

A Statistical Model of the Ultimatum Game*

Kristopher W. Ramsay[†]

Curtis S. Signorino[‡]

February 15, 2009

Abstract

In this paper we derive a statistical estimator to be used when the data generating process is best described as an equilibrium to the popular ultimatum bargaining game with private information and private values. This procedure gives the analyst the ability to estimate the effect of substantively interesting covariates on equilibrium behavior in this work horse bargaining model. Using Monte Carlo analysis we explore the small sample properties of this estimator and compare how the inference one makes with this model differ from those generated by linear and generalized linear models. We end by demonstrating the real world effect of using our estimator with a re-analysis of results from ultimatum lab experiments where subject covariates are hypothesized to explain bargaining behavior. We find, contrary to the experimental claim, there is no evidence of a “national” effect on offers in this data when the bargaining estimator is used.

*We would like to thank Tim Carter, Kevin Clarke, Songying Fang, Stephen Gent, Jeremy Kedziora, Jaehoon Kim, Randy Stone, and Robert Walker for their comments and helpful discussions. Support from the National Science Foundation (Grant # SES-0213771) and from the Peter D. Watson Center for Conflict and Cooperation is gratefully acknowledged.

[†]Department of Politics, Princeton University. email: kramsay@princeton.edu.

[‡]Department of Political Science, University of Rochester. email: curt.signorino@rochester.edu.

1 INTRODUCTION

Over the last twenty years, formal models and quantitative analyses have come a long way toward explaining how strategic actors bargain in a variety of political settings. (Banks 1990, Bennett 1996, Baron 1989, Fearon 1995, Huth and Allee 2002, Laver and Schofield 1990, London 2002, Morrow 1989, Wagner 2000, Powell 1987, Powell 1996). Game theory, in particular, has proved to be a useful tool for understanding the basic logic of bargaining in the face of conflicting interests. If the frequency with which a single idea or framework is cited or used in the literature is a measure of its importance, then the importance of bargaining models cannot be denied. For example, bargaining models have been applied by political scientists to analyze everything from the effects of open and closed rules on the distributive politics of legislative appropriation to the study of war initiation and termination (Baron and Ferejohn 1989, Fearon 1995, Mansfield, Milner and Rosendorff 2000). In fact, the theoretical and empirical study of bargaining is one of the few places where the different subfields of political science can identify one phenomenon that all agree is important and worthy of attention.

Results of numerous theoretical studies of the bargaining problem have pointed to the importance of asymmetric information and the “reservation values” of players in distributional politics. Yet, as is usually the case when scholars try to bring their theoretical model to the data it is difficult to specify the link between substantive variables, theories, and outcomes. This makes many hypotheses difficult to operationalize and even more difficult to test. Additionally, it is often the case that we would like to know the effects of particular substantive variables, like a congressman’s district demographics or whether a state possesses nuclear weapons, on the bargaining process. The theoretical models tell us something about the path by which these variables may influence outcomes. However, there is no “canned” statistical estimator for examining these effects.

As an alternative to the “theory down” approach to understanding bargaining, an increasingly sophisticated body of work has looked directly at the empirical relationship between substantive variables of interest, such as regime type, economic interdependence, institutional rules, legislative composition, and bargaining outcomes (Bennett 1996, McCarty and Poole 1995, Milner 1997, Werner 1999). However, lacking an explicit model of the process that generates the empirical data, and leaving out the choice-based path by which these variables influence decisions, it is often the case that selection and omitted variable bias plague the analysis (King, Keohane and Verba 1994). In particular, Signorino demonstrates that traditional linear and categorical estimation techniques

can lead to faulty inferences when the strategic data generating process is ignored during estimation (Signorino 1999, Signorino 2002, Signorino and Yilmaz 2003). It is unclear how reliable the inferences from these empirical models are given these findings.

What is called for in the bargaining literature is an integration of formal theoretical models and statistical methods. In particular, analysts need a statistical tool that permits them to make theoretically consistent inferences about the relationship between substantive variables, the bargain struck, and the probability of bargaining failure.¹ In other words, we need an estimator that explicitly models the strategic data generating process.

To move in this direction, we derive a statistical model for ultimatum bargaining games. After verifying that our theoretical model satisfies the minimum criteria for structural estimation, we derive an econometric estimator for the bargaining model. This model explicitly captures the relationship between the variables that affect the players's utilities and the outcomes of the bargaining in a strategic setting. Next, we conduct a Monte Carlo experiment by generating strategic bargaining data and then estimating the relationships between the regressors and the dependent variable(s). We estimate not only the statistical bargaining model, but also traditional OLS and censored variable models to give a point of reference for understanding the Monte Carlo results. Finally, we explore the small sample properties of our estimator to better understand its reliability and power in smaller data sets often found in political science.

2 THE MODEL

Political decision-making is often fundamentally a bargaining problem. That is, the essence of strategic decision-making between states, parties or leaders is largely about who gets what and when. One of the simplest and most popular bargaining models is the ultimatum game. In this section, we describe this bargaining model and then define a statistical model consistent with that data generating process. The estimator derived in the following section will be a straightforward structural implementation of the strategic model described below.

¹For examples of other work that has begun to deal with this problem in alternative settings see Wolpin (1987), Merlo and Wilson (1995), McKelvey and Palfrey (1996), Merlo (1997), and Merlo and Wilson (1998).

Figure 1: Ultimatum Game

2.1 The Ultimatum Game

Consider the usual bargaining arrangement, depicted in Figure 1, where two players must divide a contested prize. The issue being bargained over could be territory, a budget, or some policy in a 1-dimensional space. We represent this contested prize as a closed and bounded interval in R . Without further loss of generality, we consider normalized intervals of the form $[0, Q]$.

To start, player 1 offers some division of the prize $(Q - y, y)$, where player 1's allocation is $Q - y$ and player 2's is y . Player 2 then decides whether to *accept* or *reject* player 1's offer. If player 2 accepts, they divide the prize according to player 1's offer. If player 2 rejects the offer, they receive some reservation amount, which may differ between the players.

Our information structure follows a prominent one found in much of the applied bargaining literature. We assume that the players and the analyst have complete information about the size of the pie Q and about player 1's offer y . We also assume that each player's reservation utility has two components: one that is public R_i and one that is private ϵ_i . From a game theoretic perspective, the private component ϵ_i defines player i 's type. From an econometric perspective, the reservation utilities are random utility functions in the sense of Luce (1959) and McFadden (1976).

Finally, assume that each player i 's type ϵ_i is an independent and identically distributed random variable, drawn from a well defined probability distribution F_{ϵ_i} on R , with density f_{ϵ_i} . Let this

random variable have mean 0 and finite variance. Also assume that the players' (and analyst's) prior beliefs regarding the other player's type are F_{ϵ_i} .

It is worth noting that, while this is not the most general information structure one might consider, it is consistent with many substantively motivated models. Here players' types are private values, and do not directly influence the other players payoffs over outcomes. Private value uncertainty like this is found in a variety of common applied bargaining models, like those where the costs of war or litigation are privately known or situations where the type of a player (for example those that are "resolved" or those facing domestic audience costs) most directly affects the costs or benefits of bargaining failure. Alternatively, players could have shocks to common valued components of their utilities, implying correlation in payoffs, or independent shocks to each settlement. These different information structures imply completely different strategic and structural models and we do not consider them here.²

In this game each player's strategy can be characterized by a mapping from types into actions: $\sigma_i : \epsilon_i \rightarrow A^i$, $i = \{1, 2\}$, where A^i defines the actions available to player i . Since player 1 is making the ultimatum offer her action, or proposal, can be represented by a number in $[0, Q]$. Player 2 is then left to accept or reject the offer, so $A^2 = \{accept, reject\}$.

Assuming the players utilities are linear in the share of the pie, our random utility structure and information assumptions lead to the following simple utilities over outcomes and the game tree found in Figure 1:

$$\begin{aligned} u_1(y, accept) &= Q - y \\ u_2(y, accept) &= y \\ u_1(y, reject) &= R_1 + \epsilon_1 \\ u_2(y, reject) &= R_2 + \epsilon_2 \end{aligned}$$

In equilibrium, the ultimatum game has player 1 making an offer that balances the marginal utility of increasing the probability that an offer is accepted and the marginal utility of a larger amount of y . Player 2, knowing her own type, chooses the alternative that maximizes her utility.

