Structural Estimation of a 2 Player Game of Incomplete Information
The Prisoners Dilemma Example
Econ 625, Assignment 5

John Rust, *University of Maryland*

November, 2005

1 The Problem

You were asked to solve, simulate and estimate the identified parameters of a modified version of the prisoners dilemma game. As noted in the problem set let the two players (prisoners) be denoted by **a** and **b**. Each has two possible actions: confess (**c**) or don't confess (**d**). We will also use a numerical coding to indicate the choice of each prisoner

$$d_a = \begin{cases} 1 & \text{if prisoner a confesses} \\ 0 & \text{if prisoner a does not confess} \end{cases} \tag{1}$$

and d_b, the decision of prisoner **b**, is defined similarly.

Let $u_a(d_a,d_b,x_a,\varepsilon_a)$ denote the *ex post* payoff to prisoner **a**. This payoff (the sentence awarded) is a function of the joint decisions of the two prisoners, (d_a,d_b), a scalar variable x_a denoting prisoner **a**'s *observed type,* and a (2×1) vector ε_a denoting prisoner **b**'s *unobserved type.* Similarly, there is a utility function for prisoner **b** given by $u_b(d_a,d_b,x_b,\varepsilon_b)$, which depends on the two prisoners' decisions, player **b**'s observed type x_b and player **b**'s unobserved type ε_b.

According to the problem set, we specialize the utility functions as follows:

$$u_a(d_a,d_b,x_a,\varepsilon_a) = \theta^a_{d_a,d_b}x_a + \sigma_a\varepsilon_a(d_a) \tag{2}$$

$$u_b(d_a,d_b,x_b,\varepsilon_b) = \theta^b_{d_a,d_b}x_b + \sigma_b\varepsilon_b(d_b) \tag{3}$$

where (σ_a,σ_b) are positive *scale parameters.* Thus, each prisoner's utility function is a linear function of their observed type, x_a or x_b, with a coefficient that depends on the *joint decisions* made by the two prisoners, $(\theta^a_{d_a,d_b},\theta^b_{d_a,d_b})$, plus an additive term that depends only on the *individual decision* of each prisoner, $(\varepsilon_a(d_a),\varepsilon_b(d_b))$. We assume that these latter terms are observed only be each prisoner, but not by their opponent prisoner of by the econometrician.

Thus, we assume that the observed types of each prisoner, (x_a,x_b) are *common knowledge* (i.e. each knows their type, the other's type, and each knows that the other prisoner knows these things, etc.) but the unobserved types $(\varepsilon_a, \varepsilon_b)$ are private information. That is, prisoner **a** knows his/her value of ε_a but does not know ε_b, and vice versa for prisoner **b**. However, we assume that it is common knowledge that each prisoner believes that the unobserved type of his/her opponent prisoner is distributed according to a standardized bivariate Type III extreme value distribution, where each component of the (2×1) vectors $(\varepsilon_a, \varepsilon_b)$ are distributed independently of each other.

The *parameters* of this model are the 10 values $(\{\theta^a_{ij}\}, \sigma_a, \{\theta^b_{ij}\}, \sigma_b)$: i.e. there are 4 possible θ coefficients for each prisoner, $(\theta^a_{11}, \theta^a_{10}, \theta^a_{01}, \theta^a_{00})$ (and similarly for prisoner **b**) corresponding to the 4 different possible combinations for (d_a, d_b), plus the two scale parameters (σ_a, σ_b). We will let θ denote the overall 10×1 vector of unknown coefficients to be estimated.

We assume that θ is known by both prisoners and that this is common knowledge just as are the values of the observed types, (x_a, x_b), however we assume that while the econometrician can observe (d_a, d_b) and (x_a, x_b), the econometrician observes neither θ nor the values of $(\varepsilon_a, \varepsilon_b)$. The point of this problem set is to determine what the econometrician can learn about θ if the econometrician is able to observe the outcomes of N independent prisoners dilemma games, each treated as an *IID* realization of an equilibrium outcome of the game. Thus, our data set will consist of $\{d^i_a, d^i_b, x^i_a, x^i_b\}_{i=1}^N$ where $(d^i_a, d^i_b, x^i_a, x^i_b)$ is the realized outcome of the i^{th} prisoners dilemma game.

Since it is hard to find real data for such outcomes, you were asked to generate *artificial data* using a fixed "true value" for the unknown parameter vector θ^* and then to see what you could learn about θ^* using these simulated artificial data. I suggested that you program up the method of *maximum likelihood* to estimate θ^*, but a simpler alternative is the *semiparametric two step estimator* described in the paper by Bajari, Hong, Krainer and Nekipelov (2005) "Estimating Static Models of Strategic Interactions".

The true values of the parameters were specified to be

$$(\sigma_a, \sigma_b) = (1, 1) \tag{4}$$

$$(\theta^a_{11}, \theta^b_{11}) = (-2, -2) \tag{5}$$

$$(\theta^a_{00}, \theta^b_{00}) = (-1, -1) \tag{6}$$

$$(\theta^a_{10}, \theta^b_{01}) = (-1/2, -1/2) \tag{7}$$

$$(\theta^a_{01}, \theta^b_{10}) = (-5, -5) \tag{8}$$

Thus, I have imposed a *symmetry assumption* on the true parameter vector, i.e. each prisoner has the same utility function and faces the same incentives to confess. The problem asks which of the parameters can be identified under the symmetry assumption and which parameters can be identified if we relax the symmetry assumption. Before I answer these questions, I describe software for solving and simulating the game. This is necessary in order to create the artificial data that we will use to try to estimate θ^*.

2 Solving the Equilibrium

Let p_a denote the probability (from prisoner **b**'s perspective) that prisoner **a** will confess. Due to the presence of incomplete information, (i.e. the ε_a terms in prisoner **a**'s utility function, when $\sigma_a > 0$), prisoner **a**'s decision will be probabilistic from the standpoint of prisoner **b**. We can view p_a as prisoner **b**'s *belief* of the probability that prisoner **a** will confess. Similarly p_b represents prisoner **a**'s belief that prisoner **b** will confess.

Given these beliefs (initially we treat the beliefs as fixed and do not consider whether they are "equilibrium beliefs" yet), prisoner **b** must take a decision, $d_b \in \{0,1\}$, to maximize his/her expected utility. The expected utility associated with the decision d_b is given by

$$\begin{aligned} E\{u_b(\tilde{d}_a, d_b, x_b, \varepsilon_b)\} &= p_a u_b(1, d_b, x_b, \varepsilon_b) + (1-p_a) u_b(0, d_b, x_b, \varepsilon_b) \\ &= p_a \theta^b_{1,d_b} x_b + (1-p_a)\theta^b_{0,d_b} x_b + \varepsilon_b(d_b). \end{aligned} \tag{9}$$

The expected utility function for prisoner **b** is just a weighted average of the utilities associated with the two possible actions that prisoner **a** might choose, with the weight equal to prisoner **b**'s belief about the probability **a** will confess, p_a. Notice that since the unobserved component of **b**'s payoff associated with the decision d_b does not depend on prisoner **a**'s decision d_a, it comes out of the expectation, and remains an additive "error term"' just as in a standard single agent discrete choice model.

As we discussed in class, when the error terms have a standardized Type III extreme value distribution, the probablity (from prisoner **a**'s perspective) that prisoner **b** will confess, *if prisoner* **a** *thinks that prisoner* **b** *thinks that* **a** *will confess with probability* p_a, is given by the standard binomial logit formula

$$\begin{aligned} p_b &= Pr\{\tilde{d}_b = 1\} \\ &= Pr\{E\{u_b(1,x_b,\varepsilon_b)\} \geq E\{u_b(0,x_b,\varepsilon_b)\}\} \\ &= \frac{\exp\{(p_a\theta^b_{1,1}x_b + (1-p_a)\theta^b_{1,0}x_b)/\sigma_b\}}{\exp\{(p_a\theta^b_{1,1}x_b + (1-p_a)\theta^b_{1,0}x_b)/\sigma_b\} + \exp\{(p_a\theta^b_{0,1}x_b + (1-p_a)\theta^b_{0,0}x_b)/\sigma_b\}} \\ &= \frac{1}{1+\exp\{x_b(\theta^b_{0,0}-\theta^b_{1,0})/\sigma_b + p_a x_b(\theta^b_{0,1}-\theta^b_{0,0}+\theta^b_{1,0}-\theta^b_{1,1})/\sigma_b\}} \end{aligned} \tag{10}$$

This logit formula can be thought of as a *best response function.* That is, it is prisoner **a**'s belief about the probability prisoner **b** will confess, if prisoner **b** believes that **a** will confess with probability p_a.

