Bargaining Under Uncertainty: a Strategic Statistical Model of the Ultimatum Game

Stephen R. Haptonstahl

As of January 22, 2009

Abstract

The Ultimatum game is commonly applied in describing political interactions such as negotiations between Congress and the president. Experimental and observational results differ from those predicted: proposers give too much, and responders sometimes reject offers. Recent work suggests introducing either of two different sources of uncertainty – about actions or payoffs – to solve the "zero-likelihood" problem and permit direct statistical comparison of substantive explanations for these deviations. Using a combined model and implementing each of these sources of uncertainty, we examine original data from a laboratory experiment designed to differentiate between these sources of uncertainty. The data suggest that these two sources of uncertainty are observationally quite distinct.

1 Introduction

Bargaining games such as the Ultimatum and Principal-Agent games are commonly applied[1] in describing political interactions such as negotiations between Congress and the president

---

[1] Ferejohn and Shipan (1990), Neustadt (1990, ch. 4), Miller (1992); Brehm and Gates (1997); Waterman, Rouse, and Wright (1998); Wilson (2000); Cameron (2000); Cooper (2002); Lewis (2003); Glaeser and Shleifer (2003)

or Congressional and executive oversight of the bureaucracy. Analysts typically assume that actors know perfectly each other's payoffs or that there is some constant private information about payoffs about which one can learn over the course of play. However, real bargaining involves uncertainty about players' own payoffs, and often this uncertainty cannot be reduced during the game through learning.

Recent efforts (McKelvey and Palfrey, 1995, 1996, 1998; Fey, McKelvey, and Palfrey, 1996) on similar models with discrete action spaces have considered **quantal response equilibrium** (QRE) solutions, so termed because they consider mutual best responses when both players share uncertainty about the payoffs over a discrete set of actions (McFadden, 1976). When considering continuous action spaces in games with the same information structure, a general term is more appropriate; therefore, we define and use the term **strategic statistical model** (SSM).

There are two treatments of the ultimatum game as a strategic statistical model. Yi (2005) develops an **experimental SSM**, which assumes that utility is the same as received payment. In addition, Yi assumes that neither the proposer nor the responder derives any utility (a reservation wage) if the responder rejects the offered division of Q, the "pie". Behaviorally, the bargaining interval $[0, Q]$ is assumed to accepted by both users and uncertainty is assumed to stem from player error. This means that, regardless of how much uncertainty the proposer faces, she has no interest in proposing a division of the resource they are splitting that exceeds the bounds $[0, Q]$ of the resource, either by being negative or greater in size than the resource. The result of this assumption is that it is relatively easy to calculate the likelihood of a given offer via backward induction.

Ramsay and Signorino (2005) on the other hand take an **observational SSM** tack, which assumes that utility is an estimable function of observables. They let $X\beta$ represent the reservation wage for each player, where X can represent individual subject characteristics. This induces an optimal offer that can be any real value, which means that the proposer might actually most prefer to make an offer outside the $[0, Q]$ bounds of the resource. Restriction to

$[0, Q]$ leads to a Tobit-like structure where the model assigns positive probability to observing offers on the boundary. Ramsay and Signorino's approach offers some advantages, but it takes some work for them to show that the latent 'most preferred offer' is always finite and unique. The effect is a more theoretically reasonable model, but one that is harder to handle mathematically.

Other research with observational strategic statistical models usually substitutes $X\beta$ for the entire utility (Signorino, 1999) which means that the rationality parameter λ and vector β are not jointly identified. The common solution is to fix $\lambda = 1$ so that β is identified. For the model described here, λ and β are jointly identified,[2] which opens the door to simultaneously estimating both rationality and the effect of subject-level covariates.

This paper makes several important contributions. First, it compares the implications of the different approaches taken by Yi and Ramsay and Signorino to determine the difference in their substantive interpretations: Yi's approach corresponds to uncertainty about actions; Ramsay and Signorino's corresponds to uncertainty about payoffs. Next, it develops a combined model which allows for inclusion of subject-level covariates and the estimation of λ. Finally, it presents experimental evidence suggesting that these two sources of uncertainty are observationally quite distinct in real data and, in the analyzed experiment, uncertainty about payoffs better accounts for the behavior of real bargainers than uncertainty about actions.

2 Problems with the Ultimatum Game

2.1 Ultimatums in Theory

Theoreically, the ultimatum game is simple. There are two players, a proposer and a responder, and a "pie" Q to be divided between them. The proposer offers a share s of the pie to

---

[2] Ramsay and Signorino (2005) derive results without assuming $\lambda = 1$ but in the analysis assume $\lambda = 1$, an unnecessary restriction.

Figure 1: The Basic Ultimatum Game

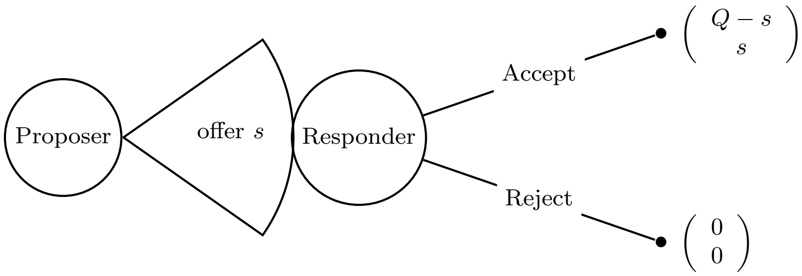

Q = pie to be divided
s = offer to responder

go to the responder. If the responder accepts the offer, then he[3] receives the share s and the proposer keeps $Q - s$. If the responder rejects the offer, both players receive nothing. (See Figure 1.)

What are the equilibria of this game? There are infinitely many Nash equilibria of the form "The responder has a minimum amount that he will accept, the proposer offers that amount, and the responder accepts." However, none of these are credible except when that minimum amount is zero; the only subgame-perfect Nash equilibrium is "The proposer offers the minimum positive amount and the responder accepts." Either way, subgame-perfection yields two hypotheses.