²Different sources of private information can induce different equilibrium behavior. For bargaining examples, see Fey and Ramsay (2009). For discrete choice games in extensive form, see Signorino (2003).

2.2 Existence and Uniqueness of Equilibrium

An equilibrium to such a “statistical” ultimatum bargaining game, where each player knows the other has random utilities, is equivalent to a perfect Bayesian Nash equilibrium of a game in which the types of the players are private information and those types reference private value components of players’ utilities. This means we can use well-known game theoretic tools to begin to specify both our theoretical predictions and our empirical estimator.³ Recent research has shown that traditional existence results do not guarantee the existence of an equilibrium in Bayesian games with unbounded type spaces and continuous actions sets (Meirowitz 2003). It is, however, easy to show that there always exists a generically unique perfect Bayesian equilibrium to this ultimatum game. Proposition 1 gives a sufficient condition for this to be true in our game.

Proposition 1. *If F_{ϵ_2} is log-concave, then there exists a unique perfect Bayesian-Nash equilibrium to the statistical ultimatum game.*

This result is not surprising, as this game is closely related to the ultimatum game of Fearon (1995), whose Claim 2 showed the existence of a generically unique equilibrium for his game. We therefore leave the rather tedious proof for the appendix, and instead sketch the logic of each player’s equilibrium choice. From this discussion uniqueness of the equilibrium will be obvious. To start, assuming player 1 has made an offer y , player 2 chooses between that offer and her reservation value $R_2 + \epsilon_2$. In any equilibrium player 2 has an easy question to answer, which is better: the settlement or disagreement? Thus in any equilibrium player two plays a simple cutpoint strategy where types who prefer the settlement to their reservation payoff accept offers and the remaining types reject.⁴

Player 1 can reason that player 2 will choose in such a way, given the offer she makes, but as player 1 does not observe ϵ_2 , she must assess the probability that player 2 will accept his offer. In equilibrium, player 1 correctly conjectures two will player her cutpoint strategy, so the relevant question for one is: what is the probability that 2 accepts an offer y ? Given F_2 and two’s strategy

³For examples of work on statistical game theory see McKelvey and Palfrey (1996), Beja (1992), van Damme (1991).

⁴The equilibrium is generically unique because there are many things player 2 could do when she is indifferent and no matter which action she takes in that special case, player 1’s offer strategy is unaffected, so they are all consistent with equilibrium.

it must be

$$\begin{aligned}
\Pr(\text{accept}|y) &= \Pr(y \geq R_2 + \epsilon_2) \\
&= \Pr(\epsilon_2 \leq y - R_2) \\
&\equiv F_{\epsilon_2}(y - R_2).
\end{aligned} \tag{1}$$

Now, player 1 has a simple concave optimization problem, given two's strategy. His expected utility from an offer $y \in [0, Q]$ is

$$Eu_1(y|Q) = F_{\epsilon_2}(y - R_2) \cdot (Q - y) + (1 - F_{\epsilon_2}(y - R_2)) \cdot (R_1 + \epsilon_1),$$

subject to the constraints that $0 \leq y \leq Q$. By the F.O.C. and the log-concavity of f_{ϵ_2} , 1's optimal offer is the unique y^* that implicitly solves

$$y^* = Q - R_1 - \epsilon_1 - \frac{F_{\epsilon_2}(y^* - R_2)}{f_{\epsilon_2}(y^* - R_2)}, \tag{2}$$

when the constraints on the optimization problem are slack, and is 0 or Q otherwise. This single optimization problem and the binary choice described above completely characterize rational play to this game and, therefore, the equilibrium is unique.

An important question concerns how the optimal offer changes with changes in the reservation values. Moreover, because the observable reservation values are specified with regressors, we would like to understand the comparative statics for two cases: (1) for regressors that appear either in X or Z , but not both; and (2) for regressors that appear in both X and Z . Let $m(x) = \frac{F_{\epsilon_2}(x)}{f_{\epsilon_2}(x)}$. We leave the proof for the appendix and simply note the following:

Proposition 2. *Suppose*

- *Regressor x appears only in $X\beta$, with associated coefficient β_x ;*
- *Regressor z appears only in $Z\gamma$, with associated coefficient γ_z ; and*
- *Regressor v appears in both $X\beta$ and $Z\gamma$, with associated coefficients β_v and γ_v , respectively.*

Then the unconstrained optimal offer y^ will be*

- *Unconditionally monotone in x with direction $\text{sign}(dy^*/dx) = \text{sign}(-\beta_x)$,*
- *Unconditionally monotone in z with direction $\text{sign}(dy^*/dz) = \text{sign}(\gamma_z)$,*

- *Conditionally monotone in v if and only if $\text{sign}(dy^*/dv) = \text{sign}[\gamma_v m'(y^* - z\gamma_z - v\gamma_v) - \beta_v]$ is constant for all v , and nonmonotonic in v otherwise.*⁵

The intuition here is relatively straightforward. Increasing player 1’s reservation value (via x) decreases the amount she is willing to offer player 2. Conversely, increasing player 2’s (observable) reservation value (via z) increases the amount player 1 thinks she will need to offer player 2. If both of these countervailing pressures are present simultaneously (e.g., when a variable like v enters both reservation values), then the relationship is more complicated. Depending on the magnitudes and signs of the coefficients β_v and γ_v , the optimal offer may be nonmonotonic in v . It is worth noting that the above results are for the unconstrained optimal offer. When the constraints are applied, the strict monotonicity results reduce to weak monotonicity due to potential censoring.

3 THE LOGIT ULTIMATUM ESTIMATOR

For the remainder of this paper we study a particular class of distribution functions for the error term and explore the statistical properties of the resulting estimator. Let us assume we have data on both player 1’s and player 2’s actions — i.e., assume we can measure and code y and Q for each observation, as well as whether player 2 accepted or rejected the offer. Let the public portion of the players’ reservation values be $R_1 = X\beta$, and $R_2 = Z\gamma$, where X and Z are sets of substantive regressors. Our interest is in estimating the effects of X and Z on y and player 2’s decision.⁶

Because the outcome of the bargaining model consists of two dependent variables — 1’s offer and 2’s decision — our probability model is a joint density over those random variables. Recall from Proposition 1 that the requirement for uniqueness (in combination with existence) is log-concavity of F_{ϵ_2} . For our estimator, we will assume that the players’ types, ϵ_1 and ϵ_2 , are drawn i.i.d logistic. The i.i.d. assumption greatly simplifies matters by reducing the joint density to the product of two univariate densities. Moreover, Bagnoli and Bergstrom (2005) demonstrate that the logistic

⁵We use the terms “unconditionally monotone” and “conditionally monotone” as defined in Signorino & Yilmaz (2003:563-64). $g(x, y)$ is said to be conditionally monotone in x if $\text{sign}(dg(x, y)/dx)$ is constant for all x , conditional on y . $g(x, y)$ is said to be unconditionally monotone in x if $\text{sign}(dg(x, y)/dx)$ is not only constant for all x but also in the same direction regardless of y . We also thank an anonymous reviewer for pointing out we can give this more general proposition than we had in previous drafts.

⁶In this case, we assume all players and the analyst share the same public information embodied in the regressors X and Z , as well as their effects β and γ on the dependent variables. This may not always be an appropriate assumption. In another paper, Kedziora, Ramsay and Signorino (2009) examine a statistical model where the analyst has more information about players than the players have about each other.

distribution is log-concave, so we have existence and uniqueness of the logit ultimatum equilibrium. Although our results require only log-concavity of ϵ_2 the most natural assumption is to treat the players symmetrically, which is what we do here.

Given the distributions on ϵ_1 and ϵ_2 , the structural estimator for this statistical bargaining game can be derived straightforwardly. First, consider player 2's decision concerning whether to accept or reject player 1's offer. Since F_{ϵ_2} is logistic, Equation 1 is just the logistic probability

$$\begin{aligned}\Pr(\text{accept}|y) &= \Lambda(y - Z\gamma) \\ &= \{1 + \exp[-(y - Z\gamma)/s_2]\}^{-1}\end{aligned}\tag{3}$$

where $\Lambda(\cdot)$ is the logistic c.d.f. and s_2 is the logistic scale parameter for ϵ_2 .⁷ Although we usually ignore scale parameters in logit and probit models because of identification, the scale parameters are identified in the logit ultimatum estimator. We will return to this later.