There is a symmetrical best response function for prisoner **b**:

$$
\begin{aligned}
p_a &= Pr\{\tilde{d}_a = 1\} \\
&= Pr\{E\{u_a(1,x_a,\varepsilon_a)\} \geq E\{u_a(0,x_a,\varepsilon_a)\}\} \\
&= \frac{\exp\{(p_b\theta^a_{1,1}x_a + (1-p_b)\theta^a_{1,0}x_a)/\sigma_a\}}{\exp\{(p_b\theta^a_{1,1}x_a + (1-p_b)\theta^a_{1,0}x_a)/\sigma_a\} + \exp\{(p_b\theta^a_{0,1}x_a + (1-p_b)\theta^a_{0,0}x_a)/\sigma_a\}} \\
&= \frac{1}{1+\exp\{x_a(\theta^a_{0,0}-\theta^a_{1,0})/\sigma_a + p_b x_a(\theta^a_{0,1}-\theta^a_{0,0}+\theta^a_{1,0}-\theta^a_{1,1})/\sigma_a\}} \qquad (11)
\end{aligned}
$$

As a shorthand, we will define the *best response probability functions* for prisoners **a** and **b** as

$$
\begin{aligned}
p_a &\equiv \lambda_a(p_b,x_a) \\
&= \frac{1}{1+\exp\{x_a(\theta^a_{0,0}-\theta^a_{1,0})/\sigma_a + p_b x_a(\theta^a_{0,1}-\theta^a_{1,1}+\theta^a_{1,0}-\theta^a_{0,0})/\sigma_a\}} \qquad (12)
\end{aligned}
$$

and

$$
\begin{aligned}
p_b &\equiv \lambda_b(p_a,x_b) \\
&= \frac{1}{1+\exp\{x_b(\theta^b_{0,0}-\theta^b_{0,1})/\sigma_b + p_a x_b(\theta^b_{1,0}-\theta^b_{1,1}+\theta^b_{0,1}-\theta^b_{0,0})/\sigma_b\}} \qquad (13)
\end{aligned}
$$

A *Bayesian Nash equilibrium* is simply a set of mutually confirming beliefs, i.e. a pair of beliefs (p^*_a, p^*_b) which are mutual best responses:

$$
\begin{aligned}
p^*_a &= \lambda_a(p^*_b,x_a) \\
p^*_b &= \lambda_b(p^*_a,x_b) \qquad (14)
\end{aligned}
$$

It is easy to see that in a Bayesian Nash equilibrium, each prisoner is taking a decision that maximizes their expected utility, with correct (in equilibrium) beliefs about the actions their opponent will take.

One can show, via the Implicit function theorem, that (p^*_a, p^*b) are implicit functions of (x_a,x_b) as well as functions of the preference parameters θ^*. We can emphasize this dependence by writing

$$
\begin{aligned}
p^*_a &= p_a(x_a,x_b,\theta^*) \\
p^*_b &= p_b(x_a,x_b,\theta^*) \qquad (15)
\end{aligned}
$$

where (p_a, p_b) are now interpreted as *equilibrium selection functions,* i.e. for any set of observed types, (x_a,x_b) and θ^*, these functions select one of the equilibria of the game. Using results from differential topology, particularly *Index Theorems,* one can show that for "almost all" (x_a, x_b, θ^*), the equilibria to this game will be *regular* and *isolated* and therefore, due to the compactness of the $[0,1]$ interval, there will be an *odd number of equilibria.*

A *regular equilibrium* is a pair (p_a^*, p_b^*) for which small perturbations in the types, (x_a, x_b) and/or θ^*, results in only a small perturbation in (p_a^*, p_b^*). That is, at a regular equilibrium, p_a and p_b, will be continuous functions of (x_a, x_b) and θ^*. From the Implicit Function Theorem we know that a sufficient condition for this to hold is that a certain Jacobian matrix is non-singular. To see what this Jacobian matrix must be, we can write the best response mappings as a mapping from the unit square $S = \{(p,q)|p \in [0,1], q \in [0,1]\}$ into itself:

$$s = \Pi(s) \tag{16}$$

where $s = (p_a, p_b)$ and $\Pi : S \to S$ is defined by

$$\Pi(s) = \begin{bmatrix} \pi_1(s_1, s_2) \\ \pi_2(s_1, s_2) \end{bmatrix} = \begin{bmatrix} \lambda_a(p_b) \\ \lambda_b(p_a) \end{bmatrix}. \tag{17}$$

Since Π is a continuous mapping from S to itself, and S is homoemorphic to the closed unit ball in R^2, the *Brouwer Fixed Point Theorem* implies that at least one Bayesian Nash equilibrium point always exists. Since the map Π is continuously differentiable, the Implicit Function Theorem implies that for any equilibrium (i.e. fixed point of Π) for which the 2×2 matrix $[I - \nabla\Pi]$ is non-singular, that any fixed $s^* = (p_a^*, p_b^*)$ will be at least locally a continuously differentiable function of (x_a, x_b) and θ^*. We can see this by converting the equilibrium problem to an equivalent problem of finding a zero to the function $F : S \to R^2$ given by

$$F(s|x_a, x_b, \theta^*) = s - \Pi(s|x_a, x_b, \theta^*). \tag{18}$$

By the Implicit Function Theorem, any solution s^* will be a continuously differentiable funtion of (x_a, x_b) and θ^* provided that the Jacobian of F with respect to s is nonsingular at the solution point s^*. Calculating the Jacobian, we have

$$\nabla F(s) = [I - \nabla\Pi], \tag{19}$$

where

$$\nabla\Pi(s|x_a, x_b, \theta^*) = \begin{bmatrix} \partial\pi_1(s_1,s_2)/\partial s_1 & \partial\pi_1(s_1,s_2)/\partial s_2 \\ \partial\pi_2(s_1,s_2)/\partial s_1 & \partial\pi_2(s_1,s_2)/\partial s_2 \end{bmatrix} = \begin{bmatrix} 0 & \partial\lambda_a(p_b, x_a)/\partial p_b \\ \partial\lambda_b(p_a, x_b)/\partial p_a & 0 \end{bmatrix} \tag{20}$$

Recall that a matrix is non-singular if and only if its determinant is non-zero. Thus, a sufficient condition for the Implicit function theorem to hold is that

$$det\,(I-\nabla\Pi) = 1-[\partial\lambda_a(p_b,x_a)/\partial p_b]\,[\partial\lambda_b(p_a,x_b)/\partial p_a] \neq 0. \quad (21)$$

Thus, a *regular equilibrium* is one where $I-\nabla\Pi(s^*)$ is nonsingular, and for such equilibria, the Implicit Function Theorem implies that (p_a,p_b) will be (locally) continuously differentiable functions of (x_a,x_b) and θ^*.

Figure 1 illustrates an equilibrium (which happens to be unique in this case) for the observed types $(x_a,x_b) = (.52,.22)$. The figure plots the two best response functions: λ_a as a function of p_b and is plotted as the dashed line (its domain is actually the vertical axis, so the graph of this function has be rotated clockwise by 90 degrees in figure 1), and λ_b as a function of p_a (the solid line in figure 1). The intersection of the two functions occurs at $(p_a^*,p_b^*) = (.7456,.6272)$. We can sense that this is a regular equilibrium, since small perturbations in x_a or x_b should result in small perturbations in λ_a and λ_b, causing only a small perturbation in the intersection point, (p_a^*,p_b^*).

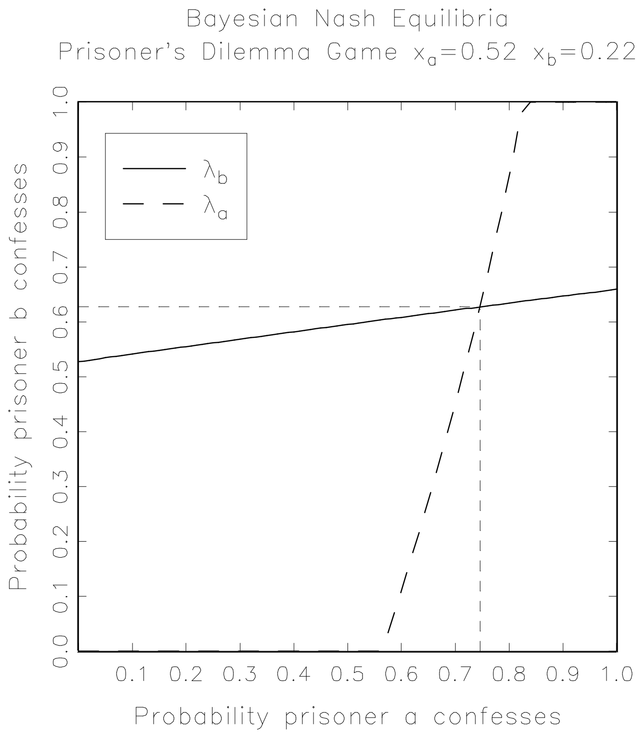

Figure 1: Bayesian Nash Equilibrium in Prisoners Dilemma when $(x_a,x_b) = (.52,.22)$

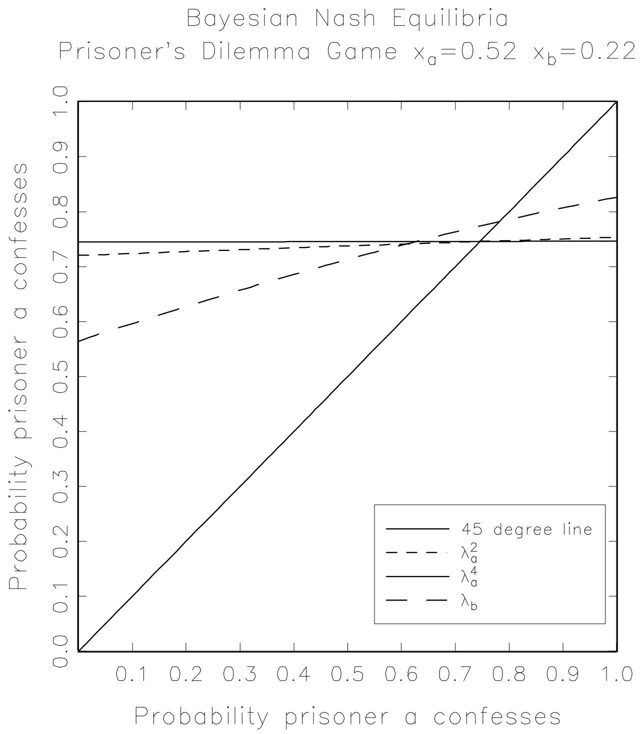

Figure 2: Computing the BNE via successive approximations, λ_a^{2i}, $i = 0,1,2,\ldots$ $(x_a, x_b) = (.52, .22)$

