Structural Approach to Bargaining and EITM

Kristopher W. Ramsay
Princeton University

Structural vs. Reduced-form approaches to empirical social science

Good points from each side

- If all I care about is the total effect of X on Y and if

$$Y = E[Y|X] + \mu = F(X\beta) + \mu$$

then regression of Y on X gives a linear approximation of

$$\frac{\partial E[Y|X]}{\partial X} \text{ if } E[\mu_i|X_i] = 0$$

- Probit (plus mfx) command also gives a linear approximation (the derivative at the global mean of $\Phi(X\beta)$) of $\frac{\partial E[Y|X]}{\partial X}$ under **normal** spherical disturbances.

Structural Estimation: Arguments and Lessons

- Motivation: “Likelihood turn” in political methodology, or why structural models followed duration models?
  - ” If the data generating process is a ———, then you need a ——— model.”
- Why would you want estimates of βs in utility functions?
  - Forecasting
  - Better than linear approximation
  - Testing auxiliary hypotheses
  - We know the correct model (experiments)
- Side note: one lesson from structural models–It is a long way from TM to EI.

A Statistical Model of the Ultimatum Game

Kristopher W. Ramsay
Princeton University

Curtis S. Signorino
University of Rochester

Introduction

- **Game theoretic** bargaining models applied to
  - the the distributive politics of legislative appropriation
  - the study of war initiation and termination
  - international political economy
- Economists and anthropologists amassing quite a bit of data on how individuals "play" in bargaining experiments.
  - Are people "Nash-ty"
  - Can deviations from **egoistic** behavior and Nash play be explained systematically?
- **Quantitative analysis** of bargaining has progressed at a much slower rate – even in economics. Typically, scholars analyzing bargaining offers employ OLS, FGLS, or Tobit.

Project Objectives

- Create statistical techniques for analyzing ultimatum bargaining games.

- Use these models to analyze existing data on the following:

  1. International Relations
  2. American Politics
  3. Society and Rationality

Roadmap

- A statistical equilibrium model of the Ultimatum game.
- A logit Ultimatum estimator.
- Monte carlo analysis demonstrates
  1. that our estimator correctly recovers the parameters of a bargaining process
  2. that OLS and Tobit can yield incorrect inferences when applied to bargaining data
- Preliminary application of our statistical model to data on ultimatum experiments conducted on Russian and U.S. individuals.
- Concluding remarks

The Ultimatum Game

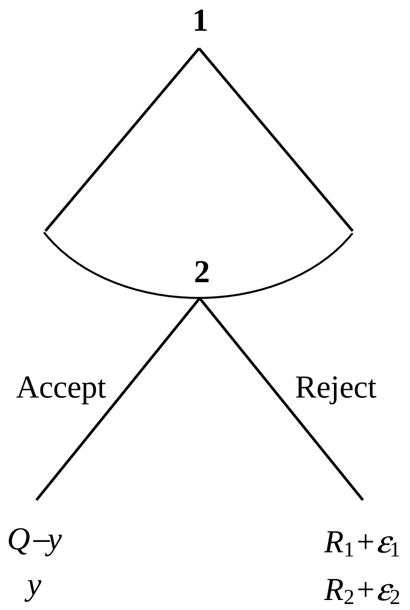

- Two players divide a prize, $\mathcal{Q}^*$.
- Player 1 offers $y \in [0, \mathcal{Q}^*]$ to player 2.
- Player 2 decides whether to accept.
- If player 2 rejects the offer, they receive some reservation amount, which may differ between the players.

- Utilities for bargaining failure have two components:
  - Public: R_i
  - Private: ϵ_i

Equilibrium

The equilibrium can be found by working "up" the tree.