For player 1, the distribution of $y^*(\epsilon_1)$, is more complicated. Equation 2 provided the implicit equilibrium condition for player 1's optimal offer. Substituting the logistic distribution for F_{ϵ_2} gives

$$y^* = Q - X\beta - \epsilon_1 - \frac{\Lambda(y^* - Z\gamma)}{\lambda(y^* - Z\gamma)},\tag{4}$$

where $\lambda(\cdot)$ is the logistic p.d.f.

At this point, we must solve for y^* in terms of player 1's type ϵ_1 . Then, given that the analyst does not observe ϵ_1 , we must derive the probability of observing a given offer y^* , based on the assumed distribution for ϵ_1 . Solving for y^* as a function of ϵ_1 produces

$$\begin{aligned}y^*(\epsilon_1) &= Q - X\beta - \epsilon_1 - \frac{1}{1 + e^{-(y^* - Z\gamma)/s_2}} \cdot \frac{s_2 [1 + e^{-(y^* - Z\gamma)/s_2}]^2}{e^{-(y^* - Z\gamma)/s_2}} \\ &= Q - X\beta - \epsilon_1 - s_2 \left[1 + e^{(y^* - Z\gamma)/s_2}\right]\end{aligned}\tag{5}$$

$$= Q - X\beta - \epsilon_1 - s_2 \left[1 + \mathcal{W}\left(e^{\frac{Q - X\beta - \epsilon_1 - s_2 - Z\gamma}{s_2}}\right)\right].\tag{6}$$

where \mathcal{W} is Lambert's \mathcal{W} .⁸ Simple differentiation of y^* with respect to ϵ_1 shows that y^* is a monotonic function of ϵ_1 and, therefore, we can derive the density function for equilibrium offers.

⁷A logistically distributed random variable ϵ with scale parameter s has variance $V(\epsilon) = \pi^2 s^2 / 2$.

⁸Lambert's \mathcal{W} solves transcendental functions of the form $z = we^w$ for w . For a discussion of Lambert's \mathcal{W} and the algorithm to solve the transcendental equation see Corless, Gonnet, Hare, Jeffrey and Knuth (1996), Fitsch, Shafer and Crowley (1973), Barray and Culligan-Hensley (1995), and Valluri, Jeffrey and Corless (2000). Lambert's \mathcal{W} has nice properties and makes the probability distribution of y^* easy to characterize. First, Lambert's \mathcal{W} is single valued on R_+ , and since $e^\alpha \geq 0$ for all $\alpha \in R$, it is single-valued where we need to use it. Second, \mathcal{W} 's first and second derivatives exist and are well behaved.

To derive the density for the optimal offer $f_{y^*}(y^*)$, we apply the method of monotonic transformation (Casella and Berger 2002, Thrm 2.1.5), producing⁹

$$f_{y^*}(y^*) = \frac{e^{-\{Q-y^*-X\beta-s_2[1+e^{(y^*-Z\gamma)/s_2}]\}/s_1}}{s_1 \left\{1 + e^{-\{Q-y^*-X\beta-s_2[1+e^{(y^*-Z\gamma)/s_2}]\}/s_1}\right\}^2} \cdot \left[1 + e^{(y^*-Z\gamma)/s_2}\right] \quad (7)$$

with cumulative distribution function

$$F_{y^*}(y^*) = \frac{1}{1 + e^{\{Q-y^*-X\beta-s_2[1+e^{(y^*-Z\gamma)/s_2}]\}/s_1}} \quad (8)$$

The constraint on the action space of player 1, however, implies that the observed y^* is censored both from above and below. This censored distribution of offers leads to the following likelihood. Let y be the observed offer and define a set of dummy variables δ_k $k \in \{0, y, 1\}$ such that $\delta_0 = 1$ if $y = 0$, $\delta_y = 1$ if $0 < y < Q$, and $\delta_1 = 1$ if $y = Q$. That is, much like a censored (Tobit) model, we can think of there being a “latent” best offer that we only observe when its realization is in the constraint set. Otherwise we see the best feasible offer, i.e., a boundary point. Next, code player 2’s acceptance as $\delta_{accept} = 1$ if she accepted the offer and $\delta_{accept} = 0$ if she rejected the offer.¹⁰

Assuming we have data on both player 1’s and player 2’s actions (i.e., y and δ_{accept}), then the likelihood would be

$$L = \prod_{i=1}^n \left\{ F_{y^*}(0)^{\delta_0} \cdot f_{y^*}(y)^{\delta_y} \cdot [1 - F_{y^*}(Q)]^{\delta_1} \times \Pr(accept|y)^{\delta_{accept}} \cdot [1 - \Pr(accept|y)]^{1-\delta_{accept}} \right\} \quad (9)$$

where the observation index has been omitted. Deriving the log-likelihood from this is straightforward. We then have a log-likelihood function for our data in terms of distributions already derived, which are functions of our regressors, and an explicit model of the ultimatum game. Estimates of β , γ , s_1 , and s_2 may be obtained using maximum likelihood estimation (or, with simple extension, via Bayesian MCMC).

It is interesting to note that all of the parameters — β , γ , s_1 , and s_2 — are individually identified in this model. Typically, in logit or probit models, we can only estimate the regression parameters and the variance parameter to scale. In this case, identification of γ and s_2 are driven by our assumptions concerning what the players know and how that enters their payoffs. Consider Equation 3, player 2’s probability of accepting. If no regression parameter is associated with the

⁹Derivation using the method of transformations is shown in the appendix.

¹⁰For a similar random utility motivation for censoring see Amemiya (1984).

offer y , then the usual logit (or probit) identification issue is not present here. It is important to note that this is not an ad hoc assumption made to ensure identification. Rather, we started by assuming that the players were bargaining over a known pie Q and that player 1 makes a known offer y . It is this shared observability of the offer y that identifies γ and s_2 in Equation 3.

The regression and variance parameters are also identified through player 1's choice. Equations 7 and 8 show that s_1 and s_2 interact with terms (e.g., Q , y^* , and 1) with which no other regression parameters are associated. In other words, a change in s_1 or s_2 cannot be "compensated" with a change in β or γ to produce the same probability value. The noteworthy aspect of these two equations is that player 2's parameters are identified through our assumptions about player 1's information concerning player 2.

Finally, suppose we had data only on player 2's or player 1's actions, but not both. For a given dependent variable (i.e., player 1 or player 2's action), we could simply use the appropriate density (already derived) as the basis of our log-likelihood equation.


4 MONTE CARLO ANALYSIS

In this section, we present two sets of Monte Carlo analyses. In the first, we conduct an analysis of the small sample properties of the estimator. In the second, we demonstrate the bias that can occur when typical regression techniques (e.g., OLS, FGLS, and Tobit) are used to analyze bargaining data, even when the underlying relationships are unconditionally monotonic.


4.1 Small Sample Properties

Many researchers are faced with data sets of relatively small sample size. Moreover, it may be difficult and costly to collect additional data. Given this predicament, analysts are often interested in how a particular estimator performs in small to medium-sized samples. With this concern in mind, we report a series of Monte Carlo experiments where we set the number of observations at levels consistent with the size of smaller data sets found in experimental and field work.


For the Monte Carlo analysis, the ultimatum game in Figure 1 was used as the data generating process. For a given observation, the value of the disputed good, Q , was drawn from a uniform distribution on $[0, 10]$. The public reservation values consisted of a single variable for each player, x for player 1 and z for player 2, with associated parameters β_u and γ_u , respectively, set to one.