Figure 2 describes the equilibrium in terms of the *second order best response function* $\lambda_a^2(p_a, x_a, x_b)$. This function allows makes it easy to verify that the non-zero determinant condition, (21) does in fact hold in this case. We define the second order best response function λ_a^2 by substituting λ_b into λ_a resulting in a mapping from p_a to p_a, (i.e. $\lambda_a^2 : [0,1] \to [0,1]$).

$$\lambda_a^2(p_a, x_a, x_b) = \lambda_a(\lambda_b(p_a, x_b), x_a). \tag{22}$$

Similarly, we could define the 4^{th} *order best response function* λ_a^4 as

$$\lambda_a^4(p_a, x_a, x_b) = \lambda_a^2(\lambda_a^2(p_a, x_a, x_b), x_a, x_b), \tag{23}$$

and so forth. All (even) higher order best response functions can be defined similarly, $\lambda_a^{2i}(p_a, x_a, x_b)$, $i = 0,1,2,\ldots$. It is easy to see from Figure 2 that all of these higher order best response functions cross the 45 degree line at the equilibrium value for $p_a^* = .7456$. Indeed, with a little thought you should be able to see that the sequence $\{\lambda_a^{2i}(p_a, x_a, x_b)\}$ is simply the sequence of *successive approximations to the equilibrium p_a^* starting from the initial point p_a*. In this case we see that the sequence of successive approximations converges rapidly to the (unique) BNE of this game: by the $i = 2$nd successive approximations iteration, the iterated value is nearly equal to $p_a^* = .7456$ for any starting value p_a. This is shown in figure 2 by the flat graph of $\lambda_a^4(p_a, x_a, x_b)$ as a function of x_a: this represents the results of doing only 2 successive approximation steps from any starting point p_a. The fact that it is nearly a horizontal line of height $p_a^* = .7456$ means that the method of successive approximations converges remarkably quickly in this example. However since the function λ_a^2 is not necessarily a contraction mapping, the method of successive approximations need not always converge, especially in cases where there are multiple equilibria.

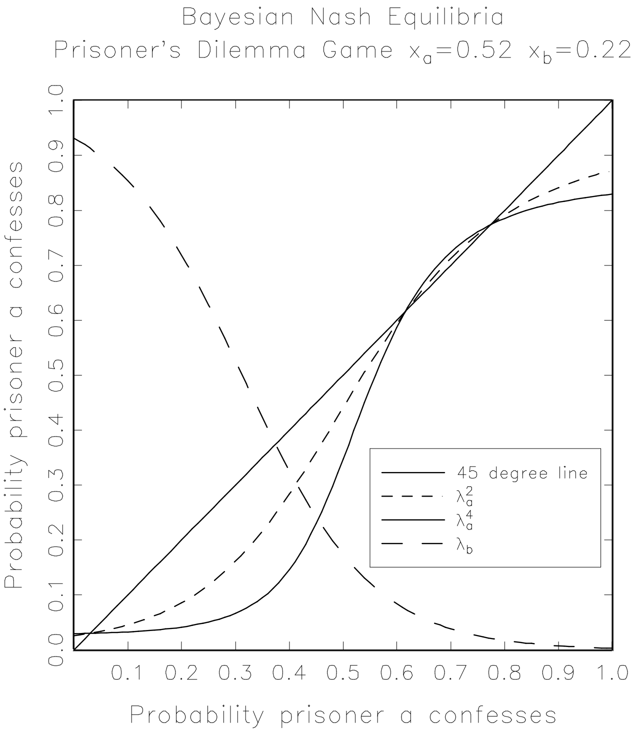

Figure 3: Example of Multiple Equilibria in the Prisoners Dilemma Game, $(x_a, x_b) = (.52, .22)$

Figure 3 illustrates a case where there are multiple equilibria to the Prisoners dilemma game, and in this case, we can verify that there are an odd number (3) of isolated equilibrium points. This equilibrium was calculated for a different value of the θ^* parameters, where the penalty to the prisoner who does not confess when the other does is not as severe as it is in the problem we were to solve for this problem set. In fact, for the equilibrium in figure 3, I set penalties for the prisoner who did not confess when the other did to be only .9 years in jail compared to the original value of 5 years in jail. With this lower "threat" to not confessing, we now see 3 equilibria in figure 3: a "good" equilibrium where both prisoner's don't confess with about 98% probability, and two "bad" equilibria where both confess with very high probabilities.

Parameters used to compute equilibria displayed in Figure 3

$$(\sigma_a,\sigma_b) = (.1,.1) \tag{24}$$

$$(\theta^a_{11},\theta^b_{11}) = (-2,-2) \tag{25}$$

$$(\theta^a_{00},\theta^b_{00}) = (-1,-1) \tag{26}$$

$$(\theta^a_{10},\theta^b_{01}) = (-1/2,-1/2) \tag{27}$$

$$(\theta^a_{01},\theta^b_{10}) = (-.9,-.9) \tag{28}$$

It is easy to see that the "middle" equilibrium in Figure 3 is not stable under successive approximations: the iterates will tend to go towards either the good "low confession" equilibrium, or the worst "high confession" equilibrium. If one wanted to compute the middle equilibrium, then some algorithm other than successive approximations would be required.

Note that all 3 of the equilibria of the problem illustrated in figure 3 are regular, i.e. whenever the 2nd order best response function crosses the 45 degree line, it does so from either above or below. This ensures that the slope of the second order best response function is either strictly greater or less than 1, and you should be able to see that this implies that the regularity condition in equation (21) holds. It follows that all 3 of these equilibria are "regular" and as a result, small shifts in (x_a,x_b) will result in small shifts in λ^2_a, and thus, small shifts in any of the equilibrium points p^*_a. In such cases "comparative static" calculations are possible. We can ask, for example, by how much small changes in (x_a,x_b) or θ^* will affect a particular equilibrium point. Thus, for example, if we wish to compute the derivative of p^*_a with respect to θ^*, we have, via the Implicit Function Theorem,

$$\partial p^*_a/\partial\theta^* = \frac{\partial\lambda^2_a(p^*_a,x_a,x_b,\theta^*)/\partial\theta^*}{1-\partial\lambda^2_a(p^*_a,x_a,x_b)/\partial p_a}. \tag{29}$$

We can evaluate this expression at any one of the 3 equilibrium points in figure 3 and determine the effect of changes in the parameters on the equilibrium point. We will see, it is important to have this to help provide analytical derivatives for the nested fixed point maximum likelihood algorithm we will describe shortly.

Figure 4 presents an example of an "irregular" equilibrium point. This example was produced by increasing slightly the penalty to the prisoner who does not confess when the other does confess. In this case I have chosen $\theta^a_{1,0} = \theta^b_{0,1} = -.98$, keeping all of the other parameters fixed at the values chosen in equations (24) to (28) above. In this case we see that there only only *two* equilibria of this game.

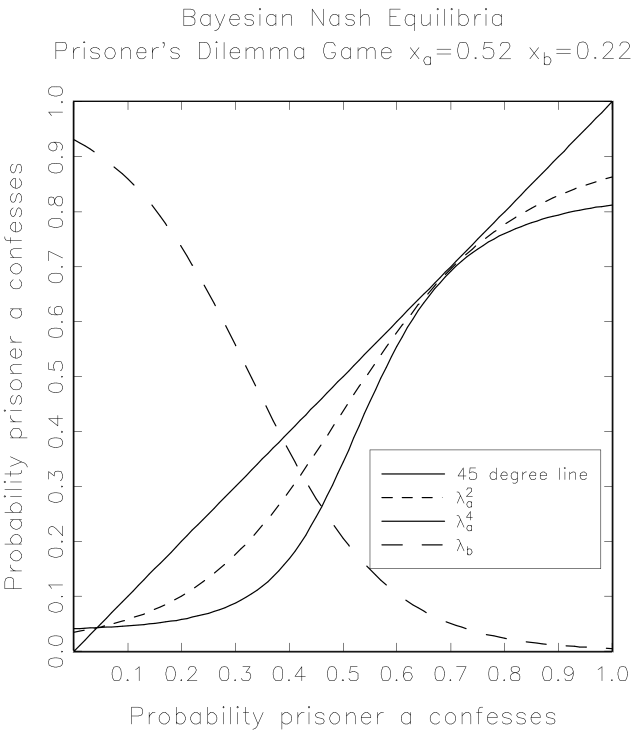

Figure 4: Example of an Irregular Equilibrium in the Prisoners Dilemma Game, $(x_a, x_b) = (.52, .22)$

Does this violate the index theorem we discussed above, which states that there should almost always be an odd number of equilibria? The answer is no, since the 2nd equilibrium, which occurs where λ_a^2 is just tangent to the 45 degree line is not a regular equilibrium. In fact, this is an example of an irregular equilibrium point, where the determinant condition in equation (21) fails to hold. This is a case where the Jacobian matrix $I - \nabla\Pi$ is singular at the second "high confession probability" equilibrium point. This implies that the sufficient condition for the Implicit Function Theorem will not hold at this second, irregular, equilibrium point, and so small changes in (x_a, x_b) or in θ^* can result in large, discontinuous changes in the equilibrium.

