Political Analysis (2003) 11:345–367
DOI: 10.1093/pan/mpg021

Revealing Preferences: Empirical Estimation of a Crisis Bargaining Game with Incomplete Information

Jeffrey B. Lewis and Kenneth A. Schultz
Department of Political Science, University of California, Los Angeles, Los Angeles, CA 90095
e-mail: jblewis@ucla.edu
e-mail: kschultz@polisci.ucla.edu

We develop an empirical estimator directly from an extensive-form crisis bargaining game with incomplete information and discuss its features and limitations. The estimator makes it possible to draw inferences about states' payoffs from observational data on crisis outcomes while remaining faithful to the theorized strategic and informational structure. We compare this estimator to one based on a symmetric information version of the same game, using the quantal response equilibrium proposed in this context by Signorino (1999, *American Political Science Review* 93:279–298). We then address issues of identification that arise in trying to learn about actors' utilities by observing their play of a strategic game. In general, a number of identifying restrictions are needed in order to pin down the distribution of payoffs and the effects of covariates on those payoffs.

1 Introduction

A large and growing literature in international relations (IR) seeks to test models of strategic choice, especially those pertaining to interstate bargaining and conflict (e.g., Bueno de Mesquita and Lalman 1992; Fearon 1994a; Smith 1996, 1999; Sartori 1998; Schultz 1999, 2001; Signorino 1999). Most of the work in this area has focused on testing the comparative-static predictions of game-theoretic models. If a defender becomes more powerful, is immediate deterrence more or less likely to fail (Fearon 1994a)? If a challenging state becomes democratic, is the target of its threat more or less likely to resist (Partell and Palmer 1999; Schultz 1999)? If a state backs down in one crisis, how does this affect the likelihood that it will get its way in a future crisis (Fearon 1994a; Sartori 1998)? The typical testing strategy is to find monotonic relationships between parameters in the model and observable outcomes, to find measurable indicators that serve as proxies for the parameters, and then use standard logit or probit models to see whether these indicators have the predicted effect on the outcome probability.

Authors' note: We gratefully acknowledge the support of the National Science Foundation (NSF-SES-0241647). We thank Michael Chwe, Mark Fey, James Honaker, Al Stam, and Anne Sartori for helpful comments and suggestions. A previous version of this article was presented at the 2001 Midwest Political Science Association meetings in Chicago, IL.

Recently, though, some analysts have begun to move beyond comparative-static analysis to test statistical models that are derived directly from the structure and payoffs of an extensive-form game (especially, Signorino 1999, 2003; see also Smith 1999). The rationale for doing so is appealing. In the first place, deriving the empirical estimator directly from the game permits a test of the full model, rather than some of its isolated implications. Moreover, in strategic settings, the relationship between variables of interest and observable outcomes may be more complicated than standard techniques force us to assume. At a minimum, the functional forms of these relationships may look very different from those imposed by logit or probit regressions. At worst, the implied relationships may be nonmonotonic, making them awkward to capture in off-the-shelf linear models. In addition, strategic interactions create selection effects that can bias inferences unless appropriate steps are taken (e.g., Fearon 1994a; Smith 1996). By embedding the game structure directly into the empirical model, the latter can more accurately capture the rich and complex relationships that emerge in strategic settings.

Although recent articles by Signorino (1999, 2003) and Smith (1999) have broken important ground in this area, the techniques that they develop are not applicable to a class of games that is central to current thinking about international conflict: games of incomplete information. In such games, state leaders have limited information about the military and political attributes that determine their opponents' preferences over the possible outcomes in a crisis. It is now well established in the IR literature that this situation is at least a contributory condition for war and perhaps even a necessary condition (Morrow 1989; Fearon 1995). Incomplete information creates uncertainty over which negotiated settlements are mutually acceptable. Overcoming this uncertainty is problematic because states generally have incentives to engage in strategic misrepresentation that make it difficult to distinguish genuine threats from bluffs. Much of states' behavior in international crises thus revolves around efforts to signal, and exploit, private information through threats, demonstrations of military force, and so forth. As a result, analysts interested in testing theoretical models of crises must find a way to capture games of incomplete information and the signaling dynamics that arise in these interactions.[1]

In this article, we develop an empirical estimator for a canonical crisis bargaining game with incomplete information and discuss some of its features and limitations. The problem we consider is as follows. Start with an extensive-form game and assumptions about the distribution of information among participants. Assume that we then collect observational data reporting the frequency of the realized outcomes in a given population, as well as data describing the characteristics of the states involved in each interaction, international conditions at the time, and so forth. On the basis of these data and assuming that the game structure is correct, what inferences can we make about the distribution of terminal node payoffs? What inferences can we make about the effect of theoretically interesting covariates (e.g., the distribution of power, alliances, and regime type) on those payoffs? What, in short, do our observations about the play of the game reveal about states' preferences in international crises?

The answers that we develop in this article suggest both the promise and pitfalls of this avenue of model testing. On the one hand, we show that it is possible to develop and implement an empirical estimator that is faithful to the nature of strategic interaction in

[1] In this issue, Signorino (2003) introduces an estimator that captures incomplete information in an unconventional way, using a "multiagent representation" in which a given player is represented at each node by different agents, with different payoff disturbances. As Signorino notes, this representation implies that neither signaling nor updating is possible.

coercive bargaining situations. Such an estimator can internalize not only the strategy sets and sequence of moves of a crisis game but also the information asymmetries and signaling dynamics that the IR literature identifies as being central to crisis behavior. On the other hand, we show that there are problems of model identification that complicate this kind of analysis. These problems are not unique to our model; instead, they arise from fundamental features of utility theory and the task of trying to infer utilities from revealed preferences. As a result, researchers interested in rigorous testing of game-theoretic models will have to keep these issues in mind as this literature moves forward.

This article complements Signorino's contribution to this issue in two ways. First, although Signorino lays out a very ambitious agenda for testing strategic choice models, our explicit incorporation of incomplete information and signaling helps to bring that agenda closer to application in a way that will appeal to the IR literature. Second, we formally address issues of model identification that are raised but not fully resolved in Signorino's article. In particular, we derive some conditions for identification and discuss the kinds of restrictions that analysts will have to employ in applying estimators of strategic choice models.

This article proceeds as follows. Section 2 presents the bargaining game that we use as the basis of the analysis and presents the equilibrium solution under incomplete information. For the purposes of comparison, this section also derives the equilibrium by using the quantal response equilibrium (QRE) solution concept employed by Signorino (1999). Section 3 then discusses the issues of estimation and identification that arise in trying to infer payoff distributions from the observed outcomes of the game. We first consider a quasi-experimental world in which the game is played a large number of times by states that are drawn from the same distribution each time—that is, the expected value of each payoff is the same across iterations. We then ask: given the relative frequency of outcomes, can we recover these mean values? After establishing some of the basic issues of estimation and identification that arise in this simple context, we use the estimators on simulated data. These results show that an estimator derived from the incomplete information model differs in fundamental ways from the estimator derived from a symmetric information version of the same game solved using QRE. Section 4 explores what happens when the mean payoffs vary across iterations with some set of independent variables. This section develops some conditions for model identification and gives practical advice for implementing the estimator on real data. Section 5 concludes.

2 The Model

The analysis in this article is based on a simple, canonical crisis bargaining game with incomplete information. Although more complicated games can and have been used to study international conflict, this game is the simplest that allows for the possibility of signaling and hence is the appropriate starting place for this endeavor. In this section, we present the basic assumptions of the model and derive its equilibria under two different assumptions about the information structure. The main presentation solves for the perfect Bayesian equilibrium (PBE) that holds when the crisis participants are assumed to have private information about their opponent's payoffs. For the purposes of comparison, we also solve the game using the method proposed by Signorino (1999): QRE.

Although both models incorporate some form of uncertainty, it will become apparent that the information structures built into these equilibria generate fundamentally different dynamics. The PBE is built on the assumption that actors are incompletely informed about some of their opponent's payoffs, and particularly their opponent's willingness to wage war. Under this kind of uncertainty, actors must try to infer the other side's "type" from its actions, such as whether or not it chooses to make or resist a threat. As a result, a given choice can affect the outcome not simply by moving the action down a certain path of the game tree but also by influencing the beliefs of the other actor. In the crisis game presented here, a challenger state first decides whether or not to present a target state with a demand backed by the threat of war. Not only does this choice determine whether or not there will be a crisis, but it also allows the target to update its prior beliefs about the challenger's resolve to fight. The target's decision whether or not to acquiesce hinges in large part on whether the threat leads the target to believe that the challenger is a type that will fight if the demand is refused. After all, not all threats must be genuine, and the challenger has strategic incentives to misrepresent its true type in order to wrest concessions from the target. In equilibrium, then, a crisis involves a complicated dance of signaling and learning: the challenger knows that its threats will influence the beliefs of the target, and the target updates its beliefs knowing that the challenger has an incentive to bluff and mislead.