(a) Estimates of β_u



(b) Estimates of γ_u



(c) Estimates of $\ln(s_1)$



(d) Estimates of $\ln(s_2)$

Figure 2: Small Sample Monte Carlo Results. *The plots display the distributions of the parameter estimates for samples of size $N=50$ (light grey), $N=100$ (grey), and $N=200$ (black). In each case, the true value of β_u , γ_u , s_1 , and s_2 was one (or, equivalently, $\ln(s_1)$ and $\ln(s_2)$ were each zero.)*

x and z were drawn from i.i.d. uniform distributions on the interval $[0, 6]$.¹¹ Lastly, the private information for players 1 and 2, ϵ_1 and ϵ_2 , respectively, were drawn from i.i.d. logistic distributions, with variance parameters $s_1 = 1$ and $s_2 = 1$, respectively. Based on this, player 1 determines his optimal offer y using Equation 6 and the constraint $0 \leq y \leq Q$. Given that, player 2 makes a decision by comparing y to her reservation $z\gamma_u + \epsilon_2$. The data for a given observation then consists of Q, y, δ_{accept}, x , and z .

The Monte Carlo analysis was conducted for three sample sizes: $N=50, 200$, and 500 observations. In each iteration of the analysis, a sample was generated (as detailed above) and the logit ultimatum estimator was used to recover estimates of β_u, γ_u, s_1 , and s_2 . This procedure was conducted 5000 times for each sample size, each time saving the parameter estimates.

Figure 4.1 depicts the distribution of the estimated model parameters for $N=50, 200$, and 500 sample sizes. The results are encouraging. Sample densities seem to approach normal distributions, even with as few as one or two hundred observations. As expected, the smaller the sample size the higher the variance of the estimated values. Somewhat surprisingly, though, the bias is relatively small for even very small sample sizes. For example with $N = 50$, the distribution of $\hat{\beta}_u$ and $\hat{\gamma}_u$ are already normal with means incredibly close to the population equation values. Only in the estimates of the variance parameters ($\ln(s_1)$ and $\ln(s_2)$) do we pick up meaningful bias, which disappears for N near 200. This suggests that the statistical bargaining model can be usefully applied to reasonably small data sets.

4.2 Bias in Alternative Methods

When analyzing bargaining data — whether offers or acceptances — it is currently common practice to employ standard techniques, such as OLS, FGLS, Tobit, Logit, and Probit. Although it is beyond the scope of this paper to analyze all the reasons for this, we highlight one very sensible and theoretically motivated argument. This argument suggests that a structural statistical model need not be derived directly from a formal theoretic model. Rather, if one can demonstrate monotonicity through comparative statics analysis, then off-the-shelf parametric models with a linear link function should be perfectly appropriate for data analysis. As we demonstrated in Section 3, the optimal offer is unconditionally monotonic in any regressor that appears in only one player’s reservation value. If we refrain from including regressors that are common to both players’ reserva-

¹¹As we will discuss in the next section, a larger reservation value makes censoring more likely (all else equal). We chose this value to represent an average or middle case.

tion values, we then have an ideal case to assess the above argument, as well as the use of standard models in general.

In order to investigate the possibility of bias in standard techniques, we conducted a second set of Monte Carlo simulations, similar to the first, but with a larger sample size and, of course, estimating alternative models. For this set of Monte Carlo experiments, the data generating process was the same as in Section 4.1 (e.g., $\beta_u = \gamma_u = s_1 = s_2 = 1$), but with the following exceptions. In each iteration of the Monte Carlo simulation, $N=1000$ observations were generated for the sample. We noticed that the larger the maximum size of the observable reservation value, the greater the number of the observations that were censored by the bargaining constraint. Therefore, to assess the effect of censoring on the inferences from the traditional models, separate Monte Carlo experiments were conducted for each value of $\text{Max R} \in \{2, 4, 6, 8, 10\}$, in which x and z were drawn from $U[0, \text{Max R}]$.

In each iteration of the Monte Carlo analysis, the full logit ultimatum model was estimated for comparison with three alternative models. For player 2’s acceptance data, a traditional logit model was estimated by regressing δ_{accept} on the offer y and player 2’s reservation variable z , with associated parameters ω_l and γ_l , respectively. For the offers data y , two models were estimated. The first was a Normal model of the form

$$y = \text{const}_n + \beta_n x + \gamma_n z + \theta_n Q + \epsilon,$$

estimated via maximum likelihood estimation.¹² The variable Q was included because the offer will depend not only on x and z , but also on the size of the prize. In addition to the Normal model, we estimated the Tobit model

$$y = \text{const}_t + \beta_t x + \gamma_t z + \epsilon$$

with censoring below at zero and above at Q . After estimating each model, the parameter estimates were saved. This was repeated for 2000 Monte Carlo iterations to form a density of the estimates.

Tables 1 displays the results of the logit ultimatum model (“Ultimatum”), estimated using the full set of data, as well as the results of a simple logit model estimated using only player 2’s data. The table is divided into five sections, one for each value of Max R . For each value of Max R , the average number of censored observations is displayed (out of 1000 observations for each

¹²The Normal model will produce results essentially identical to OLS. We chose an MLE-based Normal model rather than OLS for a more straightforward comparison with the Tobit model.

Max R	Number Censored	Ultimatum		Logit	
		$\hat{\beta}_u$	$\hat{\gamma}_u$	$\hat{\omega}_l$	$\hat{\gamma}_l$
2	253	.99 (.06)	.99 (.02)	1.00 (.07)	-1.00 (.10)
4	331	1.00 (.03)	1.00 (.02)	1.00 (.07)	-1.00 (.07)
6	421	1.00 (.02)	1.00 (.02)	1.00 (.08)	-1.01 (.07)
8	511	.99 (.02)	.99 (.02)	1.00 (.08)	-1.01 (.07)
10	593	1.00 (.02)	1.00 (.02)	1.01 (.09)	-1.01 (.08)

Each set of results is based on 2000 Monte Carlo iterations, each with samples of size N=1000. The mean parameter estimate is shown on top. The standard deviation is shown below in parentheses.

Table 1: Monte Carlo Results for Ultimatum and Logit Models.

sample). Each cell in the table reports the mean (top) and the standard deviation (bottom) of the distribution of the estimates for each group of 2000 Monte Carlo iterations.

As is clear in Table 1, $\hat{\beta}_u$ and $\hat{\gamma}_u$ are recovered by the logit ultimatum model on average, and the standard deviations of the estimates are small. Moreover, the estimates continue to be unbiased and have low variance, even as the censoring exceeds half of the observations on average per iteration (i.e., as Max R increases to 8 or 10).

It is also clear that a properly specified logit can be used alone to assess the effect of offer size and substantive regressors on the success or failure of bargaining. In fact, the logit’s estimates are consistently close to the true values. The negative sign on $\hat{\gamma}_l$ is correct and simply a result of specifying the regression equation as $\omega_l y + \gamma_l z$. As before, increased censoring due to Max R does not affect the consistency of the estimates. Not surprisingly, the effect of z is estimated more precisely when all of the data is used with the “Ultimatum” model, rather than when limited only to data on player 2’s decision.

Table 2 displays the results of using the uncensored Normal model or a Tobit model to analyze the relationship between the regressors and the size of the offer y . First consider the uncensored Normal model. A number of scholars, looking at the continuous nature of the offer variable, have chosen to use OLS and FGLS to analyze such data.¹³ For the uncensored Normal model, a constant was estimated, along with β_n and γ_n (the effects of x and z , respectively), an effect θ_n for the size of the prize Q , and the estimate of the variance σ_n^2 .

The uncensored Normal results highlight two potential problems with using OLS on bargaining data. First, the estimates of the parameters change significantly as the number of censored observations increases. In particular, $\hat{\gamma}_n$ ranges from .09 to .47. Second, every estimate looks as though it is a very good estimate, i.e., the standard deviation of the density is small.

It is already well known, however, that censoring can significantly affect estimates in a traditional OLS model. Moreover, a sophisticated empiricist would likely conclude that the strategic process would, in some way, censor the observed data. Therefore, she may account for this in her estimation, and analyze the data using a Tobit model, which is censored here above at Q and below at 0.

The Tobit results are also found in Table 2. As in the Normal case, the standard deviations of the estimates are small, leading one to believe the estimates are, in some sense, good. Perhaps most interesting, Table 2 shows that the Tobit model does not adequately control for the censoring in the

¹³See ? and Henrich, Boyd, Bowles, Camerer, Fehr, Gintis and McElreath (2001).