Player 2

Player 2 chooses between offer y and her reservation value $R_2 + \epsilon_2$. In equilibrium, player 2 chooses

$$\begin{array}{ll} accept & \text{if } y \geq R_2 + \epsilon_2 \\ reject & \text{if } y < R_2 + \epsilon_2. \end{array}$$

Player 1

Player 1 does not observe ϵ_2, but must assess the probability that player 2 will accept or reject his offer:

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

Consider the optimization problem for player 1, given player 2's strategy. His expected utility is

$$Eu_1(y|\mathcal{Q}^*) = F_\epsilon(y - R_2) \cdot (\mathcal{Q}^* - y) + [1 - F_\epsilon(y - R_2)] \cdot (R_1 + \epsilon_1),$$

With no constraints, his optimal offer (y^*) is the solution to

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

However, $0 \leq y \leq \mathcal{Q}^*$. Therefore, given the constraints, player 1 offers

$$y = \begin{cases} 0 & \text{if } \epsilon_1 \geq \mathcal{Q}^* - \xi - R_1 \\ y^* & \text{if } -\xi - R_1 < \epsilon_1 < \mathcal{Q}^* - \xi - R_1 \\ \mathcal{Q}^* & \text{if } \epsilon_1 \leq -\xi - R_1 \end{cases} \tag{3}$$

where $\xi = \frac{F_\epsilon(y^* - R_2)}{f_\epsilon(y^* - R_2)}$.

A Logit Ultimatum Model: Home brewed likelihood

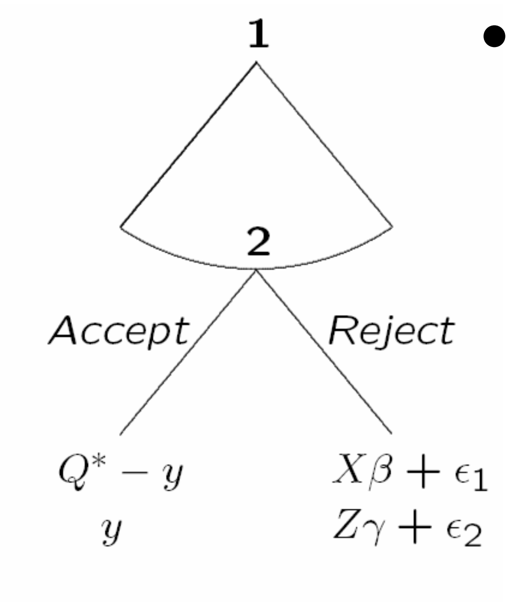

- ϵ_i i.i.d. logistic
- What affects players' decisions?
  - Public reservation values
  - Scale term in private component
  - Parts of private component observable to analyst but not other players (e.g., in experiments)

- Here: Will focus on estimating effects of substantive regressors X and Z on players' decisions

Player 2

The probability that player 2 accepts the offer y is just the logit probability

$$\Pr(accept|y) = \Lambda(y - Z\gamma)$$

Player 1

Distribution of y^* is more complicated. A logistic distribution of types implies that

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

where $\Lambda(\cdot)$ is the logit c.d.f. and $\lambda(\cdot)$ is the logit p.d.f.

Simplifies to

$$y^* = \mathcal{Q}^* - X\beta - \epsilon_1 - \left(1 + e^{y^* - Z\gamma}\right) \tag{5}$$

Solving for the unique optimal offer y^* gives

$$y^* = \mathcal{Q}^* - X\beta - \epsilon_1 - 1 - \mathcal{W}\left(e^{\mathcal{Q}^* - X\beta - Z\gamma - \epsilon_1 - 1}\right) \tag{6}$$

Given the nice properties of $\mathcal{W}$ we can apply the method of transformations and then the density $f_{y^*}(y^*)$ is

$$f_{y^*}(y^*) = \frac{\left(e^{\mathcal{Q}^* - 1 - X\beta - \exp(y^* - Z\gamma) - y^*}\right) \cdot \left(1 + e^{y^* - Z\gamma}\right)}{\left(1 + e^{\mathcal{Q}^* - 1 - X\beta - \exp(y^* - Z\gamma) - y^*}\right)^2} \tag{7}$$

and

$$F_{y^*}(y^*) = \left[1 + e^{\mathcal{Q}^* - 1 - X\beta - \exp(y^* - Z\gamma) - y^*}\right]^{-1} \tag{8}$$

y^* is the optimal offer player 1 would like to make.