Hypothesis (SPNE Offer)**.** *The proposer will minimize her offer.*

Hypothesis (SPNE Response)**.** *The responder will accept any nonnegative offer.*

2.2 Ultimatums in Practice

Experimentally, this has been tested many times in many different settings, cultures, and other variations.[4] Two things are clear from the experimental results: (1) The proposer

---

[3] I use female pronouns to refer to the proposer of ultimatums and male pronouns to refer to the responder.

[4] Roth (1995) provides a good summary of many of these. Henrich, Boyd, Bowles, Camerer, Fehr, and Gintis (2004) provide results from the ultimatum game run in fifteen different small societies around the world.

Figure 2: The Ultimatum Game with Reservation Wages

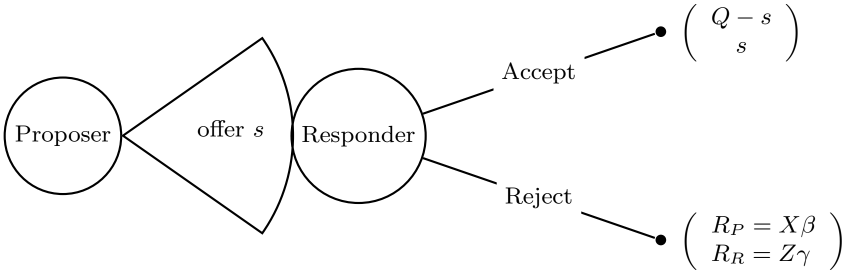

$$\begin{aligned} Q &= \text{pie to be divided} \\ s &= \text{offer to responder} \\ R_P, R_R &= \text{reservation wage for proposer, responder} \\ X, Z &= \text{covariates for proposer, responder} \\ \beta, \gamma &= \text{parameters for proposer, responder} \end{aligned}$$

almost never offers a minimal amount, and (2) the responder rejects about 15% of positive offers. This coincides with what is observed outside the lab: “lowball” offers are “insulting,” and likely to be rejected out of a sense of fairness or pride.

Clearly there is something wrong with the SPNE hypotheses. Considering some form of repeated interaction or allowing non-credible threats by the responder would account for superminimal offers. Unfortunately, experiments with one-shot ultimatums (no repetition), strict anonymity and no communication (hence no threats) eliminate these as possible explanations for the offers. Even when these might apply, neither provides much insight into the surprisingly high rate of offer rejection.

2.3 Alternate Specifications

2.3.1 The Responder’s Reservation Wage Is Not Zero

A slight refinement of the game assumes that there are some fixed amounts each players receives in case the offer is rejected. These reservation wages are denoted here R_P, R_R for the proposer and responder, respectively. (See Figure 2.)

Reservation wages provide a simple way of explaining the existence superminimal offers: the proposer knows the responder's reservation wage, and beats it my a minimal amount. However, some problems remain. If we limit the model to having homogeneous reservation wages, then we have no reason to observe heterogeneous offers. If we allow heterogeneous reservation wages, we beg the question of the source of the heterogeneity: gender? culture? Regardless, as long as reservation wages are known to the proposer, there is no reason to observe rejections, so the model is still lacking.

Allowing for reservation wages is a step in the right direction. The model developed here includes reservation wages.

2.3.2 Measurement/Regressor Error

Another possible specification of the model is to allow for researcher uncertainty about the data: the size of the pie, the size of the offer, and each players' reservation wage.[5] This has the appeal of being both plausible and commonly acknowledged in statistical settings. However, assuming researcher uncertainty is not the same as assuming player uncertainty. If the only uncertainty injected into the model is the uncertainty of the researcher, then presumably the proposer knows well the responder's reservation wage, so still we should see no offers rejected.

2.4 Zero Likelihood

Some critics[6] of game theory and rational actor models note that the predictions of such a simple game diverge completely from experimental and observational data and conclude that the assumptions of game theory are too unrealistic. Others[7] take a more measured approach and suggest that the predictions can be brought into line with observation by

---

[5] We can also allow uncertainty about the data on whether offers were taken or rejected. The only way this can account for rejections is if we assume that all rejections are data errors. This means that all offers were, in fact, accepted, which is not credible.

[6] Green (2002, section 6B) provides a good summary of critiques of rational choice theory.

[7] For example, Guth and Tietz (1990).

choosing alternate solution concepts. This later kind of thinking contributed to a focus on equilibrium refinement.[8] Recently, less effort has been placed on refinements and more has been placed on using substantive considerations to choose from among multiple equilibria, but perhaps this is still heading the wrong direction. Instead of trying to reduce the number of equilibria so that a point prediction can be identified, we suggest that increasing the number of possible predicted outcomes to the set of all possible observed outcomes may be more useful in building theories that are testable.

Signorino (2003, p. 318) differentiates between theoretical models, which are games,[9] and statistical models, which give "positive probability over all outcomes". These two types of models are not mutually exclusive, but a key problem with the way many theories are empirically tested is that they are generally set up as theoretical models that are not statistical models. Given players' reservation wages, subgame-perfect Nash equilibrium predicts with certainty one particular offer-response bundle, and no other outcomes are given any positive probability. This leads to a "zero-likelihood" problem (Signorino, 1999) in which any behavior not fitting the point prediction of the model yields a likelihood of zero, hence the likelihood of observing that data given the model is zero. This prevents straightforward analysis by maximum-likelihood or Bayesian methods.

Strategic statistical models resolve this by assuming some form of uncertainty (about payoffs or actions) which leads to observing any possible outcome with positive likelihood. This means that no matter what action is observed, a positive likelihood can be calculated, facilitating direct statistical analysis.

Note that once a researcher adopts a model such as a strategic statistical model where all possible actions are given positive likelihood, there is no way to "falsify" it. Some research (Haile, Hortaçsu, and Kosenok, 2006) has tested the assumptions of SSMs, but this does not

---

[8] Kreps and Wilson (1982); Banks and Sobel (1987); Fudenberg and Tirole (1991) provide discussion of the effort toward equilibrium refinement.