Max R	Number Censored	Normal					Tobit			
		$Const_n$	$\hat{\beta}_n$	$\hat{\gamma}_n$	$\hat{\theta}_n$	$\hat{\sigma}_n^2$	$Const_t$	$\hat{\beta}_t$	$\hat{\gamma}_t$	$\hat{\sigma}_t^2$
2	253	-.42 (.05)	-.25 (.03)	.47 (.03)	.29 (.01)	.29 (.02)	.92 (.11)	-.35 (.07)	.56 (.07)	1.69 (.08)
4	331	-.49 (.07)	-.31 (.02)	.35 (.02)	.36 (.01)	.58 (.03)	1.13 (.15)	-.48 (.05)	.43 (.05)	3.23 (.17)
6	421	-.22 (.09)	-.34 (.02)	.23 (.02)	.38 (.01)	.91 (.04)	1.52 (.18)	-.59 (.05)	.31 (.04)	4.93 (.27)
8	511	.19 (.09)	-.34 (.02)	.15 (.02)	.35 (.01)	1.17 (.06)	1.92 (.21)	-.69 (.04)	.22 (.04)	6.32 (.39)
10	593	.53 (.09)	-.30 (.01)	.09 (.01)	.31 (.01)	1.33 (.07)	2.29 (.23)	-0.78 (.04)	.16 (.04)	7.34 (.51)

Each set of results is based on 2000 Monte Carlo iterations, each with samples of size N=1000. The mean parameter estimate is shown on top. The standard deviation is shown below in parentheses.


Table 2: Uncensored Normal and Tobit Regression Models

bargaining offers. In particular, the parameter estimates change as the number of censored cases increase. The estimates for $\hat{\beta}_t$ change from $-.35$ to $-.78$, an increase in magnitude of 122 percent. Similarly, the estimates of $\hat{\gamma}_t$ change from $.56$ to $.16$, a decrease in magnitude of 71 percent.


Because the Ultimatum, Tobit, and Normal models are all based on different structural and/or distributional assumptions, it is important to compare their estimated conditional expectations as a function of x and z . Figure 2 displays the expected offer for four configurations of x and z . In each panel, either x or z is held constant at 0 or 10 and the other is varied from 0 to 10. The conditional expectations are based on the parameter estimates where $\text{Max R} = 10$ — i.e., allowing for any observed reservation level between 0 and 10. The solid, dashed, and dotted lines correspond to the expected offers from the Ultimatum, Tobit, and Normal models, respectively.

As Figure 2 shows, the Tobit model (which allows for censoring) and the Normal model (which does not) are very similar, not only in terms of the effects of x and z on the expected offer, but also on the size of the offer itself. Figures 2(a)-2(d) also suggest that we can characterize when the Tobit and Normal models will be closer to or diverge from the ultimatum model. Consider first Figures 2(a) and 2(d). These are situations where increasing a player’s reservation level has little effect on the offer. For example, in 2(a), player 2’s observed reservation level is at the minimum ($z = 0$). In this case, player 1 knows that player 2 is likely to accept if she makes even a small offer. When her own reservation is near zero, player 1 will give up a small amount of the prize. However, as player 1’s reservation (x) increases, she offers less. In 2(d), player 1’s reservation is at the highest level ($x = 10$), making her unwilling to give up much of the prize. Regardless of player 2’s reservation level, player 1 makes a negligible offer. In these two cases — when the reservation level has little effect on the offer — the Tobit and Normal models are fairly close approximations of the ultimatum model.


Now consider Figures 2(b) and 2(c). In these cases, the ultimatum model shows that a change in the observed reservation level has a substantial effect on the offer. For example, in Figure 2(b), player 1 knows that player 2 has a very high reservation ($z = 10$). Therefore, the size of the offer will depend on player 1’s reservation level. When player 1’s own reservation is near zero, she is willing to offer almost all of the prize. However, when player 1’s reservation is near the maximum, her offer will be negligible. In Figure 2(c), player 1’s reservation is assumed to be at the minimum ($x = 0$). In this case, player 1 will take just about anything. Therefore, as player 2’s reservation level increases, player 1’s offer also increases. In these cases, Tobit and the Normal model underestimate the effect of x and z — the slopes of their lines are attenuated. Moreover, the




(a) $z = 0$



(b) $z = 10$



(c) $x = 0$



(d) $x = 10$

Figure 3: Effect of x and z on Expected Offer. *Solid: Ultimatum Game. Dashed: Tobit. Dotted: Normal.*

Ultimatum model allows for much larger offers, while the maximum offers in the Tobit and Normal models are only about 4.

Finally, consider the question of whether player 2's reservation z has a substantively large effect on the offer. Figures 2(c) and 2(d) suggest that if one were to estimate the Tobit or Normal models, one would infer that player 2's reservation z does not have much of a substantive effect (however statistically significant) on the offer. By construction, we know this is incorrect. Indeed, the ultimatum model in Figure 2(c) shows that player 2's reservation value can have a large effect on the expected offer.

From these simulations, we see that the appropriate statistical method can depend on the question one wishes to answer. In this case, the logistic regression proved to be an acceptable — although slightly less efficient — method for investigating how different variables affect bargaining failure, given data on the offers. This is not terribly surprising, since the logit model is the right structural model for player 2's decision.¹⁴ However, if we are interested in saying something about how substantive variables affect the kind of bargain that is struck — and, in particular, the size of the offer — then traditional techniques are inappropriate. The Monte Carlo experiments show that (1) these techniques produce incorrect estimates of β and γ , (2) that the estimates change as a function of the distribution of the reservation values of the players, and (3) that inferences based on the marginal effects will at times be misleading. Since we cannot know, *ex ante*, the empirical distribution of the reservation payoffs, independent of the estimation of $\hat{\beta}$ and $\hat{\gamma}$, OLS and Tobit results for real world data are not reliable estimates. Moreover, this analysis was conducted for a model where all relationships were unconditionally monotonic. Given the specific structure of the strategic interaction in the ultimatum game, researchers should be wary of using off-the-shelf parametric methods, unless those methods are demonstrated to be consistent with the assumptions of the structural model.

5 AN APPLICATION TO EXPERIMENTAL BARGAINING DATA

One thing that makes the study of bargaining special is the interdisciplinary interest in the process. Everything from bargaining between multinational cooperations and states over terms of foreign

¹⁴If, on the other hand, we had reason to believe that player 1's private information was correlated with player 2's, then using the sample of bargaining success/failure data alone would induce a Heckman-like selection bias. Similarly, if the interaction contained multiple stages and player 2's decision was based on expectations of what player 1 might offer in the future, then the logit estimates would likely suffer from a functional form misspecification bias.

investment, to the resolution of territorial disputes, to anthropologist’s interests in how social and personal characteristics affect the “rational” behavior of individuals, has been studied in the context of bargaining models.

Here we analyze a data set on experimental ultimatum games conducted by Bothelho, Harrison, Hirsch and Ruström (2005). Their work and the experimental data they present provides an effect foil and allows us to point to the substantive differences between using our logit bargaining estimator versus OLS.

In the experiments that generated this data, the researchers sought to isolate the effects of demographic variables on the bargaining process. They start from the widely known experimental observation that, when people “play” the ultimatum game, they do not play the (complete information) game’s subgame perfect Nash equilibrium. In particular, ultimatum experiments consistently show that proposers giving receivers larger shares of the pie than is predicted by strict income maximizing behavior. In a number of previous studies, most noticeably Roth, Prasnikar, Okuno-Fujiwar and Zamir (1991) and Henrich et al. (2001), economists and anthropologists have joined together to conduct experiments in various countries. They found that, not only do people not play the subgame perfect Nash equilibrium, but there also appears to be systematic differences in the way people play the game across countries and cultures.

Bothelho et al. (2005) enter the debate by claiming that variance across countries does not necessarily imply a “cultural” or “national” effect. They correctly point out that these studies fail to control for the potential effects of demographic variables. In their paper, they report on two sets of experiments, one in the US (at the University of Southern California) and one in Russia (at the Moscow Institute of Electronic Technology). In these experiments, the authors collect information on the demographic characteristics of the players, and attempt to assess their influence on the bargaining process.

For our purposes, this data provides a perfect opportunity to use the logit bargaining model, to compare its results with other relevant research, and to compare some of the substantive implications of our approach with the results obtained by using traditional statistical analysis.