However, the offer must be between 0 and $\mathcal{Q}^*$.

Let y be the observed offer.

$$y = \begin{cases} 0 & \text{if } \; y^* \leq 0 \\ y^* & \text{if } \; 0 < y^* < \mathcal{Q}^* \\ \mathcal{Q}^* & \text{if } \; y^* \geq \mathcal{Q}^* \end{cases}$$

Define a set of dummy variables

$$\begin{aligned} \delta_0 &= 1 \quad \text{if } \; y = 0 \\ \delta_y &= 1 \quad \text{if } \; 0 < y < \mathcal{Q}^* \\ \delta_1 &= 1 \quad \text{if } \; y = \mathcal{Q}^* \end{aligned}$$

Finally, code player 2's acceptance as $\delta_{accept} = 1$ if she accepted the offer and $\delta_{accept} = 0$ if she rejected the offer.

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

$$L = \prod_{i=1}^{n} \Big[\quad \Pr(y^* < 0)^{\delta_0} \cdot \Pr(y^* = y)^{\delta_y} \cdot \Pr(y^* > \mathcal{Q}^*)^{\delta_1} \cdot \\ \Pr(accept)^{\delta_{accept}} \cdot \Pr(reject)^{1-\delta_{accept}} \Big]$$

- Log-likelihood function for our data in terms of distributions already derived, which are functions of our regressors. Explicitly models the Ultimatum game.

- Estimates of β and γ may be obtained using maximum likelihood estimation.

Monte Carlo Analysis

- Demonstrate that (on average) the Ultimatum estimator correctly recovers the parameters of the bargaining model as a data generating process.
- Analyze inferences from other (commonly used) estimators such as OLS and Tobit.

Monte Carlo Analysis

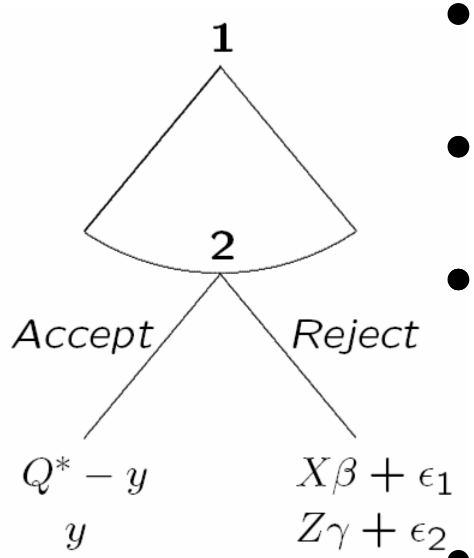

- Ultimatum game used as DGP.
- Value of the pie: $\mathcal{Q}^* \sim U(0, 10)$
- Public reservation values: $X\beta$ and $Z\gamma$
  $X, Z \sim U(0, \text{Max R})$
  $\beta = \gamma = 1$
- Private information: $\epsilon_1, \epsilon_2 \sim$ logistic

- Based on this, player 1 determines his optimal offer y using Equation 6 and the constraint $0 \leq y \leq \mathcal{Q}^*$. Given that, player 2 makes a decision by comparing y to her reservation $Z\gamma + \epsilon_2$.

- Data for a given observation:
  $\mathcal{Q}^*$, y, δ_{accept}, X, and Z.