Figure 5 illustrates the discontinuous impact of a slight change in θ^* on the set of equilibria to this game. I decreased $\theta^a_{1,0}$ and $\theta^b_{0,1}$ from $-.98$ to -1.1. We see that there is only a single "low confession" equilibrium now. It is possible to create similar examples where small changes in x_a or x_b lead to discontinuous changes in the set of equilibria, in some cases causing one of the equilibria to "disappear". Intuitively, irregular equilibria are "rare" in the sense that any small perturbation in the parameters causes the irregular equilibria to disappear, resulting in a case where there are an odd number of regular equilibria. Thus, the index theorem is correct, and it tells us that "'almost all" equilibria are such that each equilibrium point is a regular equilibrium. The idea of "almost all" can be be formalized by specifying how the equilibria are parameterized or shifted by parameters such as θ^* or variables such as (x_a, x_b). If there are values of θ^* or (x_a, x_b) that happen to result in an irregular equilibrium, then such values have measure zero relative to Lebesgue measure, which is another way of formalizing the statement that "almost all" equilibria are regular.

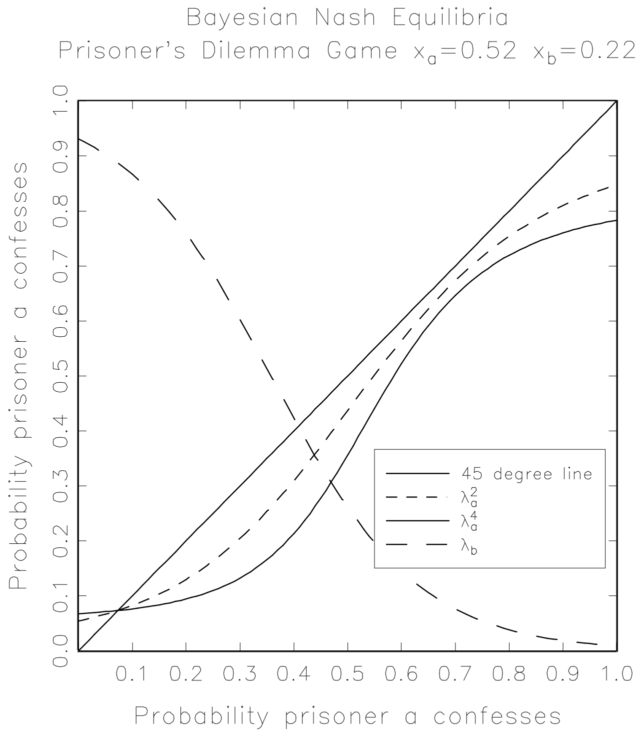

Figure 5: Effect of Small Parameter Change on an Irregular Equilibrium, $(x_a, x_b) = (.52, .22)$

Does this mean that we need not worry about discontinuities in the equilibrium mappings $p_a^*(x_a, x_b)$ and $p_b^*(x_a, x_b)$? The answer is no: *we need to be careful and use a consistent equilibrium selection rule, since our econometric estimation algorithm will be computing equilibria for many different* (x_a, x_b) *and* θ *values.* Even though the set of parameter values that result in "irregular" equilibria has measure zero, it is possible that the estmation algorithm, in the course of searching over different θ values, may jump from one region of the parameter space where there is only 1 equilibrium, to another where there are 3 or more equilibria. Or alternatively, it might happen that there is a unique equilibrium for certain values of the "covariates" (x_a, x_b), but multiple equilibria for other values. In such circumstances it is important to employ an consistent "equilibrium selection rule" otherwise problems can result from "extraneous" discontinuities in the likelihood function due to from arbitrary jumps from one equilibrium to another as the parameter vector θ, or the covariates (x_a, x_b) vary in the course of the estimation.

In this example, I have used the following equilibrium selection rule: *always choose the equilibrium that results in the highest probability of confession for prisoner* **a**. This is a well defined equilibrium selection rule, since the set of equilibria (for p_a^*) can always be ordered from the lowest equilibrium p_a^* to the highest. If the estimation algorithm happens to move from a part of the parameter space where there are three equilibria to a parameter where there is only one equilibrium, then this equilibrium selection rule will always select a consistent equilibrium point (i.e. the one with the highest value of p_a^*) and there will be no "extraneous" and discontinous jumping from one equilibrium to another. An example of such "extraneous discontinuities" would be an equilibrium algorithm that might (apparently at random) result in convergence to the equilibrium with the lowest value of p_a^* for some values of (x_a, x_b) or θ, and to the equilibrium with the highest value of p_a^* for other values. These discontinuities can result in an "irregular" likelihood function, and destroy the traditional properties of maximum likelihood estimation (e.g. consistency, asymptotic normality, asymptotic efficiency), all of which depend on general regularity conditions for the likelihood function as a function of θ, in particular, a key assumption is that the likelihood function is continuously differentiable in θ.

I do note that even in cases where there are "non extraneous" discontinuities resulting from values of the covariates (x_a, x_b) which result in unique equilibria and other values where there are multiple equilibria, typically the set of (x_a, x_b) and θ values where such discontinuities occur can be show to have measure zero. If this is the case, and if there is sufficient variation in (x_a, x_b) so that one can show *expected likelihood function* $E\{L(\theta)\}$ is continuously differentiable in θ, then the standard properties of the maximum likelihood estimator can typically be shown to hold, *even though in finite samples, the finite sample likelihood function* $L_N(\theta)$ *can be discontinuous in* θ. Ideas from the literature on *empirical processes* and *simulation estimation* can be adapted to show that in the limit, the maximum likelihood estimator enjoys desirable properties, despite the discontinuities in the finite sample likelihood functions, $L_N(\theta)$. The technical details associated with establishing the necessary regularity and convergence conditions are beyond the scope of this problem set, but we will provide evidence shortly, that such regularity conditions are in fact satisfied in this prisoners dilemma example.

3 Solution/Simulation Algorithm Details

In the directory containing this answer sheet, I have left a number of Gauss programs that I used to compute equilibria for this problem. I solved for equilibria by searching for the fixed point of the second order best response function λ_a^2. Recall that my equilibrium selection rule is to use the equilibrium with the highest value of p_a^*. Since it is typically the case that $\lambda_a^2(1,x_a,x_b) < 1$ if $\sigma_a > 0$, it is not hard to show that a sequence of successive approximations starting from $p_a = 1$ will be decreasing and will converge to the largest equilibrium point p_a^*. Of course, successive approximations is given by

$$p_a^{t+1} = \lambda_a^2(p_a^t, x_a, x_b) = \lambda_a^{2(t+1)}(1, x_a, x_b). \tag{30}$$

However to speed up the solution, and produce formulas for the derivatives of (p_a^*, p^*b) with respect to θ, after getting sufficiently close to p_a^* using successive approximations, I switch to Newton's method to produce a super accurate estimate of the equilibrium point. Newton's algorithm is given by

$$p_a^{t+1} = p_a^t - \frac{p_a^t - \lambda_a^2(p_a^t, x_a, x_b)}{1 - \partial\lambda_a^2(p_a^t, x_a, x_b)/\partial p_a}, \tag{31}$$

where, by the chain rule, we have

$$\partial\lambda_a^2(p_a, x_a, x_b) = [\partial\lambda_a(\lambda_b(p_a, x_b), x_a)/\partial p_b]\,[\partial\lambda_b(p_a, x_b)/\partial p_a]\,. \tag{32}$$

Using (12) and (13), the derivatives are given by

$$\begin{aligned} \partial\lambda_a(p_b, x_a)/\partial p_b &= -\lambda_a(p_b, x_a)[1 - \lambda_a(p_b, x_a)]x_a\beta_a \quad (33) \\ \partial\lambda_b(p_a, x_b)/\partial p_a &= -\lambda_b(p_a, x_b)[1 - \lambda_b(p_a, x_b)]x_b\beta_b, \quad (34) \end{aligned}$$

where

$$\begin{aligned} \beta_a &= (\theta^a_{0,1} - \theta^a_{1,1} + \theta^a_{1,0} - \theta^a_{0,0})/\sigma_a \quad (35) \\ \beta_b &= (\theta^b_{1,0} - \theta^b_{1,1} + \theta^b_{0,1} - \theta^b_{0,0})/\sigma_b. \quad (36) \end{aligned}$$

As a nice by-product of the Newton iterations, we can use the converged probability values to compute expressions for the partial derivatives $\partial p_a^*(x_a, x_b, \theta)/\partial\theta$ and $\partial p_b^*(x_a, x_b, \theta)/\partial\theta$. The formula for the former is given in equation (29) above, and for the latter, we employ the chain rule:

$$\partial p_b^*(x_a, x_b, \theta)/\partial\theta) = [\partial\lambda_b(p_a^*, x_b)/\partial p_a]\,\partial p_a^*(x_a, x_b, \theta)/\partial\theta. \tag{37}$$

The Gauss procedure `equil.g` computes the equilbrium pair (p_a^*, p_b^*) via the combined successive approximations, Newton iteration algorithm outlined above, and the Gauss procedure `dequil.g` does this plus it computes the derivatives of (p_a^*, p_b^*) with respect to θ using the formulas given above.