For illustrative purposes, we replicated the Bothelho et al. (2005) results for their least squares regression on offers, also with no interaction terms.¹⁵ Noting that the comparative static on $X\beta$ in the logit bargaining model is linear, we might suspect that using least squares regression would

¹⁵The reference to these results are found on page 374 of Bethelo et. al. (2005) and figure prominently in the working paper commonly cited from 2001.

OfferP	β	S.E.	t
Russia	-.08914	1.6632	-0.75
SlavS	3.13	.7623	4.107
ParInc 50-100kS	-1.1783	.6187	-1.905
Years Worked	-.0157	.0740	-0.212
fbusS	3.2226	.5821	5.536
HHInc 50-100k	5.0542	.6355	7.953
HHInc 100k+	5.9341	.6737	8.809

Table 3: Bargaining Model

OfferP	β	S.E.	t
Russia	-7.5605	3.4552	-2.01
SlavS	5.5880	3.1346	1.78
ParInc 50-100kS	-2.8885	2.4920	-1.16
Years Worked	-1.397	.504275	-2.77
fbusS	-2.6599	2.390862	-1.11
HHInc 50-100k	-2.2123	1.8711	-1.18
HHInc 100k+	-2.4205	1.7601	-1.38

Table 4: OLS Model

Note: The β coefficients for the bargaining model estimator are the affect of the variable in question on the reservation value of the player and is comparable to $(-\beta)$ in a linear regression of offers on player characteristics. Differences between the two modeling procedures are highlighted with bold.

produce the same results as the estimation of the bargaining model. It turns out, however that is not the case.

Consider the variables presented in Tables 3 and 4. These are variables selected from the total set of demographic variables (24 plus a constant) used for each player in the analysis. Cells in bold highlight differences between the two models. What becomes clear is that there are stark substantive differences between the effects inferred from the least squares regression and those of the bargaining model.

First, in the least squares regression we see that the so called “country” or cultural effect is both substantively and statistically significant. The result is interpreted as demonstrating that Russian participants tend to make smaller offers than Americans. The estimate from the bargaining model finds there is no such effect. Similarly, we find no statistically significant effect years worked. It is also the case that the bargaining model finds demographic variables that are not significant in the least squares regression, significant. In fact, the variables: parents’s income in the \$ 50-100k (ParInc 50-100kS), father’s occupation is business (fbusS), household income \$ 50-100k (HHInc 50-100k), and household income over \$ 100k (HHInc 100k+) are all significant in the bargaining model, but not significant in Table 4. Finally, the effect of being of Slavic culture, while marginally significant and positive in the Bothelho et al. (2005) study, in the bargaining model estimation it has a strongly significant effect, *in the other direction*. That is, people of Slavic decent make smaller offers, not larger ones.

While not a full fledged analysis of this bargaining data, the results form this comparison point to some potentially interesting conclusions. We see that in the experimental data, the structural bargaining model finds substantively different effects of demographic characteristics than one finds using a least squares approach. Where intuitively we may believe that parents's income in the \$ 50-100k range (ParInc 50-100kS), father's occupation is business (fbusS), household incomes of \$ 50-100k (HHInc 50-100k), and household income over \$ 100k (HHInc 100k+) would influence player's reservation values,i.e., their marginal utility of the experiment, if we believe the OLS estimation, they do not. The logit bargaining model, however, finds that these variables have the same intuitive directional effect as OLS, but also finds them to be statistically significant.

We also see that it may not be true that if the comparative statics of a model are linear, that an OLS model will accurately infer the effects of substantive variables of interest. As an exercise, suppose that the bargaining model were simplified and assumed a uniform distribution of types. Then, suppose we find the equilibrium offer. It is easy to see that such an offer would be linear in $X\beta$. Given the linearity of the relationship, we may think it appropriate to use a classical linear regression model to analyze the effect of variables in X on the offers. We see here, however, that it is not the case.

6 CONCLUSION

This manuscript derives a statistical estimator that can be used when the data generating process is best described as an equilibrium to an ultimatum bargaining game. The model is shown to have a number of nice properties. It allows the analyst to estimate, under the assumptions of the theory, players' utility functions and equilibrium quantities of interest—such as how equilibrium offers and the probability of bargaining failure react to changes in the independent variables. Monte Carlo experiments show that substantive inferences regarding the effect of variables will be different if this data generating process is ignored in the process of estimation. We also show that the statistical bargaining model allows the analyst to estimate the variance of the underlying logistic distribution of errors and that the estimator is well behaved in small samples.

Furthermore we can compare this model to other likelihood modes often used by analysts. First we see traditional logit models work well to answer some questions, particularly about bargaining failure, if properly specified. That is, in a take it or leave it setting the logit model is the right

structural model. Second, a criticism of structural estimation in political science has been that it is hard to deal with theories, or data generating processes, where players choose from more than two or three alternatives. The logit ultimatum model also shows that the structural estimation approach is not limited to games where players have finite action spaces. In particular, there are number of games — such as the Romer-Rosenthal setter model and the Rubenstein bargaining model — that may be estimable in similar ways. Finally, the broad interest in bargaining games leaves open a number of substantive areas where this estimator can be applied. From lab experiments to the study of territorial disputes, a simple bargaining framework underlies many theoretical arguments. The model described above allows analysts to obtain statistical estimates a step closer to theory and open up new possibilities for testing hard to operationalize hypotheses.

A ANALYTIC DERIVATIONS

First a useful definition.

Definition 1. A continuously differentiable function $f : R \rightarrow R^+$ is log-concave on an interval (a, b) if and only if $(\ln f(x))'' \leq 0$.

Also note the following useful fact from calculus.¹⁶

Fact 1. If a continuously differentiable function $f : R \rightarrow R^+$ is log-concave on an interval (a, b) , then $\frac{f'(x)}{f(x)}$ is a non-increasing function of $x \in (a, b)$

Proof. The function f is log-concave if and only if, for all $x \in (a, b)$,

$$(\ln f(x))'' = \frac{d}{dx} \frac{f'(x)}{f(x)} \leq 0. \tag{A-1}$$

□

Proposition 1. If F_{ϵ_2} is log-concave, then there exists a unique perfect Bayesian-Nash equilibrium to the statistical ultimatum game as described above.

Proof. The proof of this result is a straight forward construction of the equilibrium. It is a simple extension to types drawn from R of the result in Fearon (1995).

That player 2 rejects if and only if $y < R_2 + \epsilon_2$ is immediate from sequential rationality and, as f_2 is continuous $y = R_2 + \epsilon_2$ with probability zero, without loss of generality we can assume that player 2 accepts when she is indifferent between a settlement and disagreement. Then, in any equilibrium player 2 plays the cutpoint strategy:

$$s_2(y, \epsilon_2) = \begin{cases} \text{accept} & \text{if } y \geq R_2 + \epsilon_2 \\ \text{reject} & \text{if } y < R_2 + \epsilon_2. \end{cases}$$

Note also that the $\Pr(\text{accept}|y) = \Pr(y > R_2 + \epsilon_2) = \Pr(\epsilon_2 < y - R_2)$, and $\Pr(\epsilon_2 < y - R_2) \equiv F_{\epsilon_2}(y - R_2)$.