[9] Signorino (2003) writes that "from a choice-theoretic perspective, a theoretical model" is a game in the sense that one need specify the players, decision nodes, options at each node, information sets, and incentives over outcomes.

impeach SSMs as models. Models are assumed; hypotheses are tested and perhaps falsified *given a model.* It is appropriate to think of this as analogous to linear regression. One may question whether a linear model is appropriate for a given data set, but one can generally run the regression. Even more important, if a linear model is not appropriate, then a restriction of the linear model is similarly not appropriate. By analogy, there is a subgame-perfect Nash equilibrium which is a restriction of a SSM,[10] so if no SSM is appropriate then comparative statics derived from a subgame-perfect Nash equilibrium are not, either.[11]

3 Modeling Strategic Uncertainty

3.1 Quantal Response Equilibria and Strategic Statistical Models

A solution adopted by Signorino and others[12] to these problems is to consider a quantal response equilibrium.[13] Substantively, this modeling approach relies on three assumptions. First, before taking an action each player has uncertainty, either about her own and others' payoffs for each outcome *or* about the actions that each player will choose in a game. Second, large uncertainty (or errors) are less likely than small uncertainty. Third, players take into account others' uncertainty (or errors) when choosing their best responses. Together, these yield a QRE whose exact structure depends on the assumed distribution of the errors. If one assumes that these errors have a Type I Extreme Value distribution,[14] this leads one to the

---

[10] More specifically, a logit QRE parameterized by learning parameter λ has a subgame-perfect Nash equilibrium, the unique limiting logit QRE, obtained by restricting λ to be arbitrarily large (100 is practically sufficient.) This yields predicted behavior that is the same as predicted by a subgame-perfect Nash equilibrium.

[11] Comparative statics on the limiting logit QRE (a subgame-perfect Nash equilibrium) are captured by the associated SSM; if there are other subgame-perfect Nash equilibria, they might conceivably provide separate insights.

[12] McKelvey and Palfrey (1995); Anderson, Goeree, and Holt (1998); Yi (1999, 2005); Levine and Palfrey (2007)

[13] A recent exchange (Carrubba, Yuen, and Zorn, 2007a; Signorino, 2007; Carrubba, Yuen, and Zorn, 2007b) suggests that this approach may not be necessary, and the debate is ongoing. However, it has not been suggested that QRE-based solutions give incorrect inferences, only that there might be easier ways to reach the same inferences.

[14] There are some good theoretical properties to this assumption. For example, Anderson, Goeree, and Holt (2002, p. 23) note that this error structure implies that behavioral predictions (in the form of probabilities

Logit QRE (LQRE).

Problems of this type are called “quantal response” problems because they identify mutual best responses when both players share a common uncertainty about the payoffs over a *discrete* (finite or countable) set of actions (McFadden, 1976). The models considered here generalize the informational uncertainty of a quantal response equilibrium to a game with a one-dimensional compact continuous choice space for one of the players, yet the equilibria retain much of the character of quantal response equilibria. When considering continuous action spaces in games with this information structure, a more general term – **strategic statistical model** – is more appropriate.

When used to model uncertainty over payoffs, strategic statistical models do not represent a new or different solution concept. The behavior predicted is still a mutual best response, still a Nash equilibrium, *given the different information structure.* The difference between a QRE and an SPNE lies in the information structure of the game, not in the solution concept. The QRE is the exactly the subgame-perfect Nash equilibrium (SPNE) for the modified game. The term “strategic statistical model” clarifies this by emphasizing the model over the solution concept. When used to model agent error, we are no longer looking at a Nash equilibrium: players do not consistently choose best responses. However, agent error strategic statistical models still capture imperfect strategic behavior and address the zero-likelihood problem. Therefore, strategic statistical model is a useful umbrella term.

It is reasonable to ask why we do not use an already a commonly accepted solution concept which addresses uncertainty in a strategic setting, the perfect Bayesian equilibrium. This solution concept is not suitable here for both methodological and substantive reasons. Methodologically, the perfect Bayesian equilibrium is a refinement of Nash equilibrium. As discussed in Section 2.4, refinements reduce the number of possible equilibria, which means the zero-likelihood problem still prevents statistical straightforward statistical analysis. Substantively, the source of uncertainty is not something about which the players can learn during

---

of particular actions) are unchanged by adding a constant to all payoffs.

the course of play. Given agent error, neither player knows with certainty what he or she will do until the time comes to act; under uncertainty over payoffs the proposer has no chance to learn about the responder's private information and the responder doesn't care about the proposer's private information. With no learning, there is no application of Bayes Theorem, so the priors are the same as the posteriors, hence no reason to consider a perfect Bayesian equilibrium.

3.2 Two Takes on Ultimatum Uncertainty

Yi (2005) and Ramsay and Signorino (2005) both investigate SSM versions of the ultimatum game. Both make stated and unstated assumptions which have significant effects on the predictions of the models. We explore these assumptions to determine which can be merged into a single model and which we must decide between.

The most obvious difference between the two is that Yi estimates λ, the rationality or learning parameter,[15] but Ramsay and Signorino instead estimate a vector β of parameters which describes how a set X of covariates combine to form the reservation wage for each player. SSMs generally take one of these two forms, estimating λ in experimental settings when the payoffs are known and assumed to be the same as utility, and estimating utility as $X\beta$ in observational settings. One usually does not estimate λ in observational settings because a typical observational SSM replaces the *entire* utility with $X\beta$. Everywhere λ appears it is multiplied by the utility. This means that λ and β are not jointly identified.[16]

Ramsay and Signorino take a slightly different path. They replace only part of each player's expected utility, the reservation wage, with a set of regressors $X\beta$. In this model, λ and β are jointly identified.[17] The model described here leaves λ as a parameter to be

---

[15] The rationality or learning parameter λ is inversely proportional to the standard deviation of the logistically distributed uncertainty in the model. When $\lambda = 0$, players appear to mix uniformly over all possible actions; as $\lambda \to \infty$, the observed behavior approaches that of a unique subgame-perfect Nash equilibrium called the "limiting logit QRE".

[16] To see this lack of joint identification, consider the expression $\lambda X\beta$. Replace λ with 2λ and β with $\frac{1}{2}\beta$. The product is unchanged, so any likelihood would be identical for the same data.