- N=1000 observations for each MC iteration. Then estimated

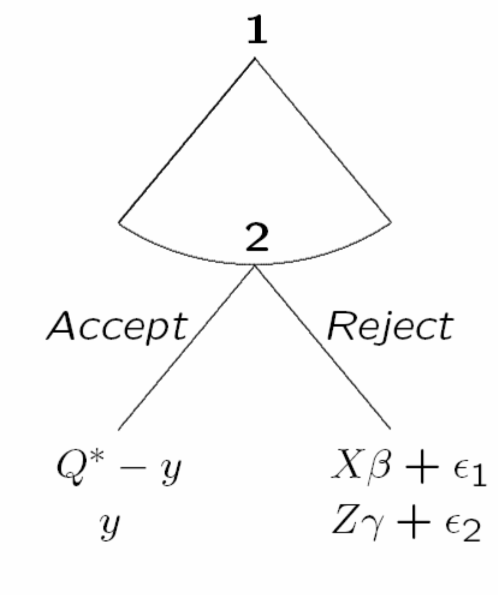

  - Logit Ultimatum model
    y, δ_{accept}

  - Traditional logit model
    $\delta^*_{accept} = y\omega + Z\gamma$

  - Normal model
    $y = Const + X\beta + Z\gamma + \mathcal{Q}^*\theta$

  - Tobit model, censored at 0 and $\mathcal{Q}^*$
    $y = Const + X\beta + Z\gamma$

- Estimates saved each iteration. Repeated 2000 iterations to form a density of the estimates.

- MCs conducted for Max R $\in \{2, 4, 6, 8, 10\}$.

Player 2's Acceptance: Ultimatum vs Logit

Max R	Number Censored	Ultimatum $\widehat{\beta}_u$	Ultimatum $\widehat{\gamma}_u$	Logit $\widehat{\omega}_l$	Logit $\widehat{\gamma}_l$
2	253	.99	.99	1.00	-1.00
		.06	.02	.07	.10
10	593	1.00	1.00	1.01	-1.01
		.02	.02	.09	.08

- Ultimatum model recovers estimates
- Strategic censoring (due to Max R) does not affect estimates
- Similar results for logit based solely on player 2's choice.
- Effect of Z is estimated more precisely with Ultimatum model.

Player 1's Offer: Ultimatum, Tobit, Normal

Max R	Num Cens	Tobit				Normal				
		$Const$	$\widehat{\beta}_t$	$\widehat{\gamma}_t$	$\widehat{\sigma}_t^2$	$Const$	$\widehat{\beta}_n$	$\widehat{\gamma}_n$	$\widehat{\theta}_n$	$\widehat{\sigma}_n^2$
2	253	.92	$-.35$	.56	1.69	$-.42$	$-.25$	.47	.29	.29
		.11	.07	.07	.08	.05	.03	.03	.01	.02
10	593	2.29	-0.78	.16	7.34	.53	$-.30$	.09	.31	1.33
		0.23	.04	.04	.51	.09	.01	.01	.01	.07

- Both Models
  - Estimates change with amount of censoring
  - Standard errors small