In the QRE, by contrast, uncertainty enters in the form of "perceptual disturbances"—unpredictable errors that the actors make as they evaluate their choices at each decision point. A crucial aspect of this information structure is that the errors are equally unpredictable to all actors. In other words, information is symmetric, not asymmetric, as it is in the PBE version. As a result, choices serve only to move the action down the game tree; they do not reveal anything about the actor's type or predict anything about its future behavior. Thus, the process of signaling and updating that is at the core of the PBE model is absent from this version of the game.[2] As we will see, this difference has important consequences for the empirical estimators built upon these two information structures. Despite the superficial similarities between the models, the QRE does not in any way "approximate" the PBE.

2.1 *Sequence of Moves*

The extensive form is depicted in Fig. 1. Two states, A and B, have a dispute over some contested good. We assume, without loss of generality, that the good belongs to B in the status quo. The game begins with a decision by A whether or not to challenge B for the good. A challenge is assumed to involve an explicit threat to use military force in the event that B does not hand over the good. If A chooses not to make a challenge, the status quo (SQ) prevails. If A does make a challenge, then B must decide whether or not to resist the demand. If B does not resist, it concedes the good to A (CD), and the crisis ends peacefully. If B resists, then A must decide whether or not to fight. If A chooses not to fight, then it backs down from its threat (BD), and the good remains in B's possession. Otherwise, A stands firm (SF) and implements its threat. We will sometimes refer to this last outcome as "war," but nothing requires that the action that ensues meet the usual requirements for full-scale war.

2.2 *Payoffs*

We let V_A denote the value that A places on getting the good without a fight, and let S_A denote A's payoff from the status quo. In the event that A makes a challenge and then backs

[2]As noted in note 1, Signorino (2003) introduces a variant on the QRE that incorporates private information in a nonstrategic way that does not permit signaling and updating. Consequently, the contrasts we draw here between our model and the original QRE apply equally to this newer variant.

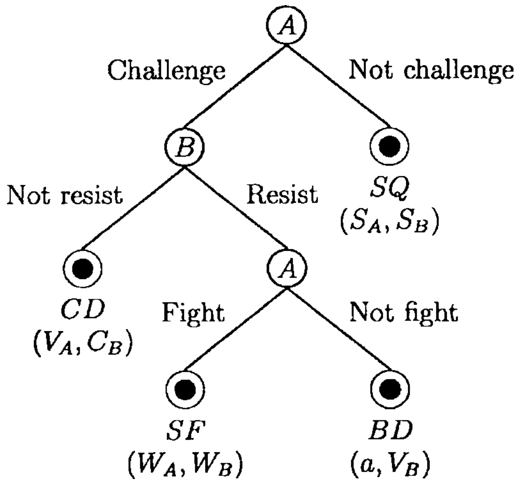

Fig. 1 Simple crisis bargaining game.

down, we allow for the possibility that A incurs some audience cost for having been caught in a bluff (Fearon 1994b; Smith 1998; Guisinger and Smith 2002). We let a denote A's payoff in this event. We do not, however, impose the restriction that a is less than zero—that is, that there is in fact a cost to making a threat and then failing to carry it out. Rather, we will leave the characteristics of this payoff, including its likely sign and magnitude, as an empirical matter.

For state B, there are two ways that it can retain possession of the good without a fight: either A declines to challenge or A makes a challenge and then backs down. For reasons that will be clear shortly, it is useful to differentiate between these two outcomes, at least notationally. Let S_B denote B's payoff in the status quo, and let V_B denote B's payoff in the event that A backs down from its challenge. Let C_B denote B's utility from conceding the good.

Finally, in the event of military action, the states receive some payoff that is a function of the probabilities of different outcomes, the costs of fighting, the value of the stakes, and their risk propensities (e.g., Morrow 1985). For simplicity, we let W_A and W_B denote the expected value of war to A and B, respectively.

2.3 *Information and Beliefs*

The game is played with two-sided incomplete information. In particular, B has private information about its payoff from war, W_B, and A has private information about both its payoff from war, W_A, and its payoff from backing down, a. In assuming that A has private information about both of these payoffs, we depart from standard practice, in which it is generally assumed that the costs of backing down from a threat are common knowledge (e.g., Fearon 1994b; Schultz 1999). In an earlier article (Lewis and Schultz 2001), we estimated a model employing this standard assumption. Introducing uncertainty about the audience costs makes the informational assumptions more realistic and also turns out to simplify estimation.

Formally, we assume that W_A, a, and W_B are generated in the following manner:

$$W_A = \bar{W}_A + \epsilon_A,$$
$$W_B = \bar{W}_B + \epsilon_B,$$
$$a = \bar{a} + \epsilon_a,$$

where each ϵ_i is independently drawn from a normal distribution with mean zero and standard deviation σ. Although the average payoffs $\bar{W}_A$, $\bar{W}_B$, and $\bar{a}$ are common knowledge, the disturbance terms are known only by the appropriate state. Hence, each payoff consists of an observable mean and an unobserved shock. The assumption of normal disturbances is made for the purposes of estimation.

2.4 *Perfect Bayesian Equilibrium*

We solve this game for a PBE in which all the strategies are sequentially rational given the actors' beliefs and the beliefs are calculated from the equilibrium strategies according to Bayes' rule.

At its final node, A fights if and only if

$$W_A > a. \tag{1}$$

Define p_F as the probability that (1) holds, given that A has made a challenge. Given its posterior belief p_F, B resists if

$$p_F W_B + (1 - p_F)V_B > C_B,$$

or

$$W_B > \frac{C_B - (1 - p_F)V_B}{p_F}. \tag{2}$$

Let p_R denote the probability that (2) holds. From the distributional assumptions, we know that the expected probability of resistance is

$$p_R = \Phi\left[\frac{p_F \bar{W}_B + (1 - p_F)V_B - C_B}{p_F \sigma}\right], \tag{3}$$

where Φ is the standard normal cumulative distribution function. Given p_R, the expected value of making a challenge for an A of type (a, W_A) is

$$EU_A(\text{CH}) = p_R \max(a, W_A) + (1 - p_R)V_A. \tag{4}$$

Thus, A prefers a challenge to the status quo if

$$\max(a, W_A) > \frac{S_A - (1 - p_R)V_A}{p_R} \equiv c^*. \tag{5}$$

The probability of a challenge, which we denote by p_C, is equal to the probability that (5) holds, which is equivalent to one minus the probability that both a and W_A are less than c^*.

Thus,

$$p_C = 1 - \Phi\left(\frac{c^* - \overline{W}_A}{\sigma}\right)\Phi\left(\frac{c^* - \overline{a}}{\sigma}\right). \tag{6}$$

We can now calculate p_F, the conditional probability that A will fight, given that it has made a challenge. This is

$$\begin{aligned} p_F &= \mathbb{P}[W_A > a \mid \max(a, W_A) > c^*] \\ &= \mathbb{P}[W_A > a \cap \max(a, W_A) > c^*]/p_C \\ &= \mathbb{P}[W_A - a > 0 \cap W_A > c^*]/p_C. \end{aligned} \tag{7}$$

Let $\Delta = W_A - a$. Given our assumptions, Δ is distributed normally with mean $\overline{W}_A - \overline{a}$ and variance $2\sigma^2$. Moreover, the joint distribution of Δ and W_A is bivariate normal with correlation $1/\sqrt{2}$. We can thus rewrite (7) as

$$p_F = \Phi_2\left(\frac{\overline{W}_A - \overline{a}}{\sigma\sqrt{2}}, \frac{\overline{W}_A - c^*}{\sigma}, \frac{1}{\sqrt{2}}\right)\bigg/ p_C. \tag{8}$$

The system of equations in (3), (6), and (8) characterizes the equilibrium choice probabilities at each node. The equations are not in closed form; nevertheless, we can confirm that there always exists a unique solution to this system.