[17] Ramsay and Signorino derive this joint identification in the appendix, but in the body of the paper

Figure 3: Comparison of the Predicted Distribution of Ultimatum Offers in Yi (Agent Error) or Ramsay and Signorino (Uncertainty Over Payoffs) Models

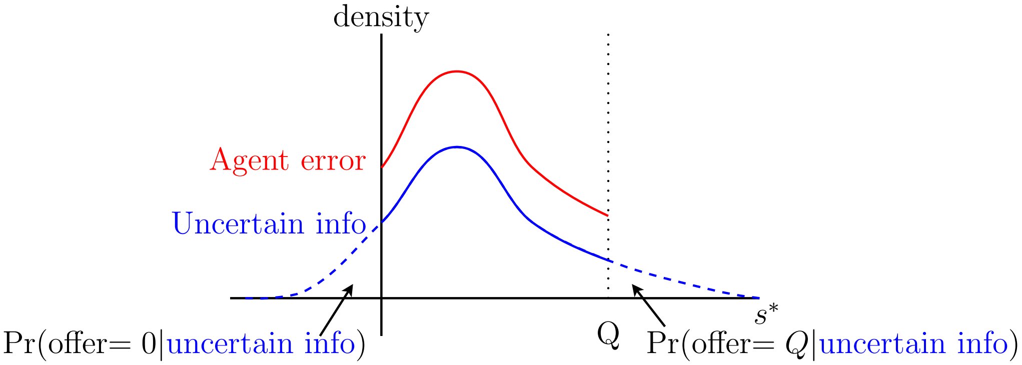

Notes: The distribution of offers under agent error is a continuous distribution on the bounded interval. The distribution of offers under uncertainty over payoffs is a mixture of a continuous distribution on the interval (proportional to the distribution under agent error) with masses at both ends. Thus the probability of an offer on the boundary is zero under agent error, but is positive under uncertainty over payoffs.

estimated, which allows us to make inferences about how players change their strategies over multiple runs of the game.

A more careful inspection reveals that Yi and Ramsay and Signorino differ on the assumed data-generating process. While Yi mentions that there may be some non-monetary component to utility, he defines the likelihood function so that there is zero probability of observing offers on the boundary of the feasible set. This implies that that uncertainty is essentially about actions, the so-called "agent error" specification of SSM. Also unstated is Yi's assumption that there is *no* reservation wage: all positive offers from the proposer are the result of (agent) error.

Ramsay and Signorino explicitly model a data-generating process assuming uncertainty over payoffs. Both this and the assumption that each player has a real-valued reservation wage are each sufficient to give the proposer an incentive to make offers outside the $[0, Q]$ range at least some of the time.[18] Observationally, this means that Ramsay and Signorino

---

assume that $\lambda = 1$.

[18] For example, if the responder has a negative reservation wage then the proposer's most-preferred offer s can be negative, even though that is not a permitted offer.

expect to see offers on the boundaries with positive probability. The distributions of offers in $(0, Q)$ are proportional for the two models, but Ramsay and Signorino expect fewer offers in the interior because of those expected on the boundaries. Figure 3 illustrates the relationship between the distributions of expected offers.

3.3 Example of Substantive Difference between Agent Error and Uncertainty Over Payoffs

An example may help in interpreting these competing sources of uncertainty. Suppose that the Environmental Protection Agency is going to propose a new policy setting the maximum level of some new pollutant that will be acceptable in drinking water. They know that an industry watchdog organization is, perhaps imperfectly, monitoring their actions, and that the industry prefers higher permissable levels. If the agency chooses a level that is too high, the watchdog might call in favors to get Congress to hold hearings, even though disrupting the regulatory process is costly (they can't do this for every policy) and might lead to more stringent regulation for this new pollutant. If there were no uncertainty, the agency could always choose a level that avoided Congressional hearings; empirically, Congressional hearings happen, so it is likely there is some form of uncertainty.

If the watchdog were not certain to notice this one new regulation on a new pollutant, or if the watchdog might not know what a safe level for this pollutant is, then this would manifest as uncertainty over the *actions* of the watchdog, hence an agent error specification. If the agency were not sure exactly what level the watchdog would find barely acceptable, or if the agency were not sure about how costs affect the watchdog's decision, then this would manifest as uncertainty over the *payoffs*.

3.4 A Combined Strategic Statistical Model

The only feature of these two models that cannot be combined into a single model is the assumed source of uncertainty. Is the uncertainty about actions, per Yi, or about payoffs, per Ramsay and Signorino? We form two versions of the combined model, one assuming each type of uncertainty. Both can be estimated and the relative fit of the models compared.

For both models we assume that there is some reservation wage, possibly zero, for each player. This reservation wage is estimated by $X\beta$ for the proposer and by $Z\gamma$ for the responder, per Figure 2.

The probability that the responder accepts an offer s is the same for both sources of uncertainty, and is

$$\begin{aligned}\Pr(\text{accept}|s) &= \frac{\exp[\lambda s]}{\exp[\lambda s] + \exp[\lambda \cdot Z\gamma]} \\ &= \frac{1}{1+\exp[\lambda(Z\gamma - s)]} \\ &= \Lambda(\lambda[s - Z\gamma]) \qquad (1)\end{aligned}$$

where Λ is the cumulative distribution function of the logistic distribution. Then the expected payoff[19] to the proposer by offering s is

$$\begin{aligned}E[u_P|s] &= (Q-s)\Pr(\text{accept}|s) + R_P \cdot \Pr(\text{reject}|s) \\ &= (Q-s)\cdot\Lambda(\lambda[s-Z\gamma]) + X\beta\cdot(1-\Lambda(\lambda[s-Z\gamma])) \qquad (2)\end{aligned}$$

3.4.1 Combined Model with Agent Error

Suppose that the source of uncertainty is agent error. This means that the proposer chooses an offer $s \in [0, Q]$ where the probability of a given offer is proportional to an exponential of

---

[19] We are not assuming that payments are equal to utility as do many experimental designs. Rather, we assume that utility is at most an affine, order-preserving transformation of payment plus some linear function of covariates. Putting no coefficient on the payment serves to identify β, hence identify the relative weight of payment and non-payment components of utility.

the expected utility for that offer. This yields a distribution

$$f_{AE}(s) = \frac{\exp\left(\lambda E[u_P|s]\right)}{\displaystyle\int_{s=0}^{s=Q} \exp\left(\lambda E[u_P|s]\right) ds} \tag{3}$$

and a likelihood

$$L_{AE} = \prod_{i=1}^{n} f_{AE}(s_i) \times \Pr\left(\text{accept}|s_i\right)^{I(a_i=\text{accept})} \cdot \left(1 - \Pr\left(\text{accept}|s_i\right)\right)^{I(a_i=\text{reject})} \tag{4}$$

where $I(x)$ is an indictor function (1 if x is true, 0 if x is false.)

