Pr(R) \times \Pr(F)]^{y_4^{(p)}}. \quad (12)$$

The functional forms of $\Pr(R)$, $\Pr(C)$, and $\Pr(F)$ in the above likelihood function are their solutions defined in equations (10), (11), and (7), respectively. The goal of estimating coefficients using equation (12) is to determine the mapping of observable covariates onto the utility space of states' assessment of crisis bargaining outcomes and the influence of private information on that assessment. This information can then be used to determine the influence of these covariates (and types) on the probability of observable outcomes, such as war. Observable covariates enter into the estimation as determinants of $V_i(a_j)$, the expected value of action j to state i . That is, $V_i(a_j) = X\beta + D\mu$, where X is a vector of observable variables, β maps these variables into utility space, D is a vector of dummy variables corresponding to states' identities, and μ maps these identities into type/utility space. The free parameters to be estimated are β and μ .

Two critical features of the likelihood function are worth noting here. First, because our quantal response solutions explicitly account for Bayesian updating and signaling dynamics, we can capture Bayesian updating in the statistical model for the crisis game when estimating the likelihood function in equation (12) on data. Second, the likelihood function defined above also accounts for unit-level heterogeneity by estimating parameter values for μ_i for all players i . Recall that these μ_i represent individual-specific differences in each player's payoff. Thus, we not only capture the role of player types in the estimation process but also alleviate concerns about omitted variable bias when estimating the likelihood in (12) on panel data sets of international conflict.

Identification can be an issue for strategic models for a variety of reasons, which are thoroughly discussed in Lewis and Schultz (2003: 359–361). To identify our model, we follow Lewis and Schultz and set the payoffs from the status quo for each state as 0; that is, S_A and $S_B = 0$. We also set the payoff to both states if A chooses not to fight after B resists as 0 and the payoff to B if it capitulates, C_B , as 0. We do so because the payoffs from capitulation for both states and victory for state B are not separately identifiable from the war payoffs (W_A and W_B). Hence, to estimate the model on data, we need to assign covariates to three of the payoffs in the crisis game: victory by state A , V_A , that is state A 's payoff for the outcome when state A threatens and state B capitulates; war for state A , W_A , that is state A 's payoff for the war outcome; and war for state B , W_B , that is state B 's payoff for the war outcome.¹³

¹³As in Lewis and Schultz (2003), these identification restrictions are nonneutral though unavoidable; they amount to estimating reduced form coefficients on V_A , W_A , and W_B that combine the influence of covariates on multiple terminal nodes.

3.5 Making the Model Dynamic

Thus far, our exposition has focused on the interaction between two states in a single play of the stage game. We must now consider how to formally capture the idea that states learn about each others' types over time as repeated rounds of the game are played. Recall that at the end of each time period of dyadic play, all actions in that time period become public information and states update their beliefs about each others' types accordingly. Let $f_t(\mu)$ be the distribution of prior beliefs about states' μ_i at time t .¹⁴ If states follow Bayes's rule (the B in BQRE), they will update according to:

$$f_t(\mu | y_t) \propto f(y_t | \mu) f_t(\mu). \quad (13)$$

That is, states' posterior beliefs about μ at the end of time t will be the product of the likelihood of other states' actions during time t and prior beliefs about μ at the beginning of time t .¹⁵ Note that equation (13) bears a close resemblance to the form of a posterior distribution of parameters from a Bayesian statistical model using the likelihood defined in equation (12). The full posterior distribution for all free parameters of our model (μ and β) with n states is

$$f_r(\mu, \beta | y_t) \propto f(y_t | \mu, \beta) \times f_t(\mu, \beta). \quad (14)$$

In equation (14), $f(y_t | \mu, \beta)$ is the likelihood function in (12). $f_t(\mu, \beta)$ is the prior distribution on state-specific type parameters and beta coefficients, respectively. Herein lies the advantage of using Bayesian estimation methods for our model: the estimation process mimics (and estimates) the year-to-year Bayesian updating of prior beliefs produced by our theory. Even though the μ_i parameters are private information unknown to the analyst and to states other than i , the analyst learns about these parameters via the same process that states are theorized to use. Hence, the model tells the analyst about the value of μ_i and about what other states believe about μ_i in a single step.

Specifying the prior is usually a quarrelsome matter for Bayesian studies, particularly when there are no previous studies' results to incorporate and/or those studies' results are inapplicable to the case at hand. But in our case, the Bayesian framework of our theoretical model makes incorporating the updating dynamics described above a reasonably simple matter: a state's prior belief $f_t(\mu)$ should be the posterior belief $f_{t-1}(\mu | y_{t-1})$ of the past year's interaction.¹⁶ That is

$$f_t(\mu) = \int f_{t-1}(\mu, \beta | y) d\beta. \quad (15)$$

If a researcher estimates coefficients on a panel year by year, this researcher can use the posterior draws of μ from time $t - 1$ to compute an empirical estimate of $f_t(\mu)$ and use this estimate to calculate the likelihood function at time t . Indeed, this process mirrors the one in which we expect the real-life state to engage: the past history of state i 's interactions in the international community provides the basis for a posterior distribution $f_{t-1}(\mu_i | y_{t-1})$

¹⁴That is, $\mu = \cup_{i=1, \dots, n} \mu_i$ when there are n states.