The probabilities associated with each outcome follow immediately from the choice probabilities derived above:

$$\begin{aligned} \mathbb{P}(\text{SQ}) &= 1 - p_C, \\ \mathbb{P}(\text{CD}) &= p_C(1 - p_R), \\ \mathbb{P}(\text{BD}) &= p_C\, p_R(1 - p_F), \end{aligned}$$

and

$$\mathbb{P}(\text{SF}) = p_C\, p_R\, p_F.$$

One point to notice at the outset is that B's payoff from the status quo, S_B, appears nowhere in the equilibrium choice probabilities. Because this payoff cannot be reached by a choice of B's, nothing in the play of the game depends on its value. As a result, we will not be able to infer anything about this payoff unless we make some additional assumptions, such as assuming $S_B = V_B$—that is, B's payoff from possessing the good is the same whether or not A made a challenge. This assumption is commonly made in crisis bargaining models. We have now shown that it is an assumption that cannot be tested empirically given the game structure.

2.5 *Agent Error/Quantal Response Equilibrium*

For the purposes of comparison, we also solve the game in the manner proposed by Signorino (1999). In this version, the mean payoffs represent the "true" payoffs, but states are assumed to make perceptual errors as they evaluate their choices. At each node, the actor makes what amounts to a probabilistic choice in which the action with the higher expected utility is more likely, but not certain, to be chosen (see McKelvey and Palfrey 1998 or Signorino 1999). The strategy in solving for such a QRE is as follows. Assume that the actor faces a choice between two actions, a_1 and a_2, and that the expected utility from a_1 is U_1 and the expected utility from a_2 is U_2. The actor is assumed to make some error in evaluating this choice, an error that is captured by adding a disturbance term to the expected value of each branch. Thus, rather than comparing U_1 and U_2, the actor compares

$$U_1^* = U_1 + \epsilon_1$$

and

$$U_2^* = U_2 + \epsilon_2,$$

where ϵ_1 and ϵ_2 are independent, normally distributed random variables with mean zero and variance σ^2. The probability that the actor chooses a_1 is then equal to

$$\mathbb{P}(U_1^* > U_2^*)$$

or

$$\Phi\left(\frac{U_1 - U_2}{\sigma\sqrt{2}}\right).$$

We can readily apply this solution concept to the present model. At its final node, state A faces a choice between $\bar{W}_A$ and $\bar{a}$. The probability of fighting is

$$p_F = \Phi\left(\frac{\bar{W}_A - \bar{a}}{\sigma\sqrt{2}}\right). \tag{9}$$

State B then faces a choice between C_B and a lottery worth

$$p_F \bar{W}_B + (1 - p_F)V_B.$$

Again perturbing each value by a normally distributed error, state B will resist with probability

$$p_R = \Phi\left(\frac{p_F \bar{W}_B + (1 - p_F)V_B - C_B}{\sigma\sqrt{2}}\right). \tag{10}$$

Finally, state A faces a choice between S_A and a lottery worth

$$p_R[p_F \bar{W}_A + (1 - p_F)\bar{a}] + (1 - p_R)V_A.$$

The probability of making the challenge is then

$$p_C = \Phi\left\{\frac{p_R[p_F \bar{W}_A + (1 - p_F)\bar{a}] + (1 - p_R)V_A - S_A}{\sigma\sqrt{2}}\right\}. \tag{11}$$

Notice that each probability can be easily rendered as a function of the payoffs by substituting (9) into (10) and then (10) into (11). As before, the probability of each observable outcome can be determined by multiplying the various choice probabilities together.

2.6 *Comparing the Perfect Bayesian Equilibrium and Quantal Response Equilibrium Models*

The observability and interpretation of the random shocks is a main point of difference between the PBE and QRE models. In the former, the shocks associated with W_A, W_B, and a are private information of the respective players. In the QRE model, the shocks represent a perceptual disturbance of the true, mean payoffs. As noted earlier, the key difference between the models is that the PBE predicts signaling and updating, whereas the QRE does not.

In both models, state B's decision turns on its belief, p_F, that A will fight at its final node. In the PBE model, p_F is a conditional probability: the probability that $W_A > a$ given that A made a challenge. It represents a posterior belief about the distribution of A's type conditional on A being the type that makes a challenge in equilibrium. In the QRE model, no such updating takes place because A's choices are driven by perceptual errors that are redrawn at each decision point rather than by its fixed type. When deciding whether or not to make a challenge, A does not know whether or not it will fight at its final node; it can only assign probabilities to its future actions based on the difference between the (commonly known) mean values. As a result, A's decision to make a challenge reveals nothing about the relative values of W_A and a that B did not already know. Therein lies the main difference between the two information structures. In the PBE model, information is asymmetric: A knows from the outset whether or not it will fight, and B must try to infer p_F from A's actions. In the QRE model, information is symmetric: A and B have the same estimate of p_F throughout.[3]

The implications of these differences can be seen in the comparative-static predictions displayed in Fig. 2. In this figure, three payoff means, $\overline{W}_A$, $\overline{W}_B$, and $\overline{a}$, are varied from -3 to 1, assuming a baseline payoff vector of $(\overline{W}_A, \overline{W}_B, \overline{a}) = (-1, -1, -0.5)$. For the purposes of this exercise, we have normalized the value of the good and the variance of the disturbance terms to one, a normalization that is discussed further below. The corresponding changes in the probability of each outcome are shown, with the solid line depicting the prediction of the PBE model and the dotted line depicting the prediction of the QRE.

Without going through every panel in great detail, it is clear from this figure that there are a number of differences in the two models' comparative-static predictions. Even when the general direction of the relationship is the same, the functional forms can be quite different. The most striking difference is in the relationship between $\overline{W}_A$ and the probability of SF. While the QRE predicts a monotonic increase in the probability of war as $\overline{W}_A$ increases, the PBE predicts a nonmonotonic relationship—with the probability of war first increasing and then decreasing.[4] The maximum probability of war occurs when the mean value of W_A is near the mean value of a, in which case there is maximum uncertainty about whether or not A will fight. In this range, the potential for B to "mistakenly" resist genuine threats is the highest. For values of $\overline{W}_A$ well below this range, the probability of war is low because it is unlikely that A will want to fight. For values of $\overline{W}_A$ well above this range, the probability of war is low because although A almost certainly wants to fight, B knows this and rarely resists.

As we will see in the next section, differences in the predicted relationships between payoffs and outcome probabilities will lead estimators based on these models to generate different inferences from the same data.

[3]Indeed, Signorino's original purpose in employing the QRE (Signorino 1999) was to generate probabilistic outcomes to the complete information game of Bueno de Mesquita and Lalman (1992).

[4]See Schultz (1999) for a similar result and discussion.