3.4.2 Combined Model with Uncertainty Over Payoffs

Suppose that the uncertainty about the action of the proposer is based on uncertainty over payoffs. Then the proposer has[20] a latent optimal offer $s^* \in \mathbb{R}$. The observed offer s is s^* truncated to the feasible region $[0, Q]$. To identify s^*, consider the first order condition for the proposer's expected utility in Equation (2); this yields

$$s^* = Q - X\beta - \epsilon_P - \frac{1}{\lambda}\big(1 + \exp[\lambda(s^* - Z\gamma)]\big) \tag{5}$$

where ϵ_P is the manifestation of the proposer's payoff uncertainty and 'logit' is the density of the logistic distribution. Substituting in the expressions for the logistic distribution and density functions and solving[21] for s^* yields

$$s^* = Q - X\beta - \epsilon_P - \frac{1}{\lambda}\big(1 + \mathcal{W}(\exp[\lambda(Q - X\beta - Z\gamma - \epsilon_P) - 1]\big) \tag{6}$$

where $\mathcal{W}$ is Lambert's W function.[22] Note that s^* is a function of ϵ_P which has a logistic distribution with mean zero and scale parameter $\frac{1}{\lambda}$. Applying the method of transformations

---

[20] Ramsay and Signorino show that this latent optimal offer exists and is finite.

[21] This calculation parallels that by Ramsay and Signorino.

[22] Lambert's W gives solutions to certain exponential equations, and is defined implicitly: $\mathcal{W}(z) = w$ if $z = we^w$. The first and second derivatives exist and are well-behaved, and it is monotonic on $\mathbb{R}_+$.

yields the density and cumulative distribution of s^*:

$$f_{s^*}(s^*) = \frac{\lambda \exp\left[-\lambda h^{-1}(s^*)\right] \cdot \left[1 + \exp\left[\lambda(s^* - Z\gamma)\right]\right]}{\left(1 + \exp\left[-\lambda h^{-1}(s^*)\right]\right)^2} \tag{7}$$

$$F_{s^*}(s^*) = \frac{1}{1 + \exp\left[-\lambda h^{-1}(s^*)\right]} \tag{8}$$

where $h^{-1}(s^*)$ comes from solving Equation (5) for ϵ_P, yielding

$$h^{-1}(s^*) = Q - X\beta - \frac{1}{\lambda}\big(1 + \exp[\lambda(s^* - Z\gamma)]\big) - s^*$$

Truncating the offers to the feasible interval yields a predicted distribution

$$f_{UP}(s) = \begin{cases} F_{s^*}(0) & s = 0 \\ f_{s^*}(s) & 0 < s < Q \\ 1 - F_{s^*}(Q) & s = Q \end{cases} \tag{9}$$

and a likelihood

$$L_{UP} = \prod_{i=1}^{n} f_{UP}(s_i) \times \Pr\left(\text{accept}|s_i\right)^{I(a_i = \text{accept})} \cdot \left(1 - \Pr\left(\text{accept}|s_i\right)\right)^{I(a_i = \text{reject})} \tag{10}$$

It is worth noting that the distribution of offers under agent error in Equation (3) can be expressed without the integral by using the closed-form expressions generated in the analysis of the uncertainty over payoffs case:

$$f_{AE}(s) = \frac{f_{s^*}(s)}{F_{s^*}(Q) - F_{s^*}(0)} \tag{11}$$

3.4.3 Distinguishing Agent Error and Uncertainty Over Payoffs

The two models are not nested, so we cannot use the likelihood ratio test. Perhaps the best way to compare non-nested models is to calculate the Bayes factors. Bayesian statistical estimation of these models is beyond the scope of this project, so the next best thing is to use one or more of the information criteria appropriate for maximum likelihood estimation. We report the Bayesian information criterion (BIC).

4 Experimental Evidence

An original experiment was designed and run to determine which of the two potential sources of uncertainty has more empirical support and to differentiate among substantive explanations for rejection of offers by responders. This paper reports the results of the comparison of the sources of uncertainty; a full analysis of the substantive results will be reported in subsequent work. This data comes from eight sessions in September 2008 using college students at a private midwestern university.

4.1 Why run an experiment?

Often, experiments give up external validity to gain internal validity. External validity is not a concern in this study: we simply want to establish that, under conditions that are reasonably close to natural, the difference between the two sources of uncertainty is large enough to affect observable results. Internal validity, is essential. In trying to distinguish between two types of uncertainty – over actions and over payoffs – we must have tight control over the information environment of the subjects. This is impossible outside the laboratory

4.2 Design

The goal is to detect the difference between behavior assumed to be influenced by one of two sources of uncertainty. Recall that the predicted distributions are proportional in the interior of the pie, but differ in that uncertainty over payoffs suggests that there will be some positive probability of observing a boundary offer. Figure 3 is suggestive, but the vertical scale is greatly compressed: the numerator of Equation 3 is an exponential, so the tails of