¹⁵This expression must be divided by a (solvable) constant in order to make the expression a proper probability density.

¹⁶Of course, an (arbitrary) prior must still be specified for $t = 1$. Our Monte Carlo simulations assume a quasi-informative prior (see the description in Section 4) and we have also performed unreported simulations using an uninformative flat prior, but others may adopt alternative prior structures.

that now, at time t , forms the prior distribution of their beliefs $f_t(\mu_i)$. A researcher using Markov chain Monte Carlo (MCMC) to estimate the model can calculate $f_t(\mu)$ by examining the draws of μ from the past year (implicitly integrating out the β parameters) without added computational costs; to estimate a panel data set, we must draw these samples anyway to obtain estimates of β and μ at time $t - 1$, meaning that no additional work needs to be performed to calculate the prior at time t .¹⁷

3.6 Advantages and Limitations of BQRE

Several characteristics of our BQRE crisis bargaining model deserve additional discussion. Our model improves on past efforts in three ways, but (like all statistical models) is limited in the scope of interactions that it can capture. We first consider our model's improvements, and then its limitations.

First, note that the static game's solution directly incorporates Bayesian updating. This can be seen in equation (10), where state B 's updating behavior in $\Pr(R)$ is captured by the term

$$f(\mu_A | Challenge) = \frac{\Pr(C | \mu_A) \times f(\mu_A)}{\int \Pr(C | \mu_A) \times f(\mu_A) d\mu_A}.$$

This expression is derived from Bayes's rule. Substantively speaking, B will learn about A 's type μ_A via Bayes's rule by observing A 's actions.

Second, observe that our solution directly incorporates signaling dynamics. That is, state A in a dyad anticipates that its actions will cause B to update its beliefs about A 's taste for conflict. This phenomenon can be seen in equation (11): the presence of $\Pr(R)$ in our quantal response solution of $\Pr(C)$ (and the fact that $\Pr(R)$ is itself dependent on $\Pr(C|\mu_A)$) indicates that state A expects state B to use Bayes's rule to update its prior about state A 's type after observing state A 's action. This is an important point: state A 's anticipation that state B updates its beliefs about μ_A influences state A 's behavior in the crisis game's first stage. Hence, state A is mindful of the signal that his challenge sends to state B about μ_A , and plays accordingly.

Finally, we point out that our statistical model mimics the Bayesian learning process that states themselves engage in. Of course, our structure is appropriate only if the theoretical model of learning that we have advanced is correct. If correct, however, our model enables a researcher to empirically track the development of state beliefs about each others' private characteristics (e.g., resolve) from a data set using equation (15). We believe that this is a noteworthy aspect of our approach.

There are two opportunities for improvement that will be of interest to researchers. First, forward-looking trigger-type strategies (as in the Folk Theorem) are not explicitly modeled in our framework. The impact of this limitation is somewhat reduced because these strategies can be captured via a modification. Recall that forward-looking dynamic game strategies can be broken down into individual stage-game strategies to be played under contingent circumstances. For example, if states are playing a Grim Trigger strategy, an element in the payoff function indicating whether a dyad partner had activated the trigger in the past (challenged, resisted, fought, or some combination thereof) can be

¹⁷To clarify, we do not expect that states use an MCMC technique to estimate and update distributions. We assume that states update their beliefs according to equation (13), without specifying how they do so. We mean to point out that a *researcher* can perform this updating process via MCMC, and that doing so reduces computational costs.

placed in the payoff functions, where it would be overwhelmingly determinative of payoffs and therefore determinative of behavior. If tit-for-tat play exists in the crisis bargaining game, it can similarly be modeled by placing lagged outcomes of the game in the payoff functions. In this case, if a state (in its role as challenger) maintained the status quo in the past (i.e., it cooperated) with a partner, that partner (in its role of challenger) would find going to war and forcing the other state's capitulation to be less attractive (i.e., it would derive less expected utility from these payoffs) in the future. If economic cooperation between states tends to make war less attractive by raising the value of the status quo (and lowering the value of war, wherein such economic cooperation will be destroyed), the inclusion of measures of economic cooperation in the payoff functions will capture this effect (Russett and Oneal 1999, 2001).

Second, as discussed above, signaling dynamics are captured by our model *within* a dyad as seen in equation (11). However, our model assumes that signaling happens only within dyads at time t , not between dyads or intertemporally. For example, our model presumes that A will not spark a crisis with B in order to influence some other state C 's beliefs about μ_A in future interactions. Although this assumption is necessary to make our model tractable, we believe that a goal of future work should be to relax the assumption to allow for intertemporal bluffing.

4 Empirical Estimation of BQRE Models

The previous section developed the BQRE model from a purely analytical perspective. In this section, we discuss how the model can be applied to data to obtain statistical estimates of the effect of utility covariates and private information types on behavior, including state behavior in crisis bargaining games.