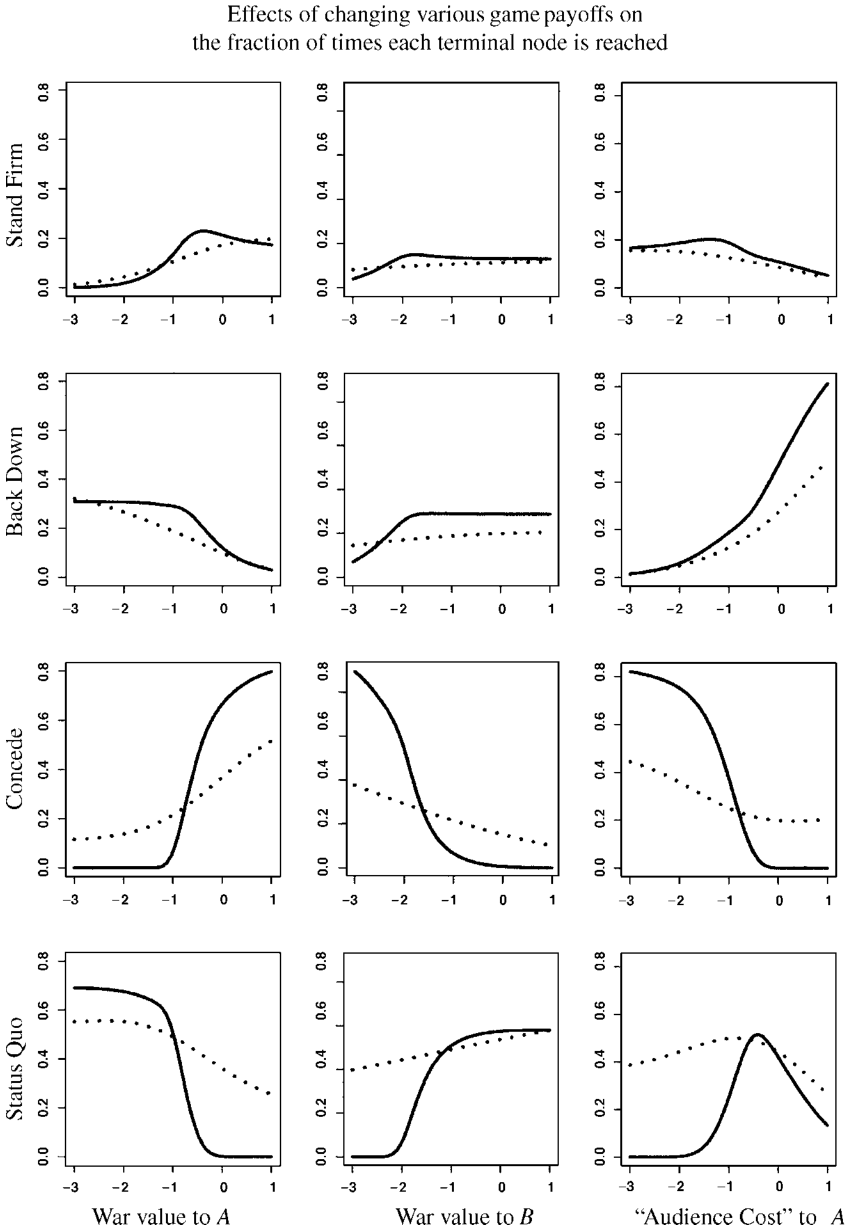

Fig. 2 Each panel shows how the fraction of outcomes of the crisis bargaining game that end at each particular node varies as a function of the given payoff. Each payoff is varied from −3 to 1. The baseline payoff vector $(\overline{W}_A, \overline{W}_B, \overline{a}) = (-1, -1, -1/2)$. The "audience cost" for A is the payoff A receives for backing down; larger values of this quantity represent lower audience costs. ——— (PBE); ·········· (QRE).

3 Estimation and Identification

The equilibrium solutions determine the probability of each outcome as a function of the payoff distributions. The estimation challenge we face is to invert these functions—that is, to make inferences about the payoff distributions given the empirical frequency of the outcomes. In particular, assume that we observe N plays of the game and have data on the outcome of each play as well as on covariates that are thought to influence the payoffs. The objective of the estimation is to recover the effects of the covariates from these data.

Formally, let $\Psi = \{\bar{V}_A, \bar{V}_B, \bar{W}_A, \bar{W}_B, \bar{a}, \bar{S}_A, \bar{S}_B, \bar{C}_B\}$ denote the vector of mean payoffs. Note that because only W_A, W_B, and a include stochastic disturbances, the distributions of the other payoffs have zero variance around their means—that is, $V_A = \bar{V}_A$, and so on. Ideally, we would like to write each element of Ψ as a linear combination of covariates, so that for each observed play of the game i,

$$\psi_i = \boldsymbol{\beta}'_\psi \mathbf{X}_i,$$

where $\mathbf{X}_i$ represents a matrix of covariates and $\boldsymbol{\beta}_\psi$ measures the effect of each covariate on a given average payoff. It is each $\boldsymbol{\beta}_\psi$ that we ultimately seek to estimate.

3.1 *A Stylized Data Generating Process*

To begin demonstrating the estimation issues that arise, we first focus on a simpler problem in which each mean payoff is constant in every iteration of the game. Imagine a highly idealized world in which identical states interact repeatedly under conditions that are observationally identical. The mean payoffs, Ψ, are the same in all N observations, and the only differences from one case to the next are the realizations of the stochastic disturbance terms. Given this data generating process, the object of the estimation is simply to infer the average payoffs from the relative frequency with which each outcome is observed.

It immediately becomes apparent that there are identification problems arising from the limited degrees of freedom in the data. As we will see, these issues are mitigated, but do not entirely disappear, when covariates are introduced. The game has four possible outcomes, and so the data consist of four observed frequencies: how often the game ended at SQ, how often at CD, how often at BD, and how often at SF. Because these frequencies must sum to one, the probability of any one outcome is completely determined by the probabilities of the other three. Hence, three parameters are sufficient to fully describe the data. On the other hand, there are nine unknown parameters in the model: the eight payoff means and the variance of the disturbance terms, σ^2. Clearly, some restrictions are needed before we can even begin.

Fortunately, because the payoffs are assumed to be von Neumann–Morgenstern utilities, which are only defined up to a linear transformation, several restrictions can be made without loss of generality. The natural restrictions involve fixing, for each player, one payoff to zero and another to one. In doing so, we effectively set the scale on which the payoffs are measured. In this context, it makes sense to fix the value of the good to one and to assume that outcomes in which the player does not possess the good are worth zero. Formally, let $\bar{V}_A = \bar{V}_B = 1$ and $\bar{S}_A = \bar{C}_B = 0$. This normalization effectively determines the utility scale for each state: one unit of utility is equal to the difference between getting the good without a fight and not possessing the good.[5] Three payoff means then remain to be estimated: $\bar{W}_A$, $\bar{W}_B$, $\bar{a}$.

[5] Although the utilities are fixed for each of the two players, these scales are not comparable between players.

After normalizing the payoffs in this way, there still remains a fourth free parameter, σ, and only three degrees of freedom in the data. To allow estimation of the payoffs, we fix $\sigma = 1$. This normalization is not innocuous as it is, for example, in probit regression models. Because we have already fixed the units of the utility scale, restricting the value of σ is a substantive restriction. In setting $\sigma = 1$, we impose the assumption that the standard deviation of the payoff shocks is equal (in utility terms) to the difference between having and not having the good. Whether or not this assumption is plausible—it most likely is not—it is not one that can be tested in the data. For any $\sigma > 0$ that we choose, we can find a set of payoffs that predict the observed distribution of outcomes. As we will see, this problem can be mitigated under some circumstances by the addition of covariates, which will allow us to remove the normalization on $\bar{V}_A$ and $\bar{V}_B$.

The general likelihood function for the data is

$$L = \prod_{P \in \{\text{SF,CD,BD,SQ}\}} \mathbb{P}(P \mid \bar{W}_A, \bar{W}_B, \bar{a})^{n_P},$$

where n_P is the number of times each outcome P occurred. In this case the model has three parameters, and the data have three degrees of freedom. Thus, the model is saturated, and the estimated mean payoffs $(\bar{W}_A^*, \bar{W}_B^*, \bar{a}^*)$ will be such that

$$\frac{n_P}{N} = \mathbb{P}(P \mid \bar{W}_A^*, \bar{W}_B^*, \bar{a}^*)$$

for all $P \in \{\text{SF, CD, BD, SQ}\}$.

3.2 *Comparing the Perfect Bayesian Equilibrium and Quantal Response Equilibrium Estimators*

Although the data generating process we have assumed is vastly oversimplified, the estimates obtained in this manner are sufficiently rich to show that a model which assumes an environment of asymmetric information leads to inferences about payoffs which differ in important ways from inferences produced by other models, such as Signorino's QRE model (Signorino 1999), that assume symmetric information.[6]

Figure 3 compares the estimates of $\bar{W}_A$, $\bar{W}_B$, and $\bar{a}$ for different outcome distributions. To make this figure, we assume the following baseline distribution: $\mathbb{P}(\text{SF}) = 0.10$, $\mathbb{P}(\text{BD}) = 0.25$, $\mathbb{P}(\text{CD}) = 0.15$, and $\mathbb{P}(\text{SQ}) = 0.50$. In each column, we vary the frequency of the specified outcome from 0.1 to 0.6, while holding the relative probabilities of the other three outcomes constant. Thus, for example, when we vary the fraction of cases that end at CD, the ratio of SF:BD:SQ remains fixed at 2:5:10. The graphs show how the PBE and QRE estimates of each parameter change in response.

As is apparent from the figure, there are some cases in which a change in the distribution of outcomes causes similar reactions by the two different estimators. For example, if the SF outcome becomes relatively more frequent, both estimators conclude that war has become more attractive to both states. At the same time, there are some very dramatic differences.

[6]We have performed a similar exercise comparing the PBE model to the private information variant of the QRE introduced in this issue (Signorino 2003). The main results are similar to those presented here.