$E(Y|X,Z)$

(a) Z=0

(b) Z=10

(c) X=0

(d) X=10

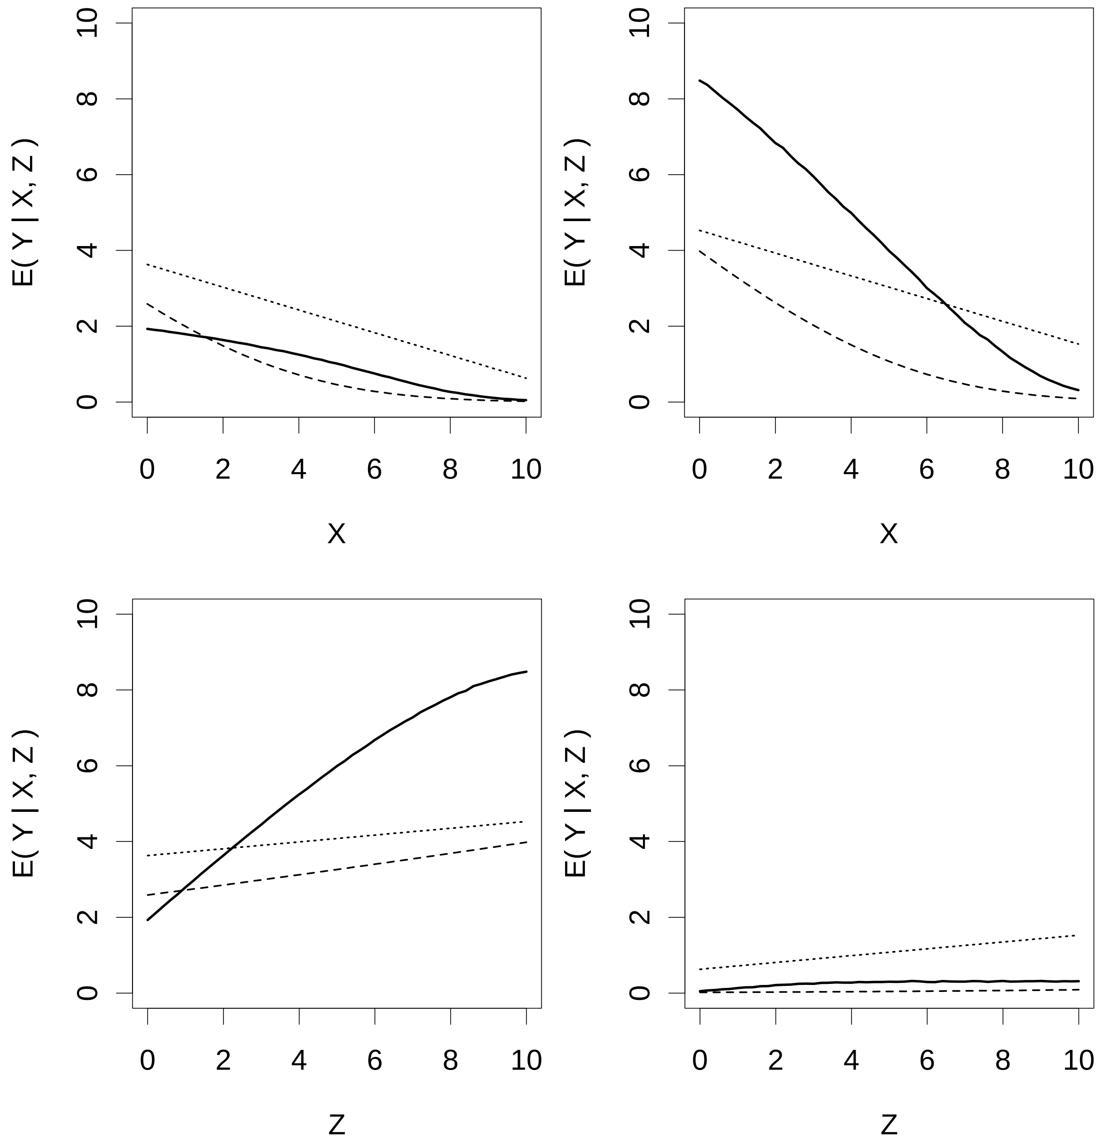

Estimator with Scale Parameters: Who says there is no free lunch!

Recall

$$y^* = Q^* - X\beta - \epsilon_1 - \frac{\Lambda(y^* - Z\gamma)}{\lambda(y^* - Z\gamma)} \tag{9}$$

where Λ is the logistic cdf

$$\Lambda(y^* - Z\gamma) = \frac{1}{1 + e^{-(y^* - Z\gamma)/s_2}} \tag{10}$$

Let s_1 and s_2 be the scale parameters in the logistic distributions for ϵ_1 and ϵ_2, respectively.

Then

$$f_{y^*}(y^*) = \frac{e^{-(Q^* - X\beta - s_2 - s_2(e^{(y^* - Z\gamma)/s_2}) - y^*)/s_1}}{s_1(1 + e^{-(Q^* - X\beta - s_2 - s_2(e^{(y^* - Z\gamma)/s_2}) - y^*)/s_1})^2} \cdot (1 + e^{(y^* - Z\gamma)/s_2}) \tag{11}$$

Can Recover Scale Parameters: s_1 and s_2

	β	γ	s_1	s_2
Parameter estimates	1.00	1.00	3.00	1.99
Standard error	.04	.03	.11	.09

Monte Carlo

- $\beta = \gamma = 1$, $s_1 = 3$, $s_2 = 2$
- N=2000 observation for each iteration
- 2000 iterations

Bargaining and Society

Many experiments conducted based on Ultimatum game:

- Botelho, Harrison, Hirsch, & Rutstrom
  — US, Russia
- Roth, Prasnikar, Okuno-Fujiwara, & Zamir
  — Japan, Israel, Yugoslavia, US
- Slonim & Roth
  — Slovak Republic
- Cameron
  — Indonesia
- Henrich et al. (Small-Scale Societies)
  — Machiquenga, Quichua, Achuar, Hadza, Ache, Tsimane, Au, Gnau, Mapuche, Torguuds, Kazakhs, Sangu, Orma, Lamalera, Shona

Bargaining and Society

- Observation: people in ultimatum experiments do not "play" the game's (complete information) Nash equilibrium.