As mentioned previously, our theoretical model likens crisis bargaining to a series of directed-dyadic games, with each state playing with every other state in the roles of challenger (state A) and defender (state B). Hence, an empirical data set on which the estimator will be applied must be in directed-dyad panel format, with n states and T time periods. The estimator starts for the first time period on the data set; at this stage, some reasonable (e.g., flat, diffuse, or empirically based) prior must be specified, as is common for Bayesian procedures. The estimator (based on MCMC methods) computes samples for the posterior distribution at time $t = 1$, which become the prior distribution for the estimation of data at time 2. This process repeats for the entire data set, until finally posterior estimates are computed for the final time period T .

To verify the accuracy of our estimator and compare it to alternatives, we estimated our model (BQRE), a standard QRE model, and a PBE model on Monte Carlo data sets that were generated from the incomplete information game in Fig. 1. The QRE model assumed homogeneity among countries (i.e., all $\mu = 0$), whereas our model estimated these parameters as described above. The data sets were generated from a BQRE process assuming that 20 countries interacted at a given point in time (a total of 380 directed-dyadic observations per Monte Carlo data set). Values for each utility payoff in each directed dyad (V_A , W_A , and W_B) were created by drawing a covariate value from the normal distribution (mean = 0, standard deviation = 1) and multiplying that by the appropriate beta coefficient. Three beta coefficients (one for V_A , W_A , and W_B , respectively) were drawn from the normal distribution for each Monte Carlo data set (mean = 0, standard deviation = 4). A value for μ_i was drawn for each country from the logistic distribution (mean = 0, scale = 3).

For the BQRE and QRE models, we calculated approximations of the mode of the posterior distribution for 1000 Monte Carlo data sets distributed across five desktop

Table 1 Squared error of Monte Carlo estimates

	<i>BQRE</i>		<i>QRE</i>		<i>PBE</i>	
	<i>Median</i> ε^2	<i>Mean</i> ε^2	<i>Median</i> ε^2	<i>Mean</i> ε^2	<i>Median</i> ε^2	<i>Mean</i> ε^2
V_A	0.05439	0.2277	0.3065	2.366	4.715	10.61
W_A	0.06097	0.1799	1.950	4.549	3.954	10.05
W_B	0.05952	0.1990	1.760	4.748	6.762	14.26

computers using maximum likelihood methods (specifically, the *optim* package in R 2.3.1). That is, maximum likelihood estimates were computed for β and μ using equation (14), with $f(y | \mu, \beta)$ as specified in equation (12). The BQRE and QRE estimates assumed quasi-informative priors ($f(\mu)$ and $f(\beta)$ in equation (14)) on all parameters with a mean equal to the true parameter plus a normally distributed random component (mean = 0, standard deviation = 1), and with prior variance equal to a random component from an inverse gamma distribution (shape = 1) implemented in *MCMCpack* (Martin and Quinn 2006).¹⁸ A L'Ecuyer random number generator (*rlecuyer*) was used to ensure that random numbers used to generate the data sets would not be correlated across the computers (Sevcikova and Rossini 2004). Replication code for these simulations is available online at the *Political Analysis* Web site.

For the PBE model, we calculated maximum likelihood estimates for 100 data sets using two multiprocessor computers. The estimates were computed using genetic optimization with derivatives as implemented in *rgenoud* (Sekhon and Mebane 1998). The *snow* package (Rossini, Tierney, and Li 2003) was employed to divide the generation of candidate solutions across processors in a multiprocessor environment. In contrast to the naive QRE model, there is no guarantee that the beta parameters mapping covariates into utility space for a BQRE will be the same for a PBE, even if the PBE was an acceptable model for the data. Therefore, we do not include a prior distribution when estimating coefficients for this model; it is unclear which values would even be appropriate for such a prior were we to include one.

The Monte Carlo simulation assumes no correlation between the individual type parameters and observable utility covariates. QRE and PBE models that failed to account for individual type parameters would produce biased beta estimates as a result of omitted variable bias if this correlation were nonzero. However, to afford alternative models their best opportunity to match the BQRE's performance, we kept this correlation at zero. Because our comparison takes place under ideal conditions for alternative estimators, our results represent a lower-bound estimate for the deviation in these three estimators' performance.

Table 1 reports the median and mean square error of estimates from the simulation, and demonstrates that the BQRE estimator outperforms the standard QRE and the PBE models on this basis (i.e., it consistently has lower MSE scores). Note that we achieve this result even though the QRE estimation enjoys the benefit of an informative prior. These results demonstrate that the BQRE model outperforms the alternatives with respect to overall fit, but does it also produce more accurate parameter estimates?

¹⁸Note that in the naive QRE model, the underlying data-generating process is similar to the true process (simply lacking the μ), which suggests that one can use the same informative priors that one would use for an BQRE estimate. The same is not true for the PBE (see below).