Figure 4: Predicted Effect of Varying Minimum Offers in the Ultimatum Game

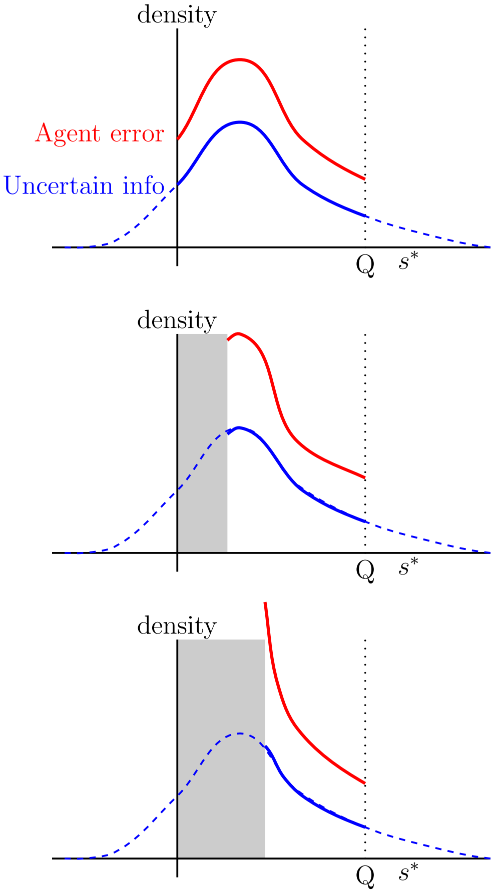

Notes: Based on data from an original experiment. $n = 310$ dyadic interactions, 62 subjects playing the ultimatum game ten times each in pairs. Bar plots on the left give the probability of a minimum offer (observed and predicted).

the densities are very thin. If the modal offer is in the middle of the pie, the probability of observing a boundary offer is very close to zero, and the two distributions are almost observationally equivalent.

Suppose that the proposer is required to offer at least some **minimum permitted offer** just below the modal offer, as in the middle of Figure 4. Under agent error, we are likely to observe an offer above the minumum offer, but under uncertainty over payoffs were have a good chance of observing a minimal offer. If the minimum permitted offer is too high, as in the bottom of Figure 4, it becomes very likely under agent error that we will observe an offer very close to the minimum. In fact, given the focal nature of a minimum offer, we may in practice observe a share of minimal offers under either of the two sources of uncertainty.

In previous experiments, offers tend to cluster around 30%-40% of the pie. To be sure of finding a minimum permitted offer that induces observationally distinct behavior, we selected minimum permitted offers at random uniformly from the bottom half of the pie.

The "pie" was set at 400 cents for each ultimatum game "round." There were 62 subjects in eight laboratory sessions who each played 10 rounds with partners who were randomly selected each round. Each of the ten rounds played by each subject involved a different anonymous partner, a different minimum permitted offer, and an independent choice of role as proposer or responder. This yielded 310 dyadic interactions which, for this analysis, were pooled across sessions and rounds.

Both proposer and responder used an slider to make and evaluate offers. The slider was unmarked except to show the minimum permitted offer, which was common knowledge for the players. (See Figure 5.) The proposer had one minute to make and offer; the responder had one minute to accept or reject the offer; none of the subjects took the full minute in any round.

4.3 Results

Two models were fit to the resulting data: the combined models under agent error (Section 3.4.1) and uncertainty over payoffs (Section 3.4.2). The responses are observationally equivalent under the two models, so only the proposers' offers were analyzed. For simplicity, in the analysis presented here the offers were pooled across the ten rounds; estimating λ_t for each round t does not affect the results. No covariates are included in this analysis; their inclusion does not affect the results. The parameters estimated were the rationality parameter λ, a constant reservation wage for the proposer β, and a constant reservation wage for the responder γ.

The observed offers are shown in Figure 6. The center plot shows the observed offers. The

Figure 5: z-Tree Slider Allowing Proposer to Make Offer

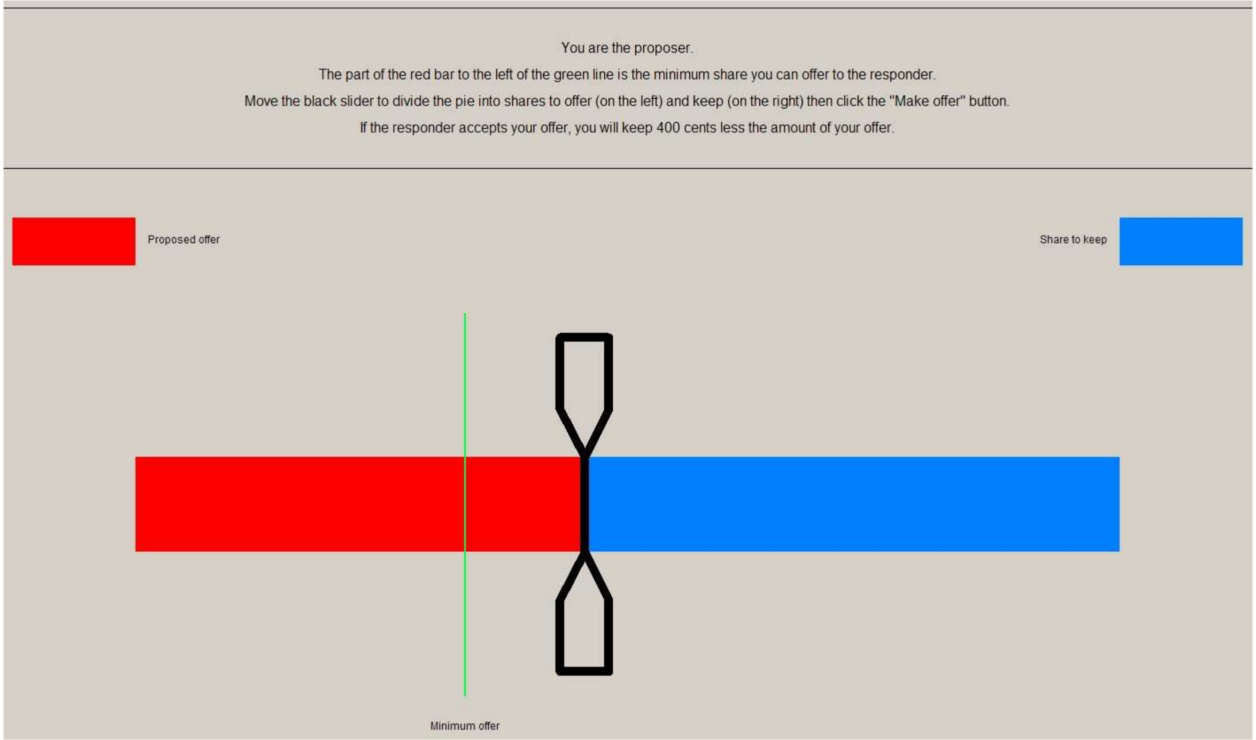

You are the proposer.
The part of the red bar to the left of the green line is the minimum share you can offer to the responder.
Move the black slider to divide the pie into shares to offer (on the left) and keep (on the right) then click the "Make offer" button.
If the responder accepts your offer, you will keep 400 cents less the amount of your offer.