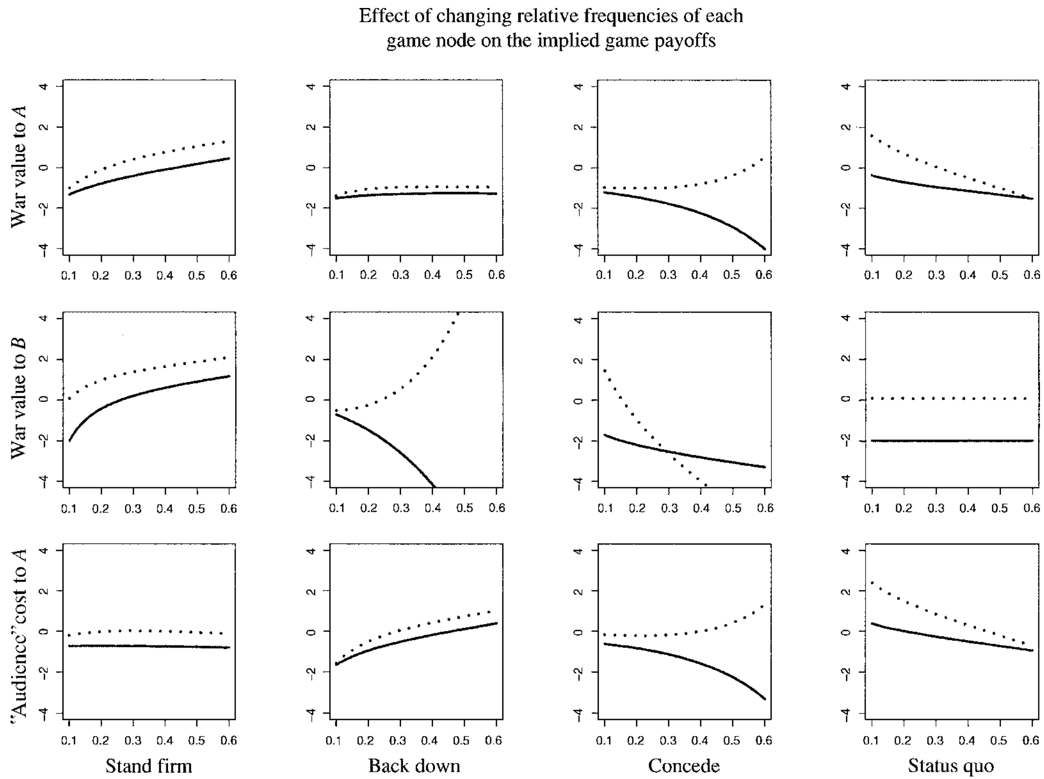

Fig. 3 Each panel shows an implied game payoff as a function of the fraction of outcomes that end at a particular node. The baseline distribution of outcomes $(\mathbb{P}(\text{SF}), \mathbb{P}(\text{BD}), \mathbb{P}(\text{CD}), \mathbb{P}(\text{SQ})) = (0.10, 0.25, 0.15, 0.50)$. The probability of each outcome was varied from 0.1 to 0.6 holding the relative probabilities of the remaining outcomes fixed. For example, when the fraction of cases that end at CD is varied, the ratio of SQ:SF is held at 5:1. The "audience cost" for A is the payoff A receives for backing down; larger values of this quantity represent lower audience costs. ——— (PBE); ·········· (QRE).

The most instructive of these differences is in how the estimators react to an increase in the rate of concession by B. When CD becomes more frequent, both estimators assume that war must have become less attractive for B; however, they disagree about what happened to A's payoffs. The PBE estimator concludes that the costs of backing down must have increased (i.e., $\bar{a}$ decreased), thereby making A's threats more credible. Such an inference follows naturally from the logic of costly signaling: the more costly it is for the coercer to back down from a threat, the more likely it is that threat will be carried out (especially, Fearon 1994b; Schultz 1999).[7] The QRE estimator, on the other hand, concludes that $\bar{W}_A$ and $\bar{a}$ increased in tandem. Such changes increase the probability that A will make a challenge in the first place—thereby giving B many more opportunities to concede—without changing the relative probabilities of the SF and BD outcomes.

These contrasting reactions underscore the differences in the behavioral and informational assumptions underlying these two equilibrium concepts. In a QRE, A's initial move has no influence on B's choice, other than to determine whether or not its decision node is even reached. To increase the frequency of concession, it is sufficient to increase the rate of challenges. In a PBE, by contrast, A's strategy at its initial node not only determines whether or not B gets a move but also influences the target's posterior beliefs and hence its willingness to resist. To increase the frequency of concessions, it is necessary to increase the rate of *credible* challenges—that is, challenges that lead B to believe that A is a type that will fight rather than back down. This can be achieved by making types that want to back down from a challenge rarer in the population, a change that influences both the target's prior beliefs and its posterior beliefs conditional on a challenge. Thus, in response to an increase in CD outcomes, the QRE estimator tries to make challenges more attractive, whereas the PBE estimator tries to make challenges more costly, and hence more credible.

Even though this result was derived using a highly stylized data generating process, we can anticipate that the differences between the solution concepts will carry over when covariates are introduced. Consider a simple experiment with a control group and a treatment group. Say, for example, that the control group is made up of cases in which state A is nondemocratic and the treatment group consists of cases in which state A is democratic. Such an experiment should permit us to estimate how the regime type of the challenger affects the payoffs. Suppose we find that the relative frequency of CD is higher in the treatment group than in the control group, but that the relative frequency of all the other outcomes remains the same.

From this, what could we infer about the effect of democracy in A on the payoffs? Clearly, our inference would depend upon the solution concept employed. If we use the QRE estimator, the increase in the relative frequency of CD will be attributed to an increase in both $\bar{W}_A$ and $\bar{a}$. We would then infer that democracy in the challenger makes both war and backing down more attractive. The natural conclusion would be that democracies generate relatively more concession outcomes because democracies make more challenges. On the other hand, if we use the PBE estimator, the increase in CD will be attributed to a decrease in $\bar{a}$. We would then infer that democracy in the challenger increases the penalty for backing down (consistent with Fearon 1994b; Smith 1998; Guisinger and Smith 2002). This inference would lead to the conclusion that democracies generate more concession outcomes because they are better able to commit themselves publicly to carrying out their threats.

[7] The slight decrease in the estimate of $\bar{W}_A$ is needed to ensure that, while $\bar{a}$ is decreasing, the relative probabilities of SF and BD do not change.

This exercise demonstrates that the information structure embedded in the estimators has important implications for the inferences one draws about payoffs based on observations of game outcomes. It is not enough for the estimation to internalize the strategy sets and sequence of moves—the information structure must also be taken into account. Using an estimator that is premised on symmetric information to estimate, for example, the effect of democracy on audience costs could yield a result that is the reverse of that which a model that incorporates incomplete information and signaling—the very features that make audience costs an interesting object of study—would reveal.

4 Identification and Estimation with Covariates

In the previous section we described the limited way in which the average payoffs of the crisis bargaining game could be empirically pinned down from repeated observations in the case where the average payoffs were fixed across observations. The problem in that case was a lack of "degrees of freedom" akin to having more independent variables than observations in a simple regression setting. Because there were (potentially) eight average payoffs and only four observationally distinguishable outcomes, there were simply more parameters in the model than could be estimated from the information in the data. A simple way around this problem would be to assume that the average payoffs vary across observations as a function of a vector of covariates, $\mathbf{X}$. In particular, let

$$\psi_i = \boldsymbol{\beta}'_{\psi}\mathbf{X}_i$$

for observations $i = 1, 2, \ldots, N$, where $\boldsymbol{\beta}_{\psi}$ is an $M \times 1$ vector of coefficients associated with payoff ψ, $\mathbf{X}_i$ is an $M \times 1$ vector of covariates, and $\psi \in \{\bar{S}_A, \bar{S}_B, \ldots, \bar{W}_A, \bar{W}_B\}$. The data are represented by an $N \times M + 1$ matrix with rows $[Y_i \; \mathbf{X}_i]_i$ for $i = 1, 2, \ldots, N$. The number of distinguishable outcomes in the data (rows in the data matrix) is now much larger. If any of the variables that comprise $\mathbf{X}$ are measured continuously, the number of distinguishable patterns would be expected to be equal to the number of observations. Thus, with the inclusion of covariates, the sorts of identification problems raised in the previous section could be overcome as N grows larger than the number of parameters in the model.

There are, however, additional limitations to identification that are distinct from this degrees of freedom problem. These limitations result from the nature of revealed preference. First, information about the relative utility that an actor receives from any pair of outcomes can only be inferred if the actor is ever in a position to take a decision that depends in some way on a comparison of the utilities associated with those two outcomes. Thus, for state B in our model, nothing can ever be learned about its utility for the status quo ($\bar{S}_B$) because the option to remain at the status quo is solely at the discretion of state A. Indeed, inspection of the QRE and PBE equilibria to the game presented in Section 2 reveals that $\bar{S}_B$ never appears in any of the choice probabilities. Thus, in these models, $\boldsymbol{\beta}_{S_B}$ is not identified.