Now, assume F_{ϵ_2} is log-concave and consider the optimization problem for player 1, given player 2's strategy. His expected utility for an offer y is:

$$Eu_1(y, Q) = F_{\epsilon_2}(y - R_2) \cdot (Q - y) + (1 - F_{\epsilon_2}(y - R_2)) \cdot (R_1 + \epsilon_1). \tag{A-2}$$

Differentiating shows that $Eu_1(y)'$ is positive when

¹⁶For a paper on log-concave functions and their applications, see -Bagnoli and Bergstrom (2005).

$$0 < f_{\epsilon_2}(y - R_2)(Q - y) - F_{\epsilon_2}(y - R_2) - f_{\epsilon_2}(y - R_2)(R_1 + \epsilon_1),$$

which implies

$$\frac{f_{\epsilon_2}(y - R_2)}{F_{\epsilon_2}(y - R_2)} > \frac{1}{Q - y - R_1 - \epsilon_1}. \quad (\text{A-3})$$

By **Fact 1** the LHS is non-increasing and, by inspection, the RHS is strictly increasing in y . Now if equation (A-3) holds evaluated at $y = Q$, then

$$\frac{f_{\epsilon_2}(Q - R_2)}{F_{\epsilon_2}(Q - R_2)} > \frac{1}{-R_1 - \epsilon_1}. \quad (\text{A-4})$$

So, as we move from $y = Q$ to $y < Q$ the LHS is non-decreasing and the RHS is strictly decreasing. Therefore, the derivative of $Eu_1(y)$ is positive over the entire interval $[0, Q]$ and $y = Q$ is the optimal offer when

$$\epsilon_1 < -R_1 - \frac{F_{\epsilon_2}(Q - R_2)}{f_{\epsilon_2}(Q - R_2)}. \quad (\text{A-5})$$

Conversely, differentiation shows that $Eu_1(y)'$ is negative when

$$0 > f_{\epsilon_2}(y - R_2)(Q - y) - F_{\epsilon_2}(y - R_2) - f_{\epsilon_2}(y - R_2)(R_1 + \epsilon_1),$$

implying

$$\frac{f_{\epsilon_2}(y - R_2)}{F_{\epsilon_2}(y - R_2)} < \frac{1}{Q - y - R_1 - \epsilon_1}. \quad (\text{A-6})$$

Again, **Fact 1** implies LHS is non-increasing. Now if equation (A-6) holds evaluated at $y = 0$, then

$$\frac{f_{\epsilon_2}(-R_2)}{F_{\epsilon_2}(-R_2)} < \frac{1}{Q - R_1 - \epsilon_1}. \quad (\text{A-7})$$

So, as we move from $y = 0$ to $y > 0$ the LHS is non-increasing and the RHS is strictly increasing. Therefore, the derivative of $Eu_1(y)$ is negative over the entire interval $[0, Q]$ and $y = 0$ is the optimal offer when

$$\epsilon_1 > Q - R_1 - \frac{F_{\epsilon_2}(-R_2)}{f_{\epsilon_2}(-R_2)}. \quad (\text{A-8})$$

It is clear from an examination of (A-5) and (A-8) that some times neither equation is satisfied. To see this note that,

$$-R_1 - \frac{F_{\epsilon_2}(Q - R_2)}{f_{\epsilon_2}(Q - R_2)} \leq Q - R_1 - \frac{F_{\epsilon_2}(-R_2)}{f_{\epsilon_2}(-R_2)}. \quad (\text{A-9})$$

Multiplying through by -1 and taking the inverse of each side, we get

$$\frac{f_{\epsilon_2}(Q - R_2)}{F_{\epsilon_2}(Q - R_2)} \leq \frac{f_{\epsilon_2}(-R_2)}{F_{\epsilon_2}(-R_2) - Qf_{\epsilon_2}(-R_2)}. \quad (\text{A-10})$$

At $Q = 0$ the LHS and RHS are equal. However as Q increases the LHS is non-increasing by **Fact 1** and the RHS is strictly increasing. Thus there are always ϵ_1 such that (A-5) and (A-8) cannot hold, given our assumption that $Q > 0$.

When neither (A-5) nor (A-8) hold, then for some (possibly multiple) $y \in [0, Q]$, the derivative of $Eu_1(y)$ is zero, implying

$$\frac{f_{\epsilon_2}(y - R_2)}{F_{\epsilon_2}(y - R_2)} = \frac{1}{Q - y - R_1 - \epsilon_1}. \quad (\text{A-11})$$

Since the LHS is non-increasing on $[0, Q]$ and the RHS is strictly increasing on the same interval, equation (A-11) can have at most one solution. Call this offer y^* and note that it implicitly solves,

$$y^* = Q - R_1 - \epsilon_1 - \frac{F_{\epsilon_2}(y^* - R_2)}{f_{\epsilon_2}(y^* - R_2)}, \quad (\text{A-12})$$

We now demonstrate that y^* maximizes player 1's expected utility. Obviously, since $Eu_1(y)$ is continuous for every ϵ_1 on the interval, the utility maximizing offer exists and must be a critical point, like y^* or a boundary point. If neither (A-5) nor (A-8) hold, then there are two cases. First, if one of the end points is the unique solution to equation (A-11) we are done. Second, if y^* is interior the derivative of $Eu_1(y)$ at $y = 0$ is positive and $Eu_1(y)'$ at $y = Q$ is negative by (A-5) and (A-8). Thus the interior critical point is a local and global maximum. □

B COMPARATIVE STATICS FOR THE OPTIMAL OFFER

Recall from Equation 2 that player 1's unconstrained optimal offer is

$$y^* = Q - R_1 - \epsilon_1 - \frac{F_{\epsilon_2}(y^* - R_2)}{f_{\epsilon_2}(y^* - R_2)}.$$

We are interested in how changes in observable reservation components R_1 and R_2 change the optimal offer. First, let us specify the reservation utilities with regressors x , z , and v as follows:

$$R_1 = \beta_x x + \beta_v v$$

$$R_2 = \gamma_z z + \gamma_v v$$

This specification allows us to characterize the comparative statics for elements that are unique to each player (e.g., x and z), as well as those that are shared (e.g., v). Next, define

$$m(x) = \frac{F_{\epsilon_2}(x)}{f_{\epsilon_2}(x)}.$$

Case 1: $\beta_v = \gamma_v = 0$. We first consider the case where R_1 and R_2 share no common regressors. By the Implicit Function Theorem we can express the first derivatives of y^* as

$$\frac{dy^*}{dx} = \left[\frac{1}{1 + m'(y^* - z\gamma_z)} \right] (-\beta_x) \quad (\text{A-13})$$

$$\frac{dy^*}{dz} = \left[\frac{m'(y^* - z\gamma_z)}{1 + m'(y^* - z\gamma_z)} \right] \gamma_z \quad (\text{A-14})$$

Because $m'(x) > 0$ by log-concavity, it follows that

$$\text{sign}(dy^*/dx) = \text{sign}(-\beta_x)$$

$$\text{sign}(dy^*/dz) = \text{sign}(\gamma_z)$$

Thus, the unconstrained optimal offer y^* is unconditionally monotone in x and in z . For $\beta_x > 0$, the unconstrained optimal offer decreases monotonically in x . For $\gamma_x > 0$, the unconstrained optimal offer increases monotonically in z . Because the constrained optimal offer has a floor (zero) and a ceiling (Q), the constrained optimal offer is weakly monotone in x and in z .

Case 2: $\beta_v \neq 0, \gamma_v \neq 0$. We next consider the case where R_1 and R_2 share a common regressor v . It is easy to check that $\text{sign}(dy^*/dx) = \text{sign}(-\beta_x)$ and $\text{sign}(dy^*/dz) = \text{sign}(\gamma_z)$ as before. However, the derivative with respect to v is

$$\frac{dy^*}{dv} = \left[\frac{1}{1 + m'(y^* - z\gamma_z - v\gamma_v)} \right] (-\beta_v) + \left[\frac{m'(y^* - z\gamma_z - v\gamma_v)}{1 + m'(y^* - z\gamma_z - v\gamma_v)} \right] \gamma_v \quad (\text{A-15})$$

As Equation A-15 shows, v has countervailing effects on the optimal offer, due to the fact that it appears in both players' reservation values. Because the denominator is positive in both terms on the RHS, it follows that $\text{sign}(dy^*/dv) = \text{sign}[\gamma_v m'(y^* - z\gamma_z - v\gamma_v) - \beta_v]$. Therefore, the optimal offer increases in v when $\gamma_v m'(y^* - z\gamma_z - v\gamma_v) > \beta_v$ and decreases in v when $\gamma_v m'(y^* - z\gamma_z - v\gamma_v) < \beta_v$. Notice, however, that $m'(y^* - z\gamma_z - v\gamma_v)$ changes with v . Therefore, we are not assured of monotonicity in v unless we can show the sign remains constant for all admissible values of v — and, indeed, this will not always be the case as one can construct counter-examples when the distribution of errors is logistic. \square

C DERIVATION OF $f_y(y)$

C.1 Method of Transformation

We assume the distribution of ϵ_1 is logistic with scale parameter s_1 . Similarly, the distribution of ϵ_2 is logistic with scale parameter s_2 . We have confirmed that the derivative of y^* is single signed (< 0), so we can apply the method of monotonic transformation to obtain the distribution of y^* .