Notes: The experiment was conducted using z-Tree software. Bar colors were randomized each round.

Figure 6: Observed and Fit Offers from an Experiment with Varying Minimum Permitted Offers

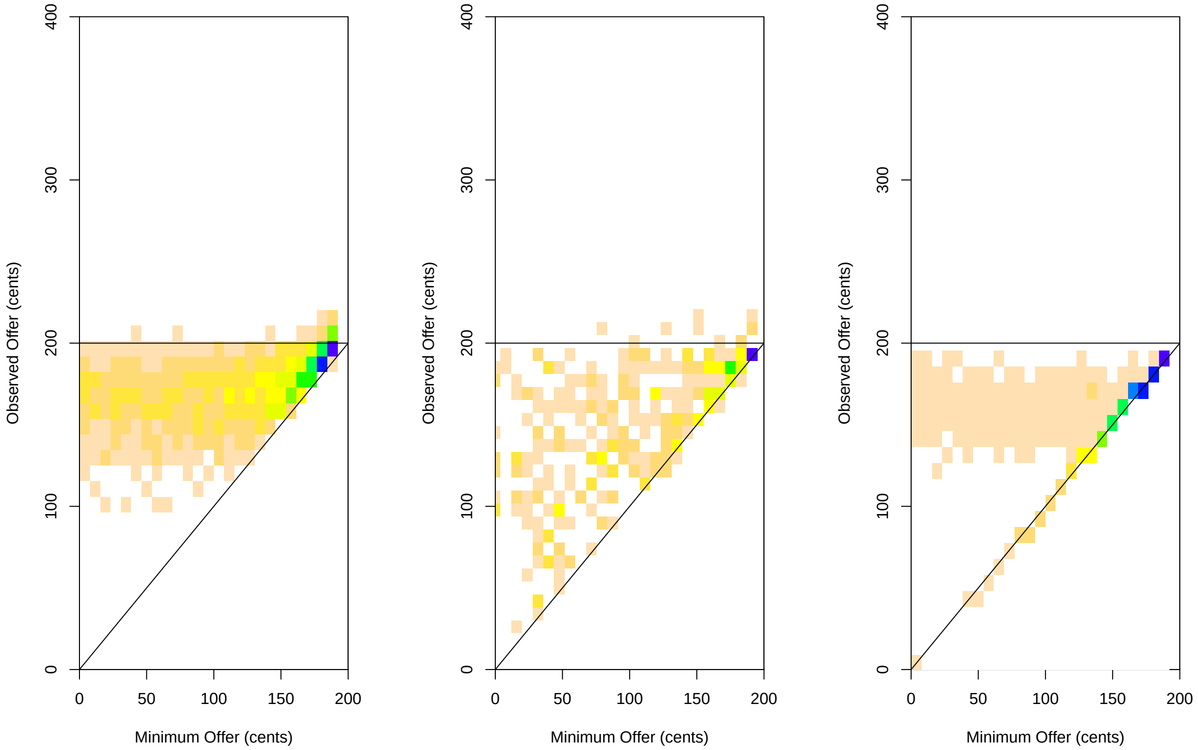

Agent Error

Observed

Uncertain Payoffs

Notes: Based on data from an original experiment. $n = 310$ dyadic interactions, 62 subjects playing the ultimatum game ten times each in pairs. The center plot shows the observed offers. The left plot shows random draws from the predicted distribution of draws assuming agent error and using estimated coefficients; the right plot is similarly constructed assuming uncertainty over payoffs. The slanted line shows the minimum permitted offer.

left plot was generated by first fitting the combined model with agent error, then taking the estimated coefficients (λ, β, γ) and generating random draws from the predicted distribution. The right plot is similarly constructed assuming uncertainty over payoffs. As expected, the fit values for agent error in the left plot show essentially no boundary offers, and the fit value for uncertainty over payoffs show noticable boundary offers.

The observed offers show noticable boundary offers, which suggests that uncertainty over payoffs may be a better explanation than agent error. However, it is difficult to make a definitive determination from these graphs.

Table 1: Comparison of Models of Based on Observed Offers and Responses from an Ultimatum Experiment with Varying Minimum Permitted Offers

Model	**BIC**	**Results**
Agent error	438.2	
Uncertainty over payoffs	**-266.1**	Lower BIC ⇒ **better fit**

Notes: Based on data from an original experiment. $n = 310$ dyadic interactions, 62 subjects playing the ultimatum game ten times each in pairs.

A comparison of the BICs (Table 1) proves more telling. The BIC for the agent error model is 438.2; the BIC for the uncertainty over payoffs model is -266.1, which strongly suggests that the uncertainty over payoffs model fits the data better.[23]

5 Conclusion

Offers and responses in the ultimatum game do not follow the predictions of a model with complete and certain information. Recent work suggests introducing either of two different sources of uncertainty – about actions or payoffs – to solve the "zero-likelihood" problem and permit direct statistical comparison of substantive explanations for these deviations. Using a combined model and implementing each of these sources of uncertainty, we examine original data from a laboratory experiment designed to differentiate between these sources of uncertainty. The data speak clearly, suggesting that these two sources of uncertainty are observationally quite distinct.

How might this manifest in an observational study? Returning to the example of Section 3.3, if the results of an observational study about the Environmental Protection Agency paralleled those of the experiment, we might conclude that the EPA has more uncertainty about the costs and preferences of watchdog organizations and less uncertainty about the ability of the watchdogs to observe the EPA and to identify their own preferred level of regulation. Normatively, we might seek to minimize the time bureaucrats spend testify-

---

[23] Other specifications of the model were fit; the results consistently suggested that the imperfect information specification fit far better than did agent error.

ing before Congress. Understanding that (still hypothetically) the agency has uncertainty about the costs and preferences of the watchdogs, we might remove statuatory or regulatory communication barriers between the watchdogs and the agency.