Note that both `equil.g` and `dequil.g` make repeated calls to the best response functions λ_a and λ_b. These best response functions are coded in the Gauss procedures `bra.g` and `brb.g`. Both of these procedures have 2 return values: the best reponse probability, and the derivative of the best reponse probability with respect to its argument, i.e. `bra.g` returns $\lambda_a(p_b, x_a)$ and $\partial\lambda_a(p_b, x_a)/\partial p_b$. The product of these two derivatives can then be used to compute the derivative of the second order best response function given in formula (32) above.

Figure 6 below shows the result of computing the equilibria of the prisoners dilemma game at all possible (x_a, x_b) combinations, for the original parameter values for θ^* given in equation (8). No cases of multiple equilibria were encountered for any values of (x_a, x_b), so the need for an equilibrium selection rule is superfluous — at least for θ^*.

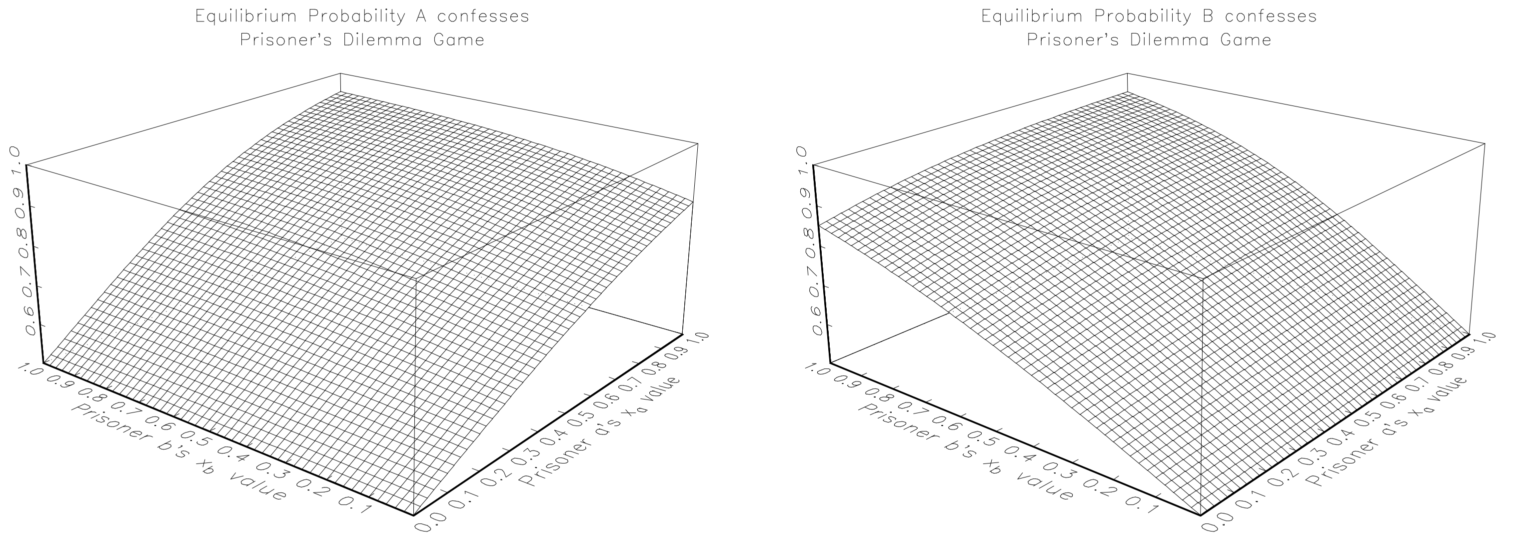

Figure 6: Set of All Equilibria to the Prisoners Dilemma Game

We see that, as we would expect, the equilbrium probability of confessing is monotonically increasing in x_a and x_b. These variables govern relative weight of the sentencing parameters, i.e. the θ_{ij}^a and θ_{ij}^b values, relative to the "idiosyncratic" random factors affecting whether to confess or not (i.e. the ε_a and ε_b vectors). However note that when $x_a = 0$, then the sentencing parameters receive no weight for prisoner **a** and his/her behavior appears completely random, regardless of the value of x_b. Thus, when $x_a = 0$, it follows that $p^*_a(0,x_b) = .5$ for all x_b. Symmetrically we have $p^*_b(x_a,0) = .5$ for all x_a.

When $x_a > 0$ and $x_b > 0$, the equilibrium probability of confessing is increasing in both observed types: the higher the other prisoner's type is, the more likely it is that the prisoner will confess. The rationale for this is straightforward: when $x_a = x_b = 1$, it is easy to see that it is a *dominant strategy* to confess. This is illustrated by the expected payoff functions (excluding ε_a and ε_b which are set to zero) in figure 7 below. We see that regardless of the probability that the other prisoner will confess, it is always in each prisoner's interest to confess. Even when $x_a < 1$ and $x_b < 1$, it is still better to confess when the ε_a and ε_b shocks are zero – except in the case where $x_a = 0$ or $x_b = 0$, in which case both prisoners are indifferent about confessing or not. Thus, the higher the value of x_b, the higher weight prisoner **b** places on confessing, and as we see from the figures below, the higher is the relative gain in expected utility for prisoner **a** to confess rather than not confess. This leads the equilibrium probabilities $p^*_a(x_a,x_b)$ and $p^*_b(x_a,x_b)$ to be strictly increasing in both of their arguments except in the case where x_a or x_b is zero.

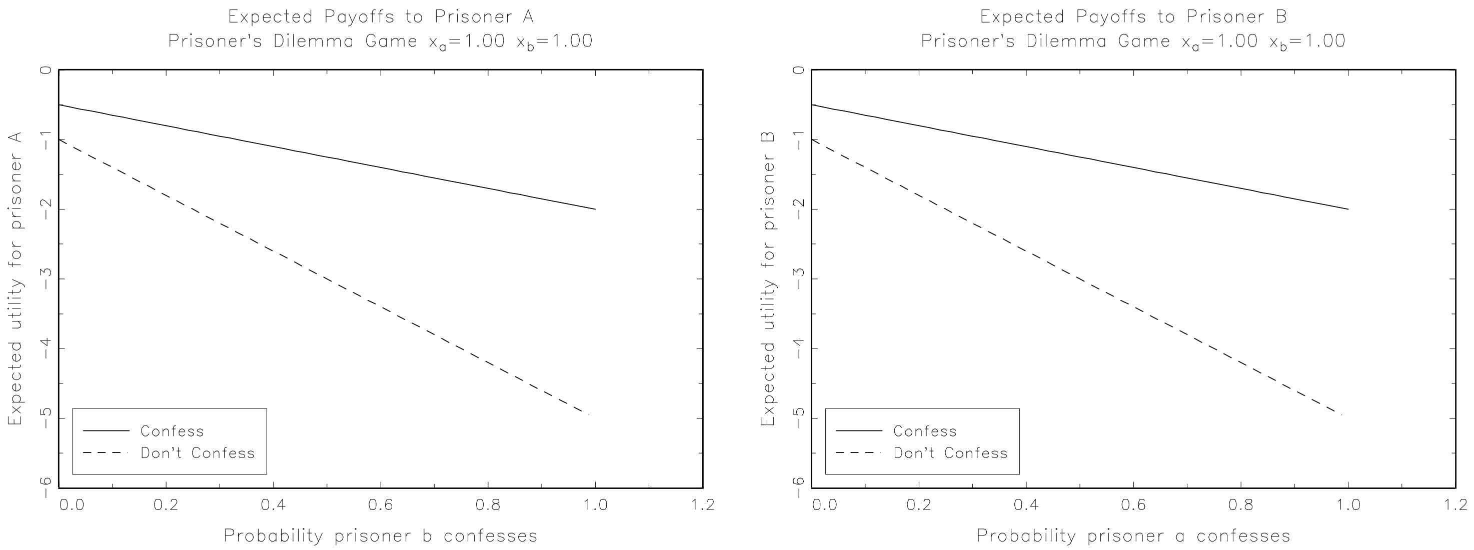

Figure 7: Expected Payoffs to Confessing as a Function of the Probability the Other Prisoner Confesses

The presence of the additive idiosyncratic errors $(\varepsilon_a, \varepsilon_b)$ implies that neither prisoner will always confess in equilibrium, however the variance of these idiosyncratic shocks when $\sigma_a = \sigma_b = 1$ is not so large to override the incentive to confess, and this is why there is a unique equilibrium for these parameter settings. In this sense, I chose an "easy" problem in order to avoid complications associated with multiple equilibria.