That is, with $y^* = h(\epsilon_1)$ and $\epsilon_1 = h^{-1}(y^*)$:

$$f_{y^*}(y^*) = f_{\epsilon_1}(h^{-1}(y^*)) \left| \frac{d(h^{-1}(y^*))}{dy^*} \right|. \quad (\text{A-16})$$

Solving Equation 5 for ϵ_1 produces

$$\epsilon_1 = h^{-1}(y^*) = Q - y^* - X\beta - s_2 \left[1 + e^{(y^* - Z\gamma)/s_2} \right] \quad (\text{A-17})$$

Taking the derivative of $h^{-1}(y^*)$ with respect to y^* gives

$$\frac{d(h^{-1}(y^*))}{dy^*} = - \left[1 + e^{(y^* - Z\gamma)/s_2} \right]. \quad (\text{A-18})$$

Substituting this into A-16 yields

$$f_{y^*}(y^*) = \frac{e^{-[h^{-1}(y^*)]/s_1}}{s_1 \left\{ 1 + e^{-[h^{-1}(y^*)]/s_1} \right\}^2} \cdot \left[1 + e^{(y^* - Z\gamma)/s_2} \right] \quad (\text{A-19})$$

$$= \frac{e^{-[Q - y^* - X\beta - s_2(1 + e^{(y^* - Z\gamma)/s_2})]/s_1}}{s_1 \left\{ 1 + e^{-[Q - y^* - X\beta - s_2(1 + e^{(y^* - Z\gamma)/s_2})]/s_1} \right\}^2} \cdot \left[1 + e^{(y^* - Z\gamma)/s_2} \right] \quad (\text{A-20})$$

□

REFERENCES

- Amemiya, Takeshi. 1984. "Tobit Models: A Survey." *Journal of Econometrics* 24:3–61.
- Bagnoli, Mark and Ted Bergstrom. 2005. "Log-concave Probability and its Applications." *Economic Theory* 26:445–469.
- Banks, Jeffrey S. 1990. "Equilibrium Behavior in Crisis Bargaining Games." *American Journal of Political Science* 34(3):579–614.
- Baron, David P. 1989. "A Noncooperative Theory of Legislative Coalitions." *American Political Science Review* 33(4):1181–1206.
- Baron, David P. and John Ferejohn. 1989. "Bargaining in Legislatures." *American Political Science Review* 33(4):1048–1084.
- Barray, D. A. and P. J. Culligan-Hensley. 1995. "Real Values of the W-Function." *ACM Transactions on Mathematical Software* 21(2):161–171.
- Beja, Avraham. 1992. "Imperfect Equilibrium." *Games and Economic Behavior* 4:18–36.
- Bennett, D. Scott. 1996. "Security, Bargaining, and the End of Interstate Rivalry." *International Studies Quarterly* 40(2):157–183.
- Bothelho, Anabela, Glenn W. Harrison, Marc A. Hirsch and Elisabeth E. Ruström. 2005. "Bargaining Behavior, Demographics and Nationality: What Can Experimental Evidence Show?" *Research in Experimental Economics* 10(4):337–372.
- Casella, Goerge and Roger L. Berger. 2002. *Statistical Inference*. Pacific Grove, CA: Duxbury.
- Corless, R. M., G. T. Gonnet, D.E. G. Hare, D.J. Jeffrey and D. E. Knuth. 1996. "On the Lambert W Function." *Advances in Computational Mathematics* 5:329–359.
- Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(3):379–414.
- Fey, Mark and Kristopher W. Ramsay. 2009. "Uncertainty about Relative Power and War." *Princeton University*.

- Fitsch, F. N., R.E. Shafer and W. P. Crowley. 1973. "Algorithm 443: Solution of the Transcendental Equation $we^w = x$ [C5]." *Communications of the AMC* 16(2):123–124.
- Henrich, Joseph, Robert Boyd, Samuel Bowles, Clin Camerer, Ernst Fehr, Herbert Gintis and Richard McElreath. 2001. "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." *American Economic Review* 91(2):73–78.
- Huth, Paul K. and Todd L. Allee. 2002. *The Democratic Peace and Territorial Conflict in the Twentieth Century*. New York, NY: Cambridge University Press.
- Kedziora, Jeremy, Kristopher W. Ramsay and Curtis S. Signorino. 2009. "Bargaining Models with Partially Observed Covariates." *University of Rochester* .
- King, Gary, Robert Keohane and Sidney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton, NJ: Princeton University Press.
- Laver, Michael and Norman Schofield. 1990. *Multiparty Government*. Ann Arbor, MI: University of Michigan Press.
- London, Tamar R. 2002. Leaders, Legislatures, and International Negotiations: A Two-Level Game with Different Domestic Conditions. Ph.d thesis University of Rochester.
- Luce, Robert D. 1959. *Individual Choice Behavior*. New York, NY: Wiley.
- Mansfield, Edward D., Helen V. Milner and B. Peter Rosendorff. 2000. "Tree to Trade: Democracies, Autocracies, and International Trade." *American Political Science Review* 94:305–321.
- McCarty, Nolan and Keith Poole. 1995. "Veto Power and Legislation: An Empirical Analysis of Executive and Legislative Bargaining from 1961 to 1986." *Journal of Law, Economics, and Organization* 11:282–312.
- McFadden, Daniel. 1976. "The Revealed Preferences of a Government Bureaucracy: Empirical Evidence." *Bell Journal of Economics* 7(1):55–72.
- McKelvey, Richard D. and Thomas R. Palfrey. 1996. "A Statistical Theory of Equilibrium in Games." *The Japanese Economic Review* 47(2):186–209.
- Meirowitz, Adam. 2003. "On the existence of equilibria to Bayesian games with non-finite type and action spaces." *Economics Letters* 78:213–218.

- Merlo, Antonio. 1997. "Bargaining over Governments in a Stochastic Environment." *Journal of Political Economy* 105(1):101–131.
- Merlo, Antonio and Charles Wilson. 1995. "A Stochastic Model of Sequential Bargaining with Complete Information." *Econometrica* 63(2):371–399.
- Merlo, Antonio and Charles Wilson. 1998. "Efficient delays in a stochastic model of bargaining." *Economic Theory* 11:39–55.
- Milner, Helen V. 1997. *Interest, Institutions, and Information: Domestic Politics and International Relations*. Princeton, NJ: Princeton University Press.
- Morrow, James D. 1989. "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining." *American Journal of Political Science* 33(4):941–972.
- Powell, Robert. 1987. "Crisis Bargaining, Escalation, and MAD." *American Political Science Review* 81(3):717–736.
- Powell, Robert. 1996. "Bargaining in the Shadow of Power." *Games and Economic Behavior* 15:255–289.
- Roth, Alvin E., Venesa Prasnikar, Msahiro Okuno-Fujiwar and Shmuel Zamir. 1991. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *American Economic Review* 81(5):10068–1095.
- Signorino, Curtis. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review* 93(2):279–297.
- Signorino, Curtis. 2002. "Strategy and Selection in International Relations." *International Interactions* 28(1):93–115.
- Signorino, Curtis. 2003. "Structure and Uncertainty in Discrete Choice Models." *Political Analysis* 11(4).
- Signorino, Curtis S. and Kuzey Yilmaz. 2003. "Strategic Misspecification in Discrete Choice Models." *American Journal of Political Science* 47(3).
- Valluri, S. R., D. J. Jeffrey and R. M. Corless. 2000. "Some Applications of the Lambert W Function to Physics." *Canadian Journal of Physics* 78:823–831.

- van Damme, Eric. 1991. *Stability and Perfection of Nash Equilibrium*. New York, NY: Springer-Verlag.
- Wagner, R. Harrison. 2000. "Bargaining and War." *American Journal of Political Science* 44(3):469–484.
- Werner, Suzane. 1999. "Choosing Demands Strategically: The Distribution of Power, the Distribution of Benefits, and the Risk of Conflict." *Journal of Conflict Resolution* 43(6):705–726.
- Wolpin, Kenneth I. 1987. "Estimating A Structural Search Model: The Transition from School to Work." *Econometrica* 55(4):801–817.