Further work using this data will examine competing explanations for why proposers offer "too much" and responders "mistakenly" reject positive offers.

References

Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt. 1998. "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction." *Journal of Political Economy* 106(4):828–853.

Anderson, S.P., J.K. Goeree, and C.A. Holt. 2002. "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies." *Southern Economic Journal* 69(1):21–47.

Banks, Jeffrey S., and Joel Sobel. 1987. "Equilibrium Selection in Signaling Games." *Econometrica* 55(3):647–661.

Brehm, John, and Scott Gates. 1997. *Working, Shirking, and Sabotage: Bureaucratic response to a Democratic Public*. University of Michigan Press.

Cameron, Charles M. 2000. *Veto Bargaining : Presidents and the Politics of Negative Power*. Cambridge University Press.

Carrubba, Clifford J., Amy Yuen, and Christopher Zorn. 2007a. "In Defense of Comparative Statics: Specifying Empirical Tests of Models of Strategic Interaction." *Political Analysis* 15(4):465–482.

Carrubba, Clifford J., Amy Yuen, and Christopher Zorn. 2007b. "Reply to Signorino." *Political Analysis* 15(4):502–504.

Cooper, Phillip J. 2002. *By Order of the President: The Use & Abuse of Executive Direct Action*. University Press of Kansas.

Ferejohn, John, and Charles Shipan. 1990. "Congressional Influence on Bureaucracy." *Journal of Law, Economics, & Organization* 6(Special Issue: [Papers from the Organization of Political Institutions Conference, April 1990]):1–20.

Fey, Mark, Richard D. McKelvey, and Thomas R. Palfrey. 1996. "An Experimental Study of Constant-sum Centipede Games." *International Journal of Game Theory* 25:269–287.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics* 10(2):171–178.

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. The MIT Press.

Glaeser, Edward L., and Andrei Shleifer. 2003. "The Rise of the Regulatory State." *Journal of Economic Literature* 41(2):401–425.

Green, Steven L. 2002. "Rational Choice Theory: An Overview." Prepared for the Baylor University Faculty Development Seminar on Rational Choice Theory.

Guth, Werner, and Reinhard Tietz. 1990. "Ultimatum bargaining behavior: A survey and comparison of experimental results." *Journal of Economic Psychology* 11(3):417–449.

Haile, Philip A., Ali Hortaçsu, and Grigory Kosenok. 2006. “On the Empirical Content of Quantal Response Equilibrium.” Cowles Foundation Discussion Paper No. 1432R.

Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis, editors. 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford University Press.

Kagel, John H., and Alvin E. Roth, editors. 1995. *The Handbook of Experimental Economics*. Princeton University Press.

Kreps, David M., and Robert Wilson. 1982. “Sequential Equilibria.” *Econometrica* 50(4):863–894.

Levine, David K., and Thomas R. Palfrey. 2007. “The Paradox of Voter Participation? A Laboratory Study.” *American Political Science Review* 101(1):143–158.

Lewis, David E. 2003. *Presidents and the Politics of Agency Design: Political Insulation in the United States Government Bureaucracy, 1946-1997*. Stanford University Press.

Long, J. Scott. 1997. *Regression Models for Categorical and Limited Dependent Variables*, volume 7 of *Advanced Quantitative Techniques in the Social Sciences*. Sage Publications.

McFadden, Daniel. 1976. “Quantal Choice Analysis: A Survey.” *Annals of Economic and Social Measurement* 5(4):363–390.

McKelvey, Richard D., and Thomas R. Palfrey. 1995. “Quantal Response Equilibria for Normal Form Games.” *Games and Economic Behavior* 10(1):6–38.

McKelvey, Richard D., and Thomas R. Palfrey. 1996. “A Statistical Theory of Equilibrium in Games.” *Japanese Economic Review* 47(2):186–209.

McKelvey, Richard D., and Thomas R. Palfrey. 1998. “Quantal Response Equilibria for Extensive Form Games.” *Experimental Economics* 1(1):9–41.

Miller, Gary J. 1992. *Managerial Dilemmas: The Political Economy of Hierarchy*. The Political Economy of Institutions and Decisions. Cambridge University Press.

Neustadt, Richard E. 1990. *Presidential Power and the Modern Presidents: The Politics of Leadership from Roosevelt to Reagan*. Simon & Schuster.

Ramsay, Kristopher W., and Curtis S. Signorino. 2005. “Bargaining and Society: A Statistical Model of the Ultimatum Game.” Working paper.

Roth, Alvin E. 1995. “Bargaining Experiments.” In *The Handbook of Experimental Exonomics* (John H. Kagel, and Alvin E. Roth, editors), Princeton University Press, chapter 4, pp. 253–348.

Signorino, Curtis S. 1999. “Strategic Interaction and the Statistical Analysis of International Conflict.” *American Political Science Review* 93(2):279–297.

Signorino, Curtis S. 2003. “Structure and Uncertainty in Discrete Choice Models.” *Political Analysis* 11(4):316–344.

Signorino, Curtis S. 2007. “On Formal Theory and Statistical Methods: A Response to Carrubba, Yuen, and Zorn.” *Political Analysis* 15(5):483–501.

Waterman, Richard W., Amelia Rouse, and Robert Wright. 1998. “The Venues of Influence: A New Theory of Political Control of the Bureaucracy.” *Journal of Public Administration Research and Theory* 8(1):13–38.

Webster, Jr., Murray, and Jane Sell, editors. 2007. *Laboratory Experiments in the Social Sciences*. Elsevier.

Wilson, James Q. 2000. *Bureaucracy: What Government Agencies Do and Why They Do It*. Basic Books.

Yi, Kang-Oh. 1999. “A Quantal Response Equilibrium Model of Order Statistic Games.” UCSD working paper 99-17.

Yi, Kang-Oh. 2005. “Quantal-response equilibrium models of the ultimatum bargaining game.” *Games and Economic Behavior* 51(2):324–348.