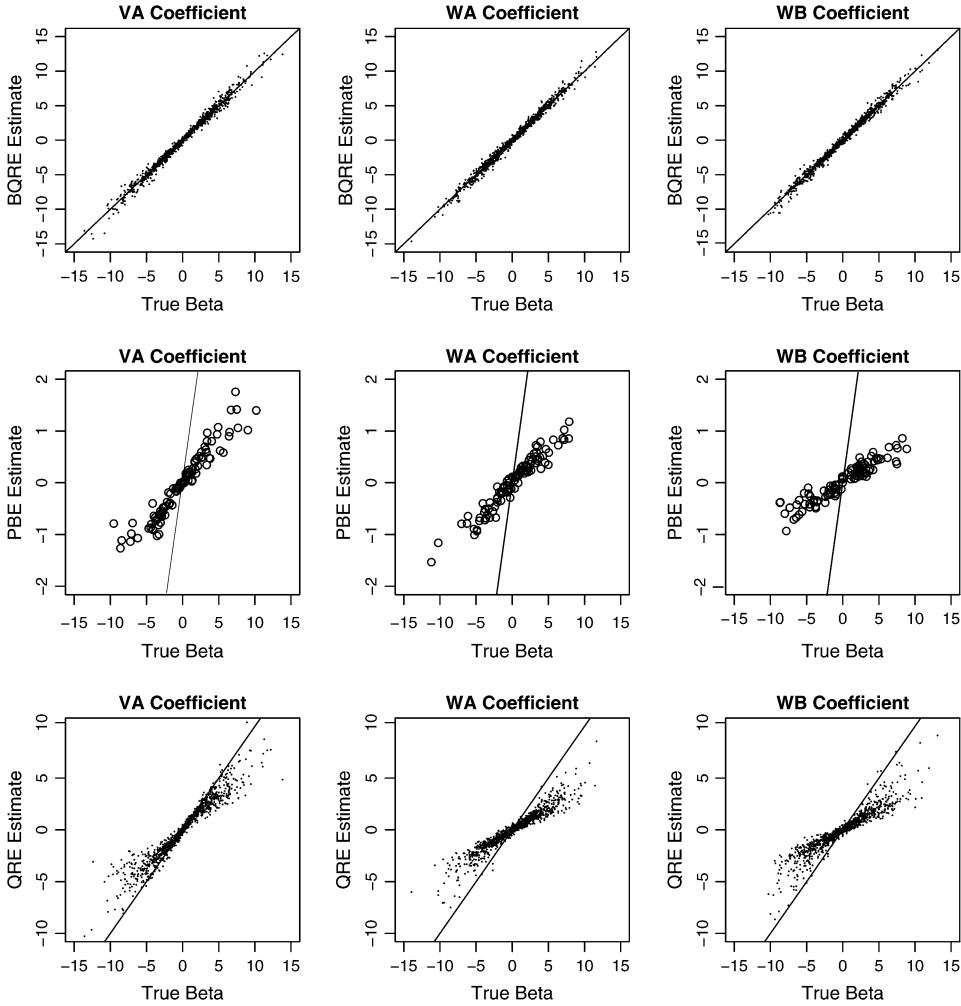


Fig. 2 BQRE, PBE, and QRE estimates versus true beta for Monte Carlo data sets. The line in each panel indicates $y = x$, or a perfect estimate.

Figure 2 provides plots of the estimates versus true beta values for the three covariates for all three models. The figure demonstrates that the BQRE estimates are superior to the QRE and PBE estimates. Note that the BQRE estimates cluster around the true value. On the other hand, both the QRE and PBE estimators consistently deflate¹⁹ coefficient estimates from their true value. This deflation has a systematic impact on the predicted probability of each outcome, causing QRE and PBE to misjudge the magnitude of a covariate's effect on the probability of an outcome, the range over which that effect is realized, and the rate of change of the probability.

Coefficient estimates are of less direct interest than other quantities that can be calculated from the coefficients. For example, Fig. 3 depicts a plot of predicted probability of

¹⁹The PBE estimator deflated coefficients by a median factor of $V_A = 0.1851$, $W_A = 0.1497$, and $W_B = 0.09786$. The QRE estimator deflated coefficients by a median factor of $V_A = 0.7597$, $W_A = 0.4359$, and $W_B = 0.4149$.