- Experiments tend to show that proposers give receivers **larger** shares of the pie than is predicted by strict income maximizing behavior.

- Previous claim that there are systematic differences in results across countries and cultures. Different social and **cultural** groups deviate in systematically different ways.

Botelho, Harrison, Hirsch, & Rutstrom (BHHR)

- BHHR enter the debate by claiming that variance across countries does not necessarily imply a "cultural" or "national" effect.

- Previous studies failed to control for the potential confounding effects of **demographic** variables.

- To examine the cultural hypotheses, they conduct a series of experiments where they collect information on the participants' demographic characteristics.

- Subjects are students at US and Russian universities. They analyze the data using OLS, Logit, and variants thereof.

BHHR Experiment

- Setup
  - 60 subjects in Russia and then 60 subjects in US
  - Half designated "Sellers" and half "Buyers"
  - One practice round and then five rounds played.
  - Sellers remain sellers throughout. Buyers remain buyers.
  - Sellers and buyers randomly paired. Anonymous.
  - Buyer makes offer for some fictitious object being sold. Seller can accept or reject it.
  - At the end of five rounds, each receives randomly chosen payoff from their five rounds.
- Collect data on experiment play, plus demographics

Comparison of OLS and Bargaining Models

- Dependent variable: y, the offer made by player 1
- Original OLS regression:

  Positive β: increase in $X \rightarrow$ increase in the offer y.
- Logit divide-the-dollar model: the $X\beta$s are the reservation utilities, and are negatively related to y.

  Positive β: increase in $X \rightarrow$ increase in $R \rightarrow$ a decrease in y.
  Negative β: increase in $X \rightarrow$ decrease in $R \rightarrow$ an increase in y.

OLS Model	$\widehat{\beta}$	S.E.	t
Russia	-7.56	3.45	**-2.01**
Slav	**5.58**	3.13	1.78
Parents' Income 50-100k	**-2.88**	2.49	**-1.16**
Years Worked	**-1.39**	.504	**-2.77**
Father Businessman	-2.65	2.39	**-1.11**
Household Income 50-100k	-2.21	1.87	**-1.18**
Household Income 100k+	-2.42	1.76	**-1.38**

Bargaining Model	$\widehat{\beta}$	S.E.	t
Russia	-.089	1.66	**-0.75**
Slav	**3.13**	.76	4.10
Parents' Income 50-100k	**-1.17**	.61	**-1.90**
Years Worked	**-.015**	.074	**-0.21**
Father Businessman	3.22	.58	**5.53**
Household Income 50-100k	5.05	.63	**7.95**
Household Income 100k+	5.93	.67	**8.80**

BHHR Experiment

- Results
  - Somewhat limited
  - Some demographic effects: Male, Russia
  - "Adjustments" over rounds

- ISSUES:
  - Use of OLS/FGLS to analyze offers
  - Adjustment overly simplistic
  - Are the subjects really playing an Ultimatum game?
    - "Sales Game" with unknown value for object

Fun with Variances–Player 1: Egoistic vs Fair Division

Is it possible to observe the offers under traditional egoistic play? Or is it more consistent with risk averseness or "fair division"?