Further inspection of the PBE and QRE equilibria reveals that all the choice probabilities—and therefore all the outcome probabilities—are determined by functions of probability statements of the following general form:

$$\mathbb{P}\left(\sum_{\psi\in\Psi} w_{\psi}\psi > \sum_{\psi\in\Psi} w'_{\psi}\psi\right), \tag{12}$$

where $w_{\psi} \geq 0$ and $w'_{\psi} \geq 0$ for all $\psi \in \Psi$ and $\sum_{\psi\in\Psi} w_{\psi} = \sum_{\psi\in\Psi} w'_{\psi} = 1$. That is, each probability is determined by a comparison of weighted sums of terminal node payoffs.

The weights are often endogenously determined by the choice probabilities and in any case many of the weights are zero. For example, in the QRE model the probability that state A chooses to fight at its second move is simply $\mathbb{P}(\overline{W}_A > \overline{a})$. What is then clear is that the PBE and QRE models are invariant to the addition of a constant, c, to each element of Ψ. That is,

$$\mathbb{P}\left[\sum_{\psi\in\Psi} w_\psi(\psi + c) > \sum_{\psi\in\Psi} w'_\psi(\psi + c)\right] = \mathbb{P}\left[\left(\sum_{\psi\in\Psi} w_\psi \psi\right) + c > \left(\sum_{\psi\in\Psi} w'_\psi \psi\right) + c\right]$$
$$= \mathbb{P}\left(\sum_{\psi\in\Psi} w_\psi \psi > \sum_{\psi\in\Psi} w'_\psi \psi\right).$$

In fact, this invariance will hold if constants are added only to certain subsets of Ψ. In particular, because no choice is determined by a comparison of expected utility across players, we can add a constant to all of the payoffs of only state A or state B without changing the probability of any outcome. In more general models, we can create smaller subsets of Ψ to which a constant can be added without changing the probability distribution over the outcomes. For example, in a general class of games in which choice probabilities can be written as in (12), subsets of payoffs, Ψ^l, for each player $l = 1, 2, \ldots, L$ of the game can be partitioned into sets of payoffs that are associated with terminal nodes that emanate from a particular initial choice (information set) for player l. Because outcomes that emanate from different initial choices for a given player are never compared (have $w_\psi > 0$ or $w'_\psi > 0$ in the same probability statement), a constant can be added to each payoff within any such set of payoffs without affecting any of the choice probabilities. In the context of the simple crisis bargaining model, this would include only the singleton $\{S_B\}$ to which, for the reasons described above, we can add any constant c without changing any outcome probabilities.

In the appendix, we define a general class of *translation invariant* multinomial models of strategic choice that includes the models presented above. We show that in such models, vectors of parameters $\boldsymbol{\beta}_\psi$ cannot be freely estimated for each payoff of any subset of Ψ to which we can add a constant without changing any of the outcome probabilities. In the simple crisis bargaining model, this implies that the parameters of an index function for each payoff cannot be freely estimated for each of A's payoffs simultaneously. Similarly for B, the parameters of an index function cannot be freely estimated for each of $\overline{V}_B$, $\overline{W}_B$, and $\overline{C}_B$ simultaneously, and in no event can an index function for $\overline{S}_B$ be estimated.

In essence, this is the same as the identification problem that arises in multinomial logit or probit models. In those models, in which there is only one "player," the utilities for each choice cannot be estimated. Rather, the utilities for each choice can only be defined relative to the others. Typically, the vector of parameters for a particular payoff is normalized to **0** and all of the other effects are interpreted as relative to that payoff. In the crisis bargaining game, the same approach can be taken. For example, we could normalize payoff parameters associated with the status quo payoff for A to be **0**. Similarly, we could normalize the parameters associated with B's conceding the good to be **0**. We would then be measuring the effects of covariates relative to their effect on the value of the status quo in the case of A and relative to their effect on concession in the case of player B. These two restrictions identify all of the parameters of the model except those associated with the status quo for B (which are fundamentally unidentifiable).

Alternatively, identification can proceed on a covariate-by-covariate basis. Each covariate (including the constant) must be omitted from the equation describing at least one of the payoffs for each player. Thus, if we have theoretical reasons to suppose that a particular covariate has no effect on one or more payoffs, we can estimate its effect on the others. Identification could also be created through other sorts of constraints, such as enforcing the equality of certain parameters across some number of payoffs. Finally, if the index functions describing each model payoff are nonlinear, for example if we write

$$\psi = \exp(\boldsymbol{\beta}'_{\psi}\mathbf{X}),$$

then all the βs may be identified. In this case, identification is achieved through differences in the effect that each covariate has on each payoff that arise not only from differences in the parameters but also from variation in the values of the other parameters and covariates across payoffs.

It should also be noted that degrees of freedom problems can obtain if the covariates contain insufficient variation across observations. For example, the addition of two binary covariates to the simple data generating process described in Section 3 would increase the number of degrees of freedom in the data to $2 \times 2 \times 3 = 12$; however, the number of parameters to estimate would be 15 (a constant plus two slopes for each of five average payoffs, the other three payoffs being set to 0). In this case, additional parameter restrictions must be made in order to achieve an identified model.

Thus, the inclusion of covariates does not entirely resolve the problems of identification in multinomial models of strategic choice. Limitations to identification arise not only from a lack of degrees of freedom (information) in the data, but also from the basic nature of utility theory and revealed preference. In a famous quip, Baumol (1958) reminds us that von Neumann–Morgenstern utilities are the "cardinal utilities that are ordinal." The same can be said of random utility models. Ultimately, only relative utilities can be inferred. Moreover, relative utilities can only be inferred across sets of payoffs corresponding to outcomes among which a given player is in a position to choose. This leads to some quite innocuous and self-evident restrictions, such as the inability to place the utilities of all players on the same scale. However, it also implies some less innocuous limitations: for example, the inability to recover any information about the value of the status quo for player *B* in our game. This is a quantity that we may have great interest in but is not something that can be learned from observing the crisis bargaining outcomes even if we have covariates that we believe affect no other payoff. Similarly, we may want to know if democratic institutions affect both the costs of pursuing war and the costs of backing down from a challenge; however, without loss of generality (for example, assuming the effect of democracy on some other payoff is substantively zero), we are only able to estimate the relative magnitude of the effect that democracy has on each cost.[8]

On the other hand, if one is only interested in the effects that covariates have on relative frequencies of various game outcomes, the inability to identify which payoffs the effects work through is not a major concern. In this case, the problems of identification described above are of no consequence. Thus, if we are only interested in whether democracies are less likely to fight (unconditional on a challenge), we need not be concerned with pinning down exactly which payoffs democracy affects.

[8]Kooreman (1994) notes a similar identification result in the context of a 2×2 game.

5 Conclusion

In this article, we have shown that it is possible to develop an empirical estimator directly from a crisis bargaining game with incomplete and asymmetric information. This estimator directly internalizes both the structure of the game and the choice and outcome probabilities of its (unique) PBE. Moreover, we have shown that such an estimator generates different inferences from an estimator derived from a QRE of the same extensive-form game with symmetric information. Although both models incorporate uncertainty, it is clear that whether or not the solution concept captures signaling and updating dynamics has fundamentally important implications for the conclusions that would be drawn from the same set of observational data. Thus, while both equilibrium concepts can be used to generate statistical models from the same game, our estimator will be useful to analysts interested in testing theories about crisis behavior in a manner that is faithful to the informational assumptions privileged in the literature.

At the same time, we show that analysts interested in making inferences about states' preferences from crisis outcomes inevitably face limitations that complicate model identification. These limitations stem not from any particular game tree or solution concept, but from the very task of estimating utility functions through a method that relies on revealed preferences. Because this is an exercise in revealed preference, utilities associated with outcomes that a given player is never in a position to choose among cannot be placed on comparable scales. In our simple crisis bargaining game, this meant that the value of the status quo to state B could not be pinned down relative to any of state B's other payoffs. Because both states in our game make choices that affect the probability that each of the other terminal nodes is reached, the payoffs associated with each outcome for player A can be placed on the same scale. For state B, the payoffs associated with terminal nodes other than the status quo can be made comparable.