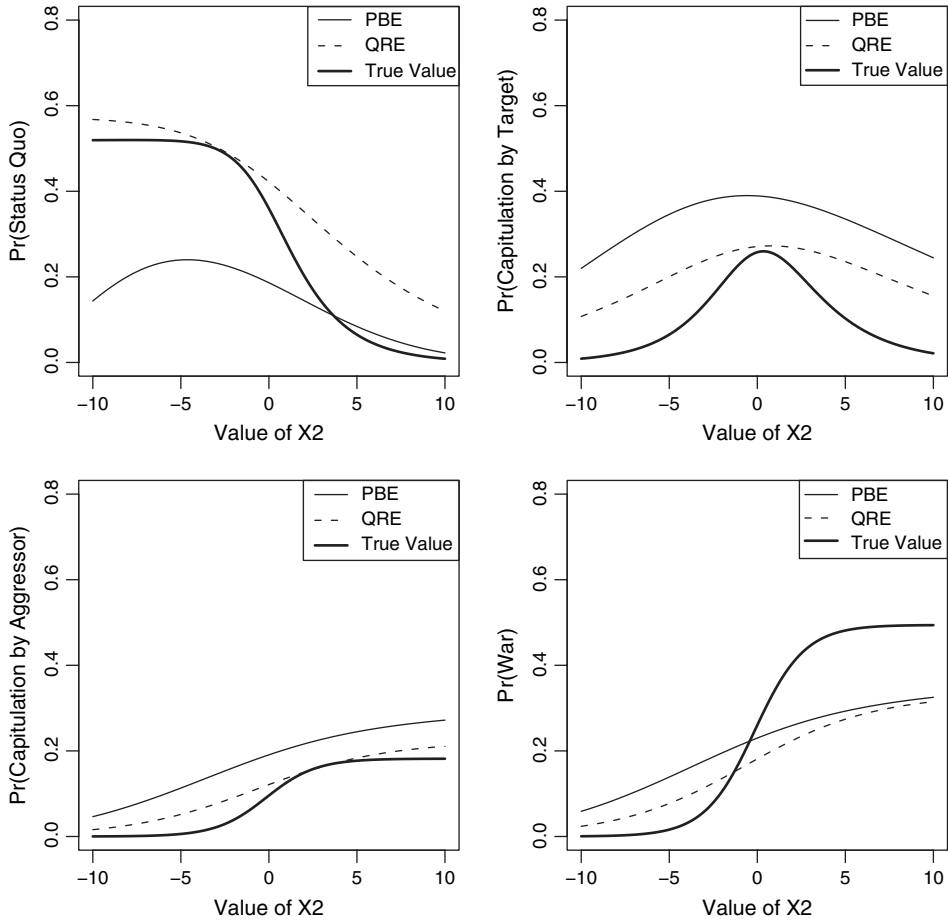


Fig. 3 Predicted probabilities of game outcomes for PBE and QRE estimation versus true BQRE value. $\beta = 1$ for all utility coefficients, $\mu = 0$ for all private information parameters in the true model, and all other covariate values are fixed at 1.

each outcome against changes in x_2 (the covariate of utility for W_A) when $\beta = 1$ for all utility coefficients, $\mu = 0$ for all private information parameters, and all covariates other than x_2 set equal to 1.²⁰ These quantities are of direct interest to researchers, as they allow us to evaluate claims about the impact of changes in utility on behavior given a specific theoretical model. As Fig. 3 demonstrates, PBE underpredicts the probability of the status quo (top left panel of Fig. 3) for values of $x_2 \leq 4$ and predicts the *wrong direction of change* for $x_2 \in [-10, 5]$. Additionally, both estimators underpredict the probability of a war outbreak by about 20 percentage points for $x_2 \geq 0$ and overly gradualize what is a relatively quick change in war probability inside $x_2 \in [-2.5, 2.5]$. In short, the PBE and QRE estimators miss tipping point effects of changes in state A's war payoff and understate

²⁰True predicted probabilities were calculated using BQRE with the (true) $\beta = 1$ for all parameters. Predicted probabilities for PBE and QRE were calculated by setting β equal to the median deflationary factor for the model.

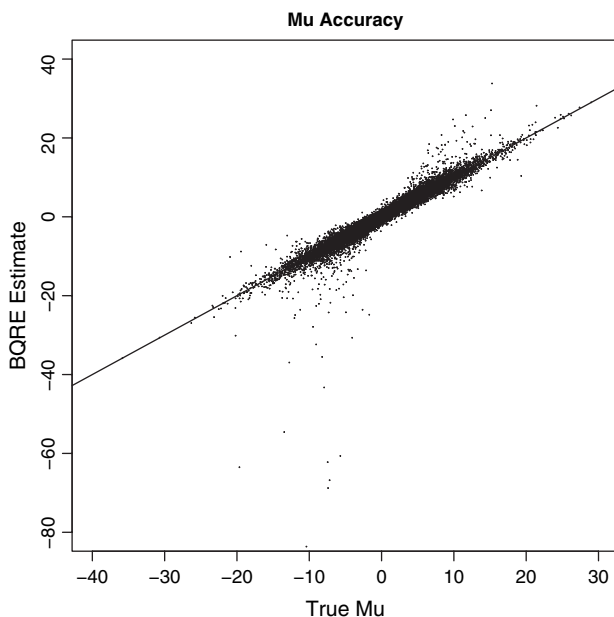


Fig. 4 BQRE estimate of μ versus its true value for 1000 Monte Carlo data sets with 20 countries each (total of 20,000 estimates). The line indicates $y = x$, or a perfect estimate.

the power of changes in this payoff to precipitate a war. A researcher using such a model to predict the path of conflict would often understate the possibility of war, whereas overstating the possibility of capitulation or back down. In the case of PBE, the researcher would also frequently predict a nonstatus quo outcome when none was forthcoming.

By contrast, our BQRE estimator accurately recovers coefficients with no deflationary pattern or bias evident. Even private information parameters are recovered with reasonable accuracy, as shown in Fig. 4. The median square error for the estimation of μ is 3.442×10^{-1} and the mean square error is 2.958, a good performance considering that each μ was estimated with only 38 observations (each simulated state participated in 19 dyads, or 38 directed dyads).