4 Identification of the Parameters

As I noted above, there are potentially 10 unknown parameters to be estimated in this problem: the eight values of $\{\theta^a_{ij}, \theta^b_{ij}\}$ plus the two scale parameters, (σ_a, σ_b). However it is clear that not all of these parameters are identified. In particular, multiplying all of these parameters by a positive scale factor k does not affect the decisions of either prisoner, and hence the equilibria are invariant to arbitrary scalings. Similarly, adding a fixed constant to all utility values shifts expected utility by the same amount and thus cannot affect the decision of either prisoner.

Definition: We say two parameter values θ and θ^* are *observationally equivalent* if they both have the same set of equilibria for all values of (x_a, x_b).

Theorem: Parameters θ and $\overline{\theta}$ are observationally equivalent if and only if the following conditions hold:

$$\begin{aligned}
\frac{1}{\overline{\sigma}_a}\left[\overline{\theta}^a_{0,0} - \overline{\theta}^a_{1,0}\right] &= \frac{1}{\sigma_a}\left[\theta^a_{0,0} - \theta^a_{1,0}\right] \\
\frac{1}{\overline{\sigma}_a}\left[\overline{\theta}^a_{0,1} - \overline{\theta}^a_{1,1}\right] &= \frac{1}{\sigma_a}\left[\theta^a_{0,1} - \theta^a_{1,1}\right] \\
\frac{1}{\overline{\sigma}_b}\left[\overline{\theta}^b_{0,0} - \overline{\theta}^b_{0,1}\right] &= \frac{1}{\sigma_b}\left[\theta^b_{0,0} - \theta^b_{0,1}\right] \\
\frac{1}{\overline{\sigma}_b}\left[\overline{\theta}^b_{1,0} - \overline{\theta}^b_{1,1}\right] &= \frac{1}{\sigma_b}\left[\theta^b_{1,0} - \theta^b_{1,1}\right].
\end{aligned} \tag{38}$$

Proof: Suppose the conditions (38) hold. Then we will show that θ and $\overline{\theta}$ have the same set of equilibria for all values of (x_a, x_b). To see this, consider the formula for the best response function, $\lambda_a(p_b, x_a)$ and $\lambda_b(p_a, x_b)$ in equations (12) and (13). It is easy to see that if (38) holds, then for any $(p_a, p_b) \in [0,1] \times [0,1]$ and any $(x_a, x_b) \in [0,1] \times [0,1]$ we have:

$$\begin{aligned}
\lambda_a(p_b, x_a, \overline{\theta}) &= \lambda_a(p_b, x_a, \theta) \\
\lambda_b(p_a, x_b, \overline{\theta}) &= \lambda_b(p_a, x_b, \theta).
\end{aligned} \tag{39}$$

This holds, for example, in the case of λ_a since it depends on its arguments (p_b, x_a) via the binomial logit formula in equation (12). λ_a depends on x_a in two places: 1) by itself, multiplied by a coefficient equal to the formula in either the right or left hand side of the first equation in (38), and 2) as a product of x_a and p_b multiplied by a coefficient equal to the difference in the left (or right) hand sides of the first equation in (38) less the left (or right) hand side formula in the second equation of (38). These are the only two ways in which x_a and p_b enter the best response probability λ_a, and since the observational equivalence condition (38) implies the coefficients of x_a and $x_a p_b$ are the same for parameters θ and $\hat{\theta}$, it follows that the top equality in equation (39) must hold for all values of (p_b, x_a). A similar argument establishes that the second equation in (39) holds for all (p_a, x_b). Since θ and $\overline{\theta}$ imply identical best response functions, it follows immediately that the equilibria implied by these two different parameter values must be the same. Conversely, it is not hard to show, using (12) and (13) that if θ and $\overline{\theta}$ have the same set of equilibria, then they must have the same best response functions, λ_a and λ_b. This holds since

$$\begin{aligned} p_a(x_a, x_b, \overline{\theta}) &= p_a(x_a, x_b, \theta) \\ p_b(x_a, x_b, \overline{\theta}) &= p_b(x_a, x_b, \theta) \end{aligned} \tag{40}$$

which implies that

$$\begin{aligned} p_a(x_a, x_b, \overline{\theta}) &= \lambda_a(p_b(p_a(x_a, x_b, \overline{\theta}), x_a, \overline{\theta}) \\ &= \lambda_a(p_b(p_a(x_a, x_b, \theta), x_a, \theta). \end{aligned} \tag{41}$$

Using the binomial logit formula for the best response probability in equation (12), it is easy to see that the equation above will hold only if the coefficients of x_a and $p_b x_a$ are the same. But these coefficients cannot be the same unless the top two equations in (38) hold. A similar argument establishes that observational equivalence of θ and $\overline{\theta}$ implies that the bottom two equations in (38) must hold as well.

Comment: From Theorem 1, we see that there are basically only four identified parameters in the prisoners dilemma model. These are the four "normalized" coefficients defined on the right or left hand sides of equation 38). We see that these coefficients are the rescaled losses from not confessing relative to confessing for both prisoners, under two cases: 1) where the other prisoner confesses, and 2) where the other prisoner does not confess. In particular, we cannot infer the absolute values of $\theta^a_{0,0}$ and $\theta^a_{1,0}$, but only the rescaled difference in their values: $(\theta^a_{0,0} - \theta^a_{1,0})/\sigma_a$. The rescaling is a normalization that is clearly needed, since as we noted above, if we were to multiply σ_a and all $\{\theta^a_{i,j}\}$ by a positive constant, we would not change prisoner **a**'s decision. The fact that we can only indentify the rescaled difference in the coefficients $(\theta^a_{0,0} - \theta^a_{1,0})/\sigma_a$ is a consequence of Theorem 1. It tells us that it is really only necessary to specify these 4 normalized coefficients, since any specification for the full set of 10 unknown coefficients that implies the same four normalized parameter values will result in the same predicted behavior in this game.

Comment: Note that Theorem 1 tells us that it is not necessary to impose symmetry in payoffs in order to identify all of the behavioral-relevant parameters of this game. In particular, symmetry implies that instead of four coefficients, it would only be necessary to estimate two coefficients, since symmetry implies that the coefficients of variables $(x_a, x_a p_b)$ in λ_a are the same as the coefficients of the variables $(x_b, p_a x_b)$ in λ_b. As such, the hypothesis of symmetry is not necessary for identification and is therefore a testable hypothesis: we can estimate an "unrestricted" specification and then test whether the symmetry conditions hold:

$$\begin{aligned} \alpha_a &= \alpha_b \\ \beta_a &= \beta_b, \end{aligned} \tag{42}$$

where β_a is the coefficient of $x_a p_b$ in the formula for λ_a in (12), and β_b is the coefficient of $x_b p_a$ in the formula for λ_b in equation (13), and α_a and α_b are the coefficients of x_a and x_b, respectively, in these two best response formulae, or explicitly,

$$\begin{aligned} \alpha_a &= (\theta^a_{0,0} - \theta^a_{1,0})/\sigma_a \\ \alpha_b &= (\theta^b_{0,0} - \theta^b_{0,1})/\sigma_b. \end{aligned} \tag{43}$$

In terms of equations (38) defining the four identified "parameters" of this model, we see that the first two equations are equal to α_a and α_b, respectively, and the last two equations equal $\beta_a - \alpha_a$ and $\beta_b - \alpha_b$, respectively.

We will treat the 4×1 identified parameter, i.e. the vector of normalized coefficients $(\alpha_a, \alpha_b, \beta_a, \beta_b)$ as the "parameters of interest" of the prisoners dilemma model. We have also imposed the symmetry condition, so it follows that there are only two parameters to estimate in this case: $\theta = (\alpha, \beta)$ where α is the coefficient for x_a and β is the coefficent of $x_a p_b$ in $\lambda_a(p_b, x_a)$, and α is also the coefficient of x_b and β is the coefficient of $x_b p_a$ in $\lambda_b(p_a, x_b)$. Given the true parameter values specified in equation (8), it follows that the true values of the identified, normalized parameters are

$$\begin{aligned} \alpha^* &= -0.5 \\ \beta^* &= -2.5. \end{aligned} \tag{44}$$

In the next section I will propose and illustrate two estimators for these parameters.

5 Estimation of the Parameters

The most efficient method of estimation is *full information maximum likeihood* (FIML). However this method is also the most computationally demanding, since it requires computing the Bayesian Nash equilibrium for every observation $i = 1, \ldots, N$ in order to take account of the different realized values of (x^i_a, x^i_b)

for each separate observation, which requires N separate solutions for the BNE each time the likelihood is evaluated at any trial θ value. Thus, computation of the FIML estimates requires a *nested fixed point algorithm.* That is, nested within each evaluation of the likelihood function is the computation of N fixed point problems to compute the Bayesian Nash equilibrium for each of the N observations in the data set.

An alternative, less efficient, but consistent and asymptotically normally distributed estimator is a *semi-parametric two step estimator.* This estimator is less computationally demanding that the full information maximum likelihood. However this estimator requires a "first stage" estimate of the equilibrium probability of confessing $p_b(x_a,x_b)$ for player **b**. As we will see, if we have such an estimate, we can estimate the unknown parameters (α,β) of prisoner **a**'s best response function λ_a (which by the symmetry assumption are the same as the coefficients of prisoner **b**'s best response function λ_b), by treating prisoner **a**'s observations "as if" he/she were playing a single agent decision problem, i.e. a "game against nature" where "nature's move" (to confess or not confess) is governed the conditional probability function $p_b(x_a,x_b)$.