Furthermore, because the payoffs are von Neumann–Morgenstern utilities, their exact values cannot be pinned down without fixing the scale within each comparable subset for each player. In practice, there are two ways to do this:

1. Normalize one payoff to zero within each comparable subset of payoffs for each player. In this case, all other comparable payoffs should be interpreted as the difference between the utility of the given outcome and the utility of the normalized outcome. Any covariates that are thought to influence the utility of the normalized outcome must appear in the equations for all other outcome payoffs.
2. Ensure that no covariate, including a constant, appears in the specification of every payoff within each comparable subset of payoffs for each player. In this case, the analyst must impose the restriction that every covariate has zero effect on at least one comparable outcome utility for each player.

Both methods in principle lead to identical inferences, though there are some reasons to prefer the second. The virtue of the second approach is that it allows the analyst to measure the direct effect of a covariate on the utility associated with a given outcome. With the first approach, by contrast, the analyst can only determine the effect of each covariate on the difference between given outcome utility and the utility of the normalized outcome. Hence, the second approach would allow us to say something like "Democracy in A increases its value for war," whereas the first approach only permits inferences like "Democracy in A increases its value for war relative to its value for the status quo." The latter conclusion leaves open the possibility that democracy decreases the value of the status quo rather than increasing the value of war. One can reclaim the first conclusion by assuming that democracy could not influence the status quo payoff, which is identical to restricting the effect of democracy on the status quo payoff to zero. Hence, whichever method is used, the cost of trying to infer the direct effects of covariates on outcome utilities is that the analyst will have to impose restrictions that will not, in principle, be testable.

In addition, depending on the specification of the game, there may be some payoffs for a given player that will have to be measured on different scales, meaning that they will not be comparable to one another. This is in addition to the well-known impossibility of interpersonal comparisons. These limitations do not invalidate the value of the exercise. Instead, they place a greater burden on our theories to tell us which covariates belong in the specification of each payoff and which do not. Thus, the role of theory in guiding our empirical models does not end once the strategic interaction and information structure have been specified.

The theoretical model must also guide one last, but crucial, element of this research project: the collection of data on which to apply the estimator. The primary challenge in applying the estimator developed here is to gather systematic information about the empirical frequency of the different outcomes to the game. Although several data sets on international crises already exist, none were collected with an explicit game structure in mind. A project is currently underway to augment existing data sets—initially, the International Crisis Behavior data set (Brecher and Wilkenfeld 1997)—with information about which terminal node of our game was reached in each observation. It should then be possible to learn about state preferences in international crises under the most suitable conditions: with both data collection and empirical estimation driven by the theoretical model of their strategic interaction.

Appendix: Identification in Empirical Models of Strategic Choice

A.1 *Identification of Multinomial Models*

We begin by characterizing the conditions under which a general class of multinomial models is identified in the sense of Rothenberg (1971).[9] This general class encompasses not only the strategic choice models described in the text, but also simple quantal response models such as probit and multinomial logit. Consider a data generating process in which one of a set of K alternatives is selected (independently) in each of a set of N trials. Let $\mathbf{Y}$ be an $N \times K$ matrix of ones and zeros such that y_{ik} is one if the kth outcome arises in the ith trial and zero otherwise (also, $\sum_k y_{ik} = 1$ for all $i = 1, 2, \ldots, N$). Let $p_{ik}(\theta)$ be the probability that alternative k occurs in trial i, and let $\theta \in \mathbb{R}^J$ be a J-dimensional vector of parameters. The function p_{ik} carries the index i to allow for variation in the probability of each outcome across observations as would typically arise from covariates. Note that $\sum_k p_{ik}(\theta) = 1$ for all $\theta \in \mathbb{R}^J$ and $i = 1, 2, \ldots, N$. The log-likelihood is of the form

$$\ln L(\mathbf{Y} \mid \theta) = \sum_{i=1}^{N} \sum_{k=1}^{K} y_{ik} \ln p_{ik}(\theta).$$

[9] Following Theorem 1 of Rothenberg (1971), we will say a likelihood function is locally identified if its information matrix $E(\frac{\partial^2 L(\theta)}{\partial\theta\partial\theta'})$ has full rank.

Differentiating the log-likelihood by θ yields

$$\frac{\partial \ln L}{\partial \theta} = \sum_{i=1}^{N} \sum_{k=1}^{K} \frac{y_{ik}}{p_{ik}(\theta)} \frac{\partial p_{ik}}{\partial \theta}.$$

Differentiating again by θ yields

$$\frac{\partial^2 \ln L}{\partial \theta \partial \theta'} = \sum_{i=1}^{N} \sum_{k=1}^{K} -\frac{y_{ik}}{p_{ik}(\theta)^2} \left(\frac{\partial p_{ik}}{\partial \theta}\right)\left(\frac{\partial p_{ik}}{\partial \theta}\right)' + \frac{y_{ik}}{p_{ik}(\theta)} \frac{\partial^2 p_{ik}}{\partial \theta \partial \theta'}.$$

Taking expectations over **Y** and canceling, we have

$$E\left(\frac{\partial^2 \ln L}{\partial \theta \partial \theta'}\right) = \sum_{i=1}^{N} \sum_{k=1}^{K} -\frac{1}{p_{ik}(\theta)} \left(\frac{\partial p_{ik}}{\partial \theta}\right)\left(\frac{\partial p_{ik}}{\partial \theta}\right)' + \frac{\partial^2 p_{ik}}{\partial \theta \partial \theta'}.$$

Because $\sum_k p_{ik}(\theta) = 1$, $\sum_k \frac{\partial p_{ik}(\theta)}{\partial \theta} = \mathbf{0}$ and $\sum_k \frac{\partial^2 p_{ik}(\theta)}{\partial \theta \partial \theta'} = \mathbf{0}$. Therefore,

$$E\left(\frac{\partial^2 \ \ln L}{\partial \theta \partial \theta'}\right) = \sum_{i=1}^{N} \sum_{k=1}^{K} -\frac{1}{p_{ik}(\theta)} \left(\frac{\partial p_{ik}}{\partial \theta}\right)\left(\frac{\partial p_{ik}}{\partial \theta}\right)'.$$

Letting **C** be a $KN \times J$ matrix with entries

$$\left[\frac{\partial p_{ik}}{\partial \theta_j}\right]_{(k-1)N+i,j}$$

and **D** be a $KN \times KN$ diagonal matrix with entries $[p_{ik}(\theta)]_{(k-1)N+i,(k-1)N+i}$, it follows that

$$E\left(\frac{\partial^2 \ \ln L}{\partial \theta \partial \theta'}\right) = \mathbf{C}'\mathbf{D}^{-1}\mathbf{C}.$$

For the probability model to be identified, $E(\frac{\partial^2 \ln L}{\partial \theta \partial \theta'})$ must have rank J (Rothenberg 1971). **D** has rank NK. The rank of a product is the minimum rank of the matrices that comprise the product. Thus, the identification of the model turns on **C** having full rank.

The row rank of **C** cannot be larger than $N(K - 1)$. Because $\sum_k p_{ik}(\theta) = 1$ for all $i = 1, \ldots, N$, rows of **C**, which share the same i, must sum to **0** and are linearly dependent:

$$\sum_k \left[\frac{\partial p_{ik}}{\partial \theta_1} \frac{\partial p_{ik}}{\partial \theta_2} \cdots \frac{\partial p_{ik}}{\partial \theta_J}\right] = [0\ 0\ \cdots\ 0].$$

Because there are N such linear dependencies, the rank of **C** cannot exceed $N(K - 1)$. In many cases the rank of **C** may be considerably smaller. For example, if $p_{ik} = p_k$, as in the stylized data generating process in Section 3,

$$\left[\frac{\partial p_{ik}}{\partial \theta_1} \frac{\partial p_{ik}}{\partial \theta_2} \cdots \frac{\partial p_{ik}}{\partial \theta_J}\right] = \left[\frac{\partial p_{i'k}}{\partial \theta_1} \frac{\partial p_{i'k}}{\partial \theta_2} \cdots \frac{\partial p_{i'k}}{\partial \theta_J}\right]$$

for all i and i', the row rank of $\mathbf{C}$ is only $K-1$. This is the "degrees of freedom" problem that arises in that context. In that model there were eight parameters associated with the two players' payoffs at each of four terminal nodes ($J=8$) and the data are simply counts of the frequency with which each of the four terminal nodes is reached ($K=4$). Thus, the rank of $\mathbf{C}$ is less than J and the model was not identified. In the text, we identified the model by treating five of the eight elements of θ as known. Although some of these restrictions were substantively innocuous, others were not.