Of course, it is not surprising that the BQRE model can recover coefficients generated from a BQRE process. However, there is an important lesson to glean from these Monte Carlo simulations: alternative models will produce relatively poor fits and will underestimate the coefficients that represent the impact that the respective utilities have at each choice node when the data-generating process resembles BQRE. As a consequence, if researchers have a theory like ours in mind but use either the PBE or QRE estimates to calculate substantive quantities of interest, they run the risk of coming to misleading conclusions. In short, these simulations reinforce the idea that a close fit between theory and empirics is necessary if empirical results are going to be useful tests of a theory and accurate descriptors of the world.

5 Conclusion

We intend this study to advance the growing literature on the development of statistical models that correspond closely to strategic theoretical models. We believe that we have advanced this literature on two fronts. First, existing efforts have had little success

specifying games with updating about player types—a process that is central to Bayesian games—and then estimating the relevant parameters while accounting for the learning/ updating process. Our model offers one way to do so. Second, previous models did not attempt to estimate private information characteristics, despite the fact that few scholars would argue that, *ceteris paribus*, all states have the same baseline propensity to threaten a rival, capitulate/back down if threatened, or respond to a threat with force. The model presented here explicitly estimates the unit-level heterogeneity we anticipate that crisis bargaining data are likely to reveal. We are able to show that a BQRE solution for the game exists, is unique under common conditions, and that it can be used to estimate private information parameters of interest (such as state resolve). We assert that the approach developed here is thus a platform on which others interested in further developing this line of inquiry will want to build.

We conclude with a suggested extension. One useful direction for further research will involve estimation of this model using crisis bargaining data. Although such an exercise will not contribute to the theoretical development of these types of models, it is important for demonstrating the value of this and other models to the wider community of scholarship and policy making. We believe that this is the next important step that we must take to further this research. Of course, although we have selected an international relations context to develop our study, there is nothing inherent in our model that limits its application to this area. The general strategy we have developed can be employed to develop BQRE models of other incomplete information games. We also believe that extensions to these areas, including behavior in laboratory experiments, may be a useful test and application of the model's power to predict behavior.

Appendix: Proof of Lemma 1

We briefly prove below both the existence and uniqueness of BQRE for a general version of the finite extensive form game of incomplete information illustrated in Fig. 1. We use the following notation for the proof. Let (i) $i \in \{1, \dots, n\}$ denote the set of players playing the game; (ii) $t = \{1, 2, 3, \dots, T\}$ denote finite time; (iii) the finite set of outcomes be defined as $(y_{1t}, y_{2t}, \dots, y_{kt}) \in \mathbf{Y}_t$; (iv) $a_i(t) \in A_i$ denote the action of each i at t where A_i is the set of all actions; and (v) $h_i(t)$ define the action history profile of player i with respect to its opponent at each t .

In the game, player i 's private information is given by the constant, nonrandom parameter μ_i , and its opponent's (i.e., j , $i \neq j$) private information is given by μ_j . The behavior strategy for each i is a function which is defined as $b_{it} = b_i(a_i(t) | \mu_i, h_i(t))$, where b_{it} denotes the probability of action a_i at t given μ_i and $h_i(t)$. Note that $b_{it} \in B_{it}$, where B_{it} is the set of all behavioral strategies available to each i . We define each player's prior about its opponent's private information (i.e., its opponent's type) as $q_{it}(\mu_i)$ and define each player's posterior belief as $\pi_i(\mu_i | a_i(t), h_i(t))$. The set of posterior beliefs for each i is thus labeled as $\alpha_{it}(\pi_i)$. Define $V_{it} = V_i(b_{it})$ as each player's conditional expected payoff at each t . Let ε_{it} denote each state's error which is distributed according to a joint distribution with density function $g(\varepsilon)$ and which in the terminology of McKelvey and Palfrey (1998: 14) is assumed to be admissible.

Let \mathbf{N}_t be the set of nodes in the game tree of the finite extensive form game of incomplete information with N_t^m and N_t^0 as the set of terminal and nonterminal nodes, respectively. Each node is labeled as $n_t \in N_t$. Finally, define $\lambda : N_t^0 \rightarrow \mathfrak{R}$ and $\lambda : N_t^m \rightarrow \mathfrak{R}$ as a probability function on N_t^0 and N_t^m and let $b_{it}(\lambda)$ denote each i 's

behavioral strategy at t given λ . We now turn to prove existence of BQRE in the finite extensive game before demonstrating uniqueness.