- Game 1: Close to True Ultimatum: R1=90, R2=10

```
       Estimate Std. Error t value p value
lns1      5.309      0.124   42.94       0
lns12     2.311      0.047   49.17       0
llik =  -1073.647
```

- Game 2: Fair Division: R1=50, R2=50

```
       Estimate Std. Error t value p value
lns1      2.255      0.055  41.172       0
lns12     2.175      0.078  27.731       0
llik =  -919.873
```

- Game 3: Estimate R1, R2, lns1, lns2

	Estimate	Std. Error	t value	p value
R1	56.627	1.594	35.522	0
R2	46.227	0.983	47.048	0
lns1	2.315	0.081	28.664	0
lns12	1.388	0.126	11.003	0

llik = -906.4914

Pretty close to 50-50 game.

Does Player 1's beliefs about Player 2 actually match Player 2's play?

- Player 2: Accept/Reject, Estimate R2 and lns2

	Estimate	Std. Error	t value	p value
R2	30.754	1.561	19.699	0
lns2	2.210	0.137	16.143	0

llik = -118.9347

Recap

- A statistical equilibrium model of the Ultimatum game.
- A logit Ultimatum estimator.
- Monte carlo analysis demonstrates
  1. that our estimator correctly recovers the parameters of a bargaining process
  2. that OLS and Tobit can yield incorrect inferences when applied to bargaining data
- Preliminary application of our statistical model to data on ultimatum experiments conducted on Russian and U.S. individuals.

What did we do here? Is it what you think EITM is about?

- Assumed data generating process
- Estimated parameters of model
- "Tested" claims about ultimatum games

Can we do more?

Yes, non-nested model tests.

Nonnested Model Testing

- We want to compare:
  - Strategic Model vs Selection Model
  - Binary Strategic Model vs Probit

- The models are nonnested in terms of their functional forms.

- Determining which of these functional forms is closest to the true, but unknown, specification requires the use of discrimination tests that are still new to most political scientists.

- Two of the easiest of these tests are
  - Vuong test (Vuong 1989)
  - Distribution-free test by Clarke (2003)

- Both of these tests are based on the Kullback-Leibler information criteria (KLIC) (Kullback and Leibler 1951).

  Both based on comparing whether the log-likelihood of one model is significantly larger than the log-likelihood of the rival model.

The Distribution-Free Test

- The Vuong test is not an exact test.
  - It is normally distributed **asymptotically**.
  - Simulations demonstrate that it is not very powerful in small samples (Clarke 2003).

- Clarke's (2003) distribution-free test applies a modified paired sign test to the differences in the individual log-likelihoods from two nonnested models.

- Let $Z_i = \ln f(Y_i|X_{fi}; \widehat{\beta}) - \ln g(Y_i|X_{gi}; \widehat{\gamma})$

- The null hypothesis of the distribution-free test is:

$$H_0 : \Pr(Z_i > 0) = .5$$
$$\Pr(Z_i < 0) = .5$$

- Under the null, half the individual log-likelihood ratios should be greater than zero and half should be less than zero.

  If model f is "better" than model g, then more individual log-likelihood ratios should be greater than zero.

  If model g is "better" than model f, then more individual log-likelihood ratios should be less than zero.

The Distribution-Free Test

- Let

$$\Psi_i = \begin{cases} 1 & \text{if } Z_i > 0 \\ 0 & \text{if } Z_i < 0, \end{cases}$$

then the test statistic is

$$B = \sum_{i=1}^{n} \Psi_i$$

The number of observations where $\ln f_i$ is greater than $\ln g_i$

- Under the null, B is distributed Binomial$(n, .5)$

- This test, like the Vuong test, may be affected if the number of coefficients in the two models being estimated is different.

  Once again, we need a correction for the degrees of freedom. The Schwarz correction is:

$$\left[\left(\frac{p}{2n}\right)\ln n - \left(\frac{q}{2n}\right)\ln n\right]$$

Conclusion

- There is still controversy over how to do empirical political science and EITM in particular.
- To some extent, reduced form analysis makes sense when looking for approximations of total effects or where theory is less well developed.
- Straight-up structural modeling is just like the other sorts of "getting the right estimator" approaches and is a natural extension.
- Limits of structural models will be discussed more later (updating anyone?).
- However, the technology exists to do the thing you were taught in scope and methods. (Theoretically informed comparison of models of the data generating process.)