A.2 *Identification in Random Utility Models*

Now consider a narrower class of multinomial models in which an "outcome" is the completion of one play of a given extensive-form game at a particular terminal node. In this class of models, the probabilities of the various terminal nodes being reached are functions of the utilities or average utilities associated with each terminal node for each player. The game is assumed to have players $l=1,\ldots,L$ with (average) payoffs ψ_{ilk} for players $l=1,\ldots,L$ and terminal nodes $t=1,2,\ldots,K$ in trial i.[10]

We define for each player a set of initial information sets, $\mathcal{I}_l$.

Definition. An initial information set is an information set such that no node that is an element of any of a given player's information sets precedes any node that is an element of an initial information set.[11]

In the game presented in the text, each player has one initial information set.

We further define a mapping T between the set of initial information sets for each player and the set of the terminal nodes $\mathcal{T}=\{1,2,\ldots,K\}$:

$$T:\mathcal{I}_l\rightarrow\mathcal{T}.$$

Because each subsequent node may, by definition, only be reached from one previous node, the set $\{T(\mathcal{O})\mid\mathcal{O}\in\mathcal{I}_l\}$ is a partition of the subset of $\mathcal{T}$ that includes those terminal nodes that are subsequent to a move by player l. Let $\mathcal{T}_l'$ be the set of terminal nodes that are not subsequent to a move by player l. $\mathcal{P}_l=\{T(\mathcal{O})\mid\mathcal{O}\in\mathcal{I}_l\}\cup\{\mathcal{T}_l'\}$ is a partition of $\mathcal{T}$. Letting $\mathbf{\Psi}_i=(\psi_{i11},\ldots,\psi_{i1K},\psi_{i21},\ldots,\psi_{i2K},\ldots,\psi_{iLK})$, we now define a class of *translation invariant* utility models of multinominal choice.

Definition. A *translation invariant* utility model of multinomial choice is one in which $p_{ik}(\mathbf{\Psi}_i+\mathbf{c})=p_{ik}(\mathbf{\Psi}_i)$ for all i and k and $\mathbf{c}$ is any $LK\times1$ vector such that

$$c_{K(l-1)+t}=\begin{cases}c & \text{if } t\in P\\ 0 & \text{otherwise}\end{cases}$$

for all l and all $P\in\mathcal{P}_l$ and $c\in\mathbb{R}$.

[10] Simple random utility models that involve only one player ($L=1$) are a special case of this class of models.

[11] In the definition of an extensive-form game the order of play must be specified. The order of play determines a partial order over the nodes of the game. That partial order formally defines the notion of one node "preceding" another. Less formally, the idea is simply that one node precedes another if it is encountered earlier in the game tree than the other node.

Less formally, this definition describes models in which a constant amount c can be added to each of the (average) payoffs for a particular player that are associated with a particular initial information set of that player without changing probabilities that the model assigns to every outcome. Both the QRE and PBE models described in the text are translation invariant.

It follows directly from the above definition that if a model is translation invariant then

$$\sum_{t \in P} \frac{\partial p_{ik}}{\partial \psi_{ilt}} = 0$$

for all i and l and all $P \in \mathcal{P}_l$. If we assume that each of the utilities for each of the players can be written as a linear function of an $M \times 1$ vector of coefficients $\boldsymbol{\beta}_{ik}$ and covariates $\mathbf{X}_i$,

$$\psi_{ilt} = \boldsymbol{\beta}'_{lt}\mathbf{X}_i,$$

then we have an $LKM \times 1$ vector of parameters $\theta = (\boldsymbol{\beta}_{11}, \ldots, \boldsymbol{\beta}_{L1}, \ldots, \boldsymbol{\beta}_{LK})$ to estimate. The elements of $\mathbf{C}$ are now

$$\left[\frac{\partial p_{ik}}{\partial \beta_{ltm}}\right]_{(K-1)i+k,(L-1)t+(M-1)l+m},$$

where

$$\frac{\partial p_{ik}}{\partial \beta_{ltm}} = \frac{\partial p_{ik}}{\partial \psi_{lt}} x_{im}.$$

Summing over the columns of $\mathbf{C}$ for a given l and m and for which $t \in P$ for any $P \in \mathcal{P}$, we find

$$\left[\sum_{t \in P} \frac{\partial p_{ik}}{\partial \psi_{lt}} x_{im}\right]_{(K-1)i+k} = \left[x_{im}\left(\sum_{k \in P} \frac{\partial p_{ik}}{\partial \psi_{lt}}\right)\right]_{(K-1)i+k} = [0]_{(K-1)i+k}.$$

Thus, if a given covariate is allowed to enter the equation for each of a given player's payoffs for all of the terminal nodes following a given initial node for that player, the column rank of $\mathbf{C}$ will not be full and the model will not be identified. Additionally, if $k \in \mathcal{T}'_l$ (if the terminal node k is achieved before player l moves) then none of the parameters associated with player l's utility for that terminal node can be identified.

This conclusion leads to the following two propositions related to the identification of transformation-invariant random-utility models of multinomial choice with linear index functions:

Proposition 1. *The effect of a given covariate* (*including a constant*) *on the utilities that a player receives cannot be freely estimated for all of the payoffs associated with the terminal nodes that follow from a particular initial information set* (*first move*) *for that player.*

Proposition 2. *Parameters characterizing a given player's utility over terminal nodes that are not subsequent to that player's initial information sets cannot be estimated.*

It is additionally possible to identify the model through equality (or inequality) restrictions and nonlinear index functions. If each index function is nonlinear,

$$\psi = g(\beta'_{\psi}\mathbf{X}),$$

where g is monotonic and nonlinear. Then, we have for each column of **C**,

$$\left[\sum_{t \in P} \frac{\partial p_{ik}}{\partial \psi_{lt}} g'(\beta_{lt}\mathbf{X}_i) x_{im}\right]_{(K-1)i+k} \neq [0]_{(K-1)i+k}.$$

Thus, if the index functions are nonlinear, inclusion of the same covariate in each of the payoffs does not create dependent columns in **C**.

References

Baumol, W. J. 1958. "The Cardinal Utility that is Ordinal." *Economic Journal* 68:665–672.
Brecher, Michael, and Jonathan Wilkenfeld. 1997. *A Study of Crisis*. Ann Arbor: University of Michigan Press.
Bueno de Mesquita, Bruce, and David Lalman. 1992. *War and Reason: Domestic and International Imperatives*. New Haven, CT: Yale University Press.
Fearon, James D. 1994a. "Signaling versus the Balance of Power and Interests." *Journal of Conflict Resolution* 38:236–269.
Fearon, James D. 1994b. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* 88:577–592.
Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49:379–414.
Guisinger, Alexandra, and Alastair Smith. 2002. "Honest Threats: The Interaction of Reputation and Political Institutions in International Crises." *Journal of Conflict Resolution* 46:175–200.
Kooreman, Peter. 1994. "Estimation of Econometric Models of Some Discrete Games." *Journal of Applied Econometrics* 9:255–268.
Lewis, Jeffrey B., and Kenneth A. Schultz. 2001. "Limitations to the Direct Testing of Extensive Form Crisis Bargaining Games." Paper presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL.
McKelvey, Richard A., and Thomas R. Palfrey. 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1:9–41.
Morrow, James D. 1985. "A Continuous Outcome Expected Utility Theory of War." *Journal of Conflict Resolution* 29:473–502.
Morrow, James D. 1989. "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining." *American Journal of Political Science* 33:941–972.
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.
Rothenberg, Thomas J. 1971. "Identification in Parametric Models." *Econometrica*, 39:577–591.
Sartori, Anne. 1998. *Deterrence by Diplomacy*. Ph.D. dissertation. Ann Arbor, MI: University of Michigan.
Schultz, Kenneth A. 1999. "Do Democratic Institutions Constrain or Inform? Contrasting Two Institutional Perspectives on Democracy and War." *International Organization* 52:233–266.
Schultz, Kenneth A. 2001. *Democracy and Coercive Diplomacy*. Cambridge, UK: Cambridge University Press.
Signorino, Curtis S. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review* 93:279–298.
Signorino, Curtis S. 2003. "Structure and Uncertainty in Discrete Choice Models." *Political Analysis* 11:316–344.
Smith, Alastair. 1996. "To Intervene or Not to Intervene: A Biased Decision." *Journal of Conflict Resolution* 40:16–40.
Smith, Alastair. 1998. "International Crises and Domestic Politics." *American Political Science Review* 92:623–638.
Smith, Alastair. 1999. "Testing Theories of Strategic Choice: The Example of Crisis Escalation." *American Journal of Political Science* 43:1254–1283.