To prove existence, we need to show that $b_{it}^* \in B_{it}$ exists as a fixed point in B_{it} . Let the index function Φ be defined as $\Phi = V_{it} : B_{it} \rightarrow B_{it}$, where Φ is continuous. By construction, V_{it} in the crisis bargaining game is continuous. Further, since for each i , $b_{it} \subseteq [0, 1]$ and satisfies the definition of continuity, it $\Rightarrow b_{it} \in B_{it}$ is continuous and bounded. Pick any sequence b_{it}^l in B_{it} since $b_{it} \in B_{it}$. Because b_{it} is continuous and bounded, it follows that $\forall t \lim_{l \rightarrow \infty} b_{it}^l \rightarrow b_{it}$ which $\Rightarrow B_{it}$ is closed and convex since $b_{it} \in B_{it}$. This means that Φ is convex valued and upper hemi-continuous, which implies by the Kakutani fixed point theorem that Φ has a fixed point $b_{it}^* \in B_{it}$, as claimed. Given that in the crisis bargaining game, $\Pr(\text{Challenge})$, $\Pr(\text{resist})$, $\Pr(\text{Fight})$ and $(1 - \Pr(\text{Challenge}))$, $(1 - \Pr(\text{resist}))$, and $(1 - \Pr(\text{Fight}))$ are each continuous and bounded in the $[0, 1]$ interval, it follows from the above existence proof of a general version of the finite extensive form incomplete information game that $\Pr(F)$, $\Pr(R)$, and $\Pr(C)$ exist as BQRE in the crisis bargaining game.

To prove uniqueness, first note that each state's prior $q_{it}(\mu_i)$ at each t by construction is fixed and deterministic. Since the prior is deterministic, by Bayes's rule $q_{it}(\mu_i) = \pi_i(\mu_i | a_i(t), h_i(t)) \forall t$. Because $q_{it}(\mu_i) = \pi_i(\mu_i | a_i(t), h_i(t)) \forall t$, it follows that $\pi_i(\mu_i | a_i(t), h_i(t)) = q_{it+1}(\mu_i)$ and from the above claim we obtain $q_{it+1}(\mu_i) = \pi_i(\mu_i | a_i(t+1), h_i(t+1)) \forall t$. Iterating the aforementioned argument yields $\pi_i(\mu_i | a_i(T-1), h_i(T-1)) = \pi_i(\mu_i | a_i(T), h_i(T))$. Hence, when the prior is fixed and deterministic, posterior beliefs $\forall t$ are unique which $\Rightarrow \alpha_{it}(\pi_i) \neq \emptyset$ and that $\alpha_{it}(\pi_i)$ is a singleton. Now consider any $n_t \in N_t^0$ where players hold beliefs about μ_i . Because $\alpha_{it}(\pi_i) \neq \emptyset$ is a singleton, it follows that V_{it} is uniquely determined at any $n_t \in N_t^0$. Since V_{it} is uniquely determined and $\alpha_{it}(\pi_i)$ is a singleton, action probabilities of logit-(A)QRE at any $n_t \in N_t^0$, that is $b_{it}^*(\lambda)$, are also uniquely determined. Repeating this argument for the nodes that follow, we obtain unique $b_{it}^*(\lambda)$ for any $n_t \in N_t^m$. We mentioned earlier that $\lambda : N_t^0 \rightarrow \mathfrak{R}$ and $\lambda : N_t^m \rightarrow \mathfrak{R}$. Since action probabilities of logit-(A)QRE are uniquely determined at any $n_t \in N_t^0$ and $n_t \in N_t^m$ —when the prior is fixed at t and the posterior beliefs are (thus) unique—it follows from Theorem 3 in McKelvey and Palfrey (1998: 16–7) that $b_{it}^*(\lambda)$ is a singleton. Note that in the crisis bargaining game in Fig. 1, a state's prior belief about $f(\mu_{it})$ is the posterior belief $f(\mu_{i(t-1)} | y)$ of the past year's interaction. Iterating this argument shows that in the crisis bargaining game, posterior beliefs $\forall t$ are unique. Hence, from the general proof of uniqueness provided above, we can thus infer that the BQRE solutions in the crisis bargaining game in Fig. 1 are unique. \square

References

- Aldrich, John, and James Alt. 2003. Introduction to the special issue. *Political Analysis* 11:309–15.
- Bas, Muhammet, Curtis Signorino, and Robert Walker. 2007. Statistical backwards induction: A simple method for estimating strategic models. *Political Analysis* doi:10.1093/pan/mpm029.
- Carrubba, Clifford J., Amy Yuen, and Christopher Zorn. 2007. In defense of comparative statics: Specifying empirical tests of models of strategic interaction. *Political Analysis* 15:465–82.
- Fearon, James. 1994a. Domestic political audiences and the escalation of international disputes. *American Political Science Review* 88:577–92.
- Fearon, James. 1994b. Signaling versus the balance of power and interests: An empirical test of a crisis bargaining model. *Journal of Conflict Resolution* 38:236–70.
- Fearon, James D. 1995. Rationalist explanations for war. *International Organization* 49:379–414.