Consider the FIML estimator first. The (log) likelihood function is given by

$$\begin{aligned} L_n(\theta) \quad &= \tfrac{1}{N}\textstyle\sum_{i=1}^N \quad d_a^i \log(p_a(x_a^i,x_b^i,\theta)) + (1-d_a^i)\log(1-p_a(x_a^i,x_b^i,\theta)) + \\ &\qquad\qquad\quad d_b^i \log(p_b(x_a^i,x_b^i,\theta)) + (1-d_b^i)\log(1-p_b(x_a^i,x_b^i,\theta)), \end{aligned} \tag{45}$$

where (d_a^i,d_b^i) are the decisions taken by the two prisoners in observation i, and (x_a^i,x_b^i) are their corresponding observed types, and (p_a,p_b) are the Bayesian Nash equilibrium probabilities (calculated separately for each of the N observations). Since these latter probabilities are the fixed point to the equilibrium problem, the FIML estimator requires a *nested fixed point algorithm.* This is a standard maximization algorithm (to maximize L_N as a function of θ), except that for each observation i, the algorithm must call a nested subroutine to compute the fixed point to determine the probabilities $(p_a(x_a^i,x_b^i,\theta), p_b(x_a^i,x_b^i,\theta))$ necessary to evaluate the likelihood of the actual observed decisions (d_a^i,d_b^i) occurring. Thus, we can summarize the FIML estimator $\hat{\theta}_{FIML}$ based on a nested fixed point maximum likelihood algorithm as

$$\begin{aligned} \hat{\theta} = \underset{\theta}{argmax}\, L_n(\theta) = \frac{1}{N}\sum_{i=1}^{N} \qquad & d_a^i \log(p_a(x_a^i,x_b^i,\theta)) + (1-d_a^i)\log(1-p_a(x_a^i,x_b^i,\theta)) + \\ & d_b^i \log(p_b(x_a^i,x_b^i,\theta)) + (1-d_b^i)\log(1-p_b(x_a^i,x_b^i,\theta)), \\ & \text{subject to:} \\ & p_a(x_a^i,x_b^i,\theta) = \lambda_a(p_b(x_a^i,x_b^i,\theta),x_a,\theta) \\ & p_b(x_a^i,x_b^i,\theta) = \lambda_b(p_a(x_a^i,x_b^i,\theta),x_b,\theta). \end{aligned} \tag{46}$$

Computation of the nested fixed point maximum likelihood estimator (or FIML-NFXP, for short) is assisted by using Newton's method to compute both the fixed points (p_a, p_b) and the derivatives of these fixed points with respect to θ, $(\partial p_a/\partial\theta, \partial p_b/\partial\theta)$. Otherwise a non-derivative maximization algorithm would need to be employed, and two extra likelihood function evaluations would need to be performed at each iteration of the FIML-NFXP algorithm in order to compute numerical derivatives with respect to (α, β). Although it is possible also to compute second derivatives, in practice I have found that the *BHHH algorithm* provides acceptably rapid convergence using only first derivative information. Recall that BHHH uses a search direction given by

$$s_t = [H_N(\theta_t)]^{-1} \partial L_N(\theta_t)/\partial\theta \tag{47}$$

where H_N is the average outer product of the gradients of each term of the likelihood function. Thus, if we have

$$L_N(\theta) = \frac{1}{N} \sum_{i=1}^{N} \log(f(x_i, \theta)) \tag{48}$$

then

$$H_N(\theta) = \frac{1}{N} \sum_{i=1}^{N} [\partial \log(f(x_i, \theta))/\partial\theta] \left[\partial \log(f(x_i, \theta))/\partial\theta'\right]. \tag{49}$$

The idea behind the semiparametric two step estimator is that if we knew prisoner **b**'s probability of confessing (i.e. the Bayesian Nash equilibrium probability $p_b(x_a, x_b, \theta^*)$), then we estimate the coefficients (α_a, β_a) from prisoner **a**'s best response probability, the logit formula in equation (12), as a simple binomial logit estimation by substituting the (known) value of $p_b = p_b(x_a, x_b, \theta^*)$ into the logit formula, as if it were "data" similar to (x_a, x_b). By doing this we would completely bypass the need to solve the game to compute (p_a, p_b). However we do not know (p_a, p_b) because these probabilities depend on the unknown parameter θ^* that we are trying to estimate. However under the hypothesis that the data $\{(d_a^i d_b^i), (x_a^i, x_b^i)\}_{i=1}^N$ are generated as realizations of equilibrium decisions by the two prisoners, it is possible to non-parametrically estimate the equilibrium probability functions (p_a, p_b). Call these non-parametrically estimated probability functions $(\hat{p}_a, \hat{p}_b)$. As $N \to \infty$, the properties of non-parametric estimators guarantee that with probability 1 we have

$$\begin{aligned} \lim_{N\to\infty} \hat{p}_a(x_a, x_b) &= p_a(x_a, x_b) \\ \lim_{N\to\infty} \hat{p}_b(x_a, x_b) &= p_b(x_a, x_b) \end{aligned} \tag{50}$$

for each $(x_a, x_b) \in [0,1] \times [0,1]$. Under appropriate regularity conditions (see Bajari, Hong, Krainer and Nekipelov (2005) for details), we can obtain consistent estimates of (α_a, β_a) by maximizing a binomial logit likelihood for prisoner **a** using the estimated equilibrium probability function $\hat{p}_b(x_a, x_b)$ and treating it as if it were the true equilibrium probability function $p_b(x_a, x_b, \theta^*)$.

Thus, the first step in the two step estimator is to estimate p_b non-parametrically. The second step is to compute estimates $(\hat{\alpha}_a, \hat{\beta}_a)$ by maximizing the binomial logit likelihood providing prisoner **a**'s best response to prisoner **b**'s behavior, as represented by the estimated equilibrium probability $\hat{p}_b$:

$$\begin{aligned}(\hat{\alpha}_a, \hat{\beta}_a) \quad = \quad & \underset{(\alpha,\beta)}{argmax} \frac{1}{N} \sum_{i=1}^{N} d_a^i \log(\lambda_a(\hat{p}_b(x_a^i, x_b^i), x_a^i, \alpha, \beta)) + \\ & (1 - d_a^i) \log(1 - \lambda_a(\hat{p}_b(x_a^i, x_b^i), x_a^i, \alpha, \beta)),\end{aligned} \tag{51}$$

where $\lambda(p_b, x_a, \alpha, \beta)$ is prisoner **a**'s best response probability, i.e. the binomial logit formula

$$\lambda_a(p_b, x_a, \alpha, \beta) = \frac{1}{1 + \exp\{x_a\alpha + x_a p_b \beta\}}. \tag{52}$$

Under appropriate regularity conditions (again see Bajari, Hong,

Krainer and Nekipelov, 2005 for details), this two step estimator $(\hat{\alpha}_a, \hat{\beta}_a)$ will be a consistent and asymptotically normal estimator of the true values, (α_a^*, β_a^*), assuming the underlying model is correctly specified.

The beauty of the two step estimator is that it make it unecessary to actually solve for the equilibrium of the game, and having to worry about the issues of "equilibrium selection" that were an important detail to be considered when estimating the parameters using the FIML-NFXP algorithm. In effect, if we hypothesize that the observations are are realizations of some equilibrium of the game, the data already provide us with the relevant equilibrium selection rule, and by estimating the equilibrium probability functions (p_a, p_b) we are able to use the same equilibrium selection that the players of the actual game "used" to determine their behavior. The drawback of the two-step estimator, however, is that it can be considerably less efficient than the maximum likelihood estimator. The reason is that there are implicit "cross equation restrictions" that the FIML-NFXP estimator uses, namely, the fact that θ appears both in the best response functions (λ_a, λ_b) and in the equilibrium probability functions (p_a, p_b). The two step estimator ignores these cross equation restrictions. Also, the two step estimator only estimates the parameters for one of the prisoners — prisoner **a** in this case — and there is a loss of efficiency in not also using the decisions of prisoner **b** as an additional source of information to estimate (α_b, β_b). Under the symmetry assumption, $(alpha_a, \beta_a) = (\alpha_b, \beta_b)$, and more efficient estimates would result from pooling the data on both prisoners and estimating the joint likelihood function

$$(\hat{\alpha}, \hat{\beta}) \quad = \quad \underset{(\alpha,\beta)}{argmax} \frac{1}{N} \sum_{i=1}^{N} d_a^i \log(\lambda_a(\hat{p}_b(x_a^i, x_b^i), x_a^i, \alpha, \beta)) +$$

$$
\begin{aligned}
&(1-d_a^i)\log(1-\lambda_a(\hat{p}_b(x_a^i,x_b^i),x_a^i,\alpha,\beta)) \\
&d_b^i\log(\lambda_b(\hat{p}_a(x_a^i,x_b^i),x_b^i,\alpha,\beta))+ \\
&(1-d_b^i)\log(1-\lambda_b(\hat{p}_a(x_a^i,x_b^i),x_b^i,\alpha,\beta)).
\end{aligned} \tag{53}
$$

It remains to specify how to estimate the equilibrium probability functions $(\hat{p}_a, \hat{p}_b)$ nonparametrically. There are many ways that this can be done, including estimating a "flexible functional form" (binary logit) model for p_b using series expansions of polynomials in the two variables (x_a, x_b). Another approach is to estimate a *local linear model.* Still another approach is to estimate $\hat{p}_b$ via non-parametric *kernel density estimation.*

I will explain the latter method first, and then explain how to estimate $\hat{p}_b$ via local linear models. We can write the true equilibrium function $p_b(x_a, x_b, \theta^*)$ as a *regression* as follows:

$$
d_b = p_b(x_a, x_b, \theta^*) + \eta \tag{54}
$$

where η is an "error term" which equals $1 - p_b(x_a, x_b, \theta^*)$ if $d_b = 1$ or $-p_b(x_a, x_b, \theta^*)$ if $d_b = 0$. The fact that the error term in this "regression" takes only two possible values for any fixed value of the covariates (x_a, x_b) is immaterial: it is still the case that $p_b(x_a, x_b, \theta^*)$ is the conditional expectation of the random variable $\tilde{d}_b$ given (x_a, x_b)

$$
p_b(x_a, x_b, \theta^*) = E\{\tilde{d}_b | x_a, x_b\}. \tag{55}
$$

Thus, we can use a huge arsenal of methods from the literature on *non parametric regression* to estimate p_b non-parametrically (i.e. without making any assumptions on its functional form, i.e. how (x_a, x_b) enter as arguments). Below I summarize some ideas from this literature, but a good introduction to these methods is Wolfgang Haerdle's (2004) *Applied Nonparametric Regression* (available online at the web site `http://www.quantlet.com/mdstat/scripts/anr/html/anrhtml.html`).

A classic and commonly used method for non-parametric regression is *kernel smoothing.* A kernel smoother is a special case of *local averaging* where we estimate $\hat{p}_b(x_a, x_b)$ as a simple average of the observations, but placing greater weight on observations (d_b^i, x_a^i, x_b^i) which are "close" to the given point of interest, (x_a, x_b) at which we are trying to evaluate $\hat{p}_b(x_a, x_b)$

$$
\hat{p}_b(x_a, x_b) = \frac{1}{N} \sum_{i=1}^{N} W_{N,i}(x_a, x_b) d_b^i \tag{56}
$$

where $\{W_{N,i}(x_a, x_b)\}_{i=1}^N$ are weights on each observation that specifies how "close" the given point (x_a, x_b) is to each observed data point (x_a^i, x_b^i). We can view this local average as the result of a weighted least squares fitting, where $\{W_{N,i}(x_a,x_b)\}$ are weights

$$\hat{p}_b(x_a,x_b) = \underset{\gamma}{argmin}\, \frac{1}{N}\sum_{i=1}^{N} W_{N,i}(x_a,x_b)[d_a^i - \gamma]^2. \tag{57}$$

What weights to use? A common technique is *kernel smoothing* where the weights are formed from a *kernel function* K. In this case, K would be a bivariate probability density function (such as a bivariate normal density) $K(u_1,u_2)$ that integrates to 1

$$1 = \int_{-\infty}^{\infty}\int_{-\infty} \infty K(u_1,u_2)du_1du_2. \tag{58}$$

We also often require the kernel to have mean zero in each variable (this is certain true it the kernel is the standardized bivariate normal density)

$$\begin{aligned} 0 &= \int_{-\infty}^{\infty} u_1 \int_{-\infty}^{\infty} K(u_1,u_2)du_1du_2 \\ 0 &= \int_{-\infty}^{\infty} u_2 \int_{-\infty}^{\infty} K(u_1,u_2)du_1du_2. \end{aligned} \tag{59}$$

Actually kernels need not be restricted to have mean zero, and they need not always even be probability density functions (i.e. be non-negative): they can sometimes have negative values, as long as they still integrate to 1: these negative valued-kernels are know as *higher order kernels* that are chosen to help reduce bias that is often present when the kernel is restricted to be nonnegative.

Using the kernel function, let h_N be a positive *bandwidth* or smoothing parameter. It is commonly required that $h_N \downarrow 0$, as $N \to \infty$, but at a sufficiently slow rate. Define

$$W_{N,i}(x_a,x_b) = \frac{K_{h_N}(x_a - x_a^i, x_b - x_b^i)}{\sum_{j=1}^{N} K_{h_N}(x_a - x_a^j, x_b - x_b^j)}, \tag{60}$$

where K_{h_N} is the scale-transformed kernel given by

$$K_{h_N}(x_a,x_b) = \frac{1}{h_N}K(\frac{x_a}{h_N},\frac{x_b}{h_N}) \tag{61}$$

Thus, h_N should be seen as a variance or *scaling factor* that tends to zero, so that K_{h_N} can be viewed as a bivariate probability density with a small variance. It follows that the weights $W_{N,i}(x_a,x_b)$ are simply the value of a probability density, centered at zero and with small variance, and evaluated at $(x_a - x_a^i, x_b - x_b^i)$. Thus, to the extent that the observation (x_a^i, x_b^i) is close to the given point (x_a,x_b) at which we want to estimate $\hat{p}_b(x_a,x_b)$ it will receive a high weight, but if it is far from (x_a,x_b) it will receive a low weight.

Local linear regression and local linear maximum likelihood are closely related to kernel smoothing methods. Just as $\hat{p}_b(x_a,x_b)$ could be viewed as the result of a weigthed least squares fit, we can estimate a "local linear regression model" given by the following 3 coefficients $(\hat{\gamma}_0(x_a,x_b),\hat{\gamma}_1(x_a,x_b),\hat{\gamma}_2(x_a,x_b))$ as

$$(\hat{\gamma}_0,\hat{\gamma}_1,\hat{\gamma}_2) = \underset{(\gamma_0,\gamma_1,\gamma_2)}{argmin} \frac{1}{N}\sum_{i=1}^{N} W_{N,i}(x_a,x_b)[d_b^i - \gamma_0 - \gamma_1 x_a^i - \gamma_2 x_b^i]^2 \tag{62}$$

and we can then form an estimate $\hat{p}_b(x_a,x_b)$ as

$$\hat{p}_b(x_a,x_b) = \hat{\gamma}_0(x_a,x_b) + \hat{\gamma}_1(x_a,x_b)x_a + \hat{\gamma}_2(x_a,x_b)x_b. \tag{63}$$

Finally, we can also estimate $\hat{p}_b(x_a,x_b)$ via a *local likelihood* procedure. We estimate a weighted likelihood function, using a binomial logit model with coefficients $(\gamma_0,\gamma_1,\gamma_2)$

$$\begin{aligned}(\hat{\gamma}_0,\hat{\gamma}_1,\hat{\gamma}_2) = \underset{(\gamma_0,\gamma_1,\gamma_2)}{argmin} \frac{1}{N}\sum_{i=1}^{N} W_{N,i}(x_a,x_b)\Big[\quad & d_b^i \log(1/(1+\exp\{\gamma_0+\gamma_1 x_a^i+\gamma_2 x_b^i\}))+ \\ & (1-d_b^i)\log(1-1/(1+\exp\{\gamma_0+\gamma_1 x_a^i+\gamma_2 x_b^i\}))\Big].\end{aligned} \tag{64}$$

Clearly the values of $(\hat{\gamma}_0,\hat{\gamma}_1,\hat{\gamma}_2)$ depend implicitly on (x_a,x_b), the point at which the local likelihood function is being evaluated, since the weights $W_{N,i}(x_a,x_b)$ depend on (x_a,x_b). Then, using these estimated coefficients we have

$$\hat{p}_b(x_a,x_b) = \frac{1}{1+\exp\{\hat{\gamma}_0(x_a,x_b)+\hat{\gamma}_1(x_a,x_b)x_a+\hat{\gamma}_2(x_a,x_b)x_b\}}. \tag{65}$$

There are many more types of non-parametric estimators. However this little survey should convince you that non-parametric estimation is not hard to do, and there are many choices available in how to do it. You should realize, however, that the two-step estimator is not entirely trivial to carry out. In any of the local averaging methods, we need to form an estimate of $\hat{p}_b(x_a^i,x_b^i)$ at each data point (x_a^i,x_b^i). However this only needs to be done once, at the beginning of the program, and stored for subsequent use, rather than being recomputed each time (α,β) is updated in the second stage of the semiparametric two step estimator.

6 Monte Carlo Results

To study the properties of the FIML-NFXP and two step estimators, instead of estimating the model just once using a single set of simulated data, it is quite useful to conduct a *monte carlo study.* In this case, we repeatedly resimulate the data and attempt to estimate the parameters in a do-loop, very many times. Then using the result set of estimated parameter values (one for each monte carlo "run") we can tabulate the empirical or "finite sample" distribution of the various estimators and see how close these distributions are to the approximate asymptotic normal distribution implied by standard asymptotic theory for maximum likelihood and semiparametric two step estimators.

In this initial monte carlo run, I have attempted to put the semiparametric two step estimator in the best possible light, by doing a monte carlo study of an "infeasible" version of this estimator, i.e. one where we use the true equilibrium probabilities $p_a(x_a,x_b,\theta^*)$ and $p_b(x_a,x_b,\theta^*)$ in place of non-parametric estimates of these, $\hat{p}_a(x_a,x_b)$ and $\hat{p}_b(x_a,x_b)$. Clearly, there will be extra "noise" if we use non-parametric estimates of these equilibrium probabilities, and as a result we would expect that the standard errors of the parameter estimates for the "feasible" semiparametric two