- Fey, Mark, and Kristopher Ramsay. 2006. The common priors assumption: A comment on bargaining and the nature of war. *Journal of Conflict Resolution* 50:607–13.
- Gelpi, Christopher F., and Michael Griesdorf. 2001. Winners or losers? Democracies in international crisis, 1918–94. *American Political Science Review* 95:633–47.
- Granato, Jim, and Frank Scioli. 2004. Puzzles, proverbs and omega matrices: The Scientific and Social Significance and Empirical Implications of Theoretical Models (EITM). *Perspectives on Politics* 2:313–23.
- Greene, William. 2003. *Econometric Analysis*. 5th ed. New York: Prentice Hall.
- Hsiao, Cheng. 2003. *Analysis of panel data*. 2nd ed. New York: Cambridge University Press.
- Keohane, Robert O. 1984. *After hegemony*. Princeton, NJ: Princeton University Press.
- Lewis, Jeffrey B., and Kenneth A. Schultz. 2003. Revealing preferences: Empirical estimation of a crisis bargaining game with incomplete information. *Political Analysis* 11:345–67.
- Martin, Andrew, and Kevin Quinn. 2006. MCMCpack. http://mcmcpack.wustl.edu/wiki/index.php/Main_Page.
- McKelvey, Richard, and Thomas Palfrey. 1995. Quantal response equilibria in normal form games. *Games and Economic Behavior* 10:6–38.
- McKelvey, Richard, and Thomas Palfrey. 1998. Quantal response equilibria in extensive form games. *Experimental Economics* 1(1):9–41.
- Morrow, James D. 1989. Capabilities, uncertainty, and resolve: A limited information model of crisis bargaining. *American Journal of Political Science* 33:941–72.
- Morrow, James D. 1999. The strategic setting of choices: Signaling, commitment, and negotiations in international politics. In *Strategic choice and international relations*, ed. D. A. Lake and R. Powell. 77–114. Princeton, NJ: Princeton University Press.
- Palfrey, Thomas R. 2007. McKelvey and quantal response equilibrium. In *Positive changes in political science: The legacy of richard d. mckelvey's most influential writings*, ed. James Alt, John Aldrich, and Arthur Lupia. Ann Arbor, MI: University of Michigan Press.
- Partell, Peter J., and Glenn Palmer. 1999. Audience costs and interstate crises: An empirical assessment of Fearon's model of dispute outcomes. *International Studies Quarterly* 43:389–405.
- Powell, Robert R. 2004. Bargaining and learning while fighting. *American Journal of Political Science* 48: 344–61.
- Reed, William. 2003. Information, power and war. *American Political Science Review* 97:633–41.
- Reiter, Dan, and Allan C. Stam, III. 1998. Democracy, war initiation, and victory. *American Political Science Review* 92:377–90.
- Reiter, Dan, and Allan C. Stam, III. 2002. *Democracies at war*. Princeton, NJ: Princeton University Press.
- Rossini, Anthony, Luke Tierney, and Na Li. 2003. Simple parallel statistical computing in R. UW Biostatistics Working Paper Series. Working paper 193. <http://www.bepress.com/uwbiostat/paper193>.
- Russett, Bruce M., and John R. Oneal. 1999. The classic liberals were right: Democracy, interdependence, and conflict, 1950–1985. *International Studies Quarterly* 41:267–94.
- Russett, Bruce M., and John R. Oneal. 2001. *Triangulating peace: Democracy, interdependence, and international organizations*. New York: Norton.
- Sartori, Anne E., and Adam Meirowitz. 2006. Secrecy and war: The origins of private information. Unpublished manuscript, Princeton University.
- Schelling, Thomas C. 1966. *Arms and influence*. New Haven, CT: Yale University Press.
- Schultz, Kenneth. 1999. Do democratic institutions constrain or inform? Contrasting two institutional perspectives on democracy and war. *International Organization* 53:233–66.
- Sekhon, Jasjeet S., and Walter R. Mebane, Jr. 1998. Genetic optimization using derivatives. *Political Analysis* 7:187–210.
- Sevcikova, Hana, and Tony Rossini. 2004. Package rlecuyer. <http://cran.r-project.org/src/contrib/Descriptions/rlecuyer.html> (accessed September 14, 2006).
- Signorino, Curtis S. 1999. Strategic interaction and the statistical analysis of international conflict. *American Political Science Review* 99:279–97.
- Signorino, Curtis S. 2003. Structure and uncertainty in discrete choice models. *Political Analysis* 11:316–44.
- Signorino, Curtis S. 2007. On formal theory and statistical methods: A response to Carrubba, Yuen, and Zorn. *Political Analysis* 15:483–501.
- Signorino, Curtis S., and Kuzey Yilmaz. 2003. Strategic misspecification in regression models. *American Journal of Political Science* 47:551–66.
- Slantchev, Branislav. 2003. The power to hurt: Costly conflict with completely informed states. *American Political Science Review* 47:123–35.
- Smith, Alastair. 1999. Testing theories of strategic choice: The example of crisis escalation. *American Journal of Political Science* 43:1254–83.

- Smith, Alastair, and Allan Stam. 2004. Bargaining and the nature of war. *Journal of Conflict Resolution* 48: 783–813.
- Snyder, Glenn, and Paul Deising. 1977. *Conflict among nations*. Princeton, NJ: Princeton University Press.
- Wand, Jonathan. 2006. Comparing models of strategic choice: The role of uncertainty and signaling. *Political Analysis* 14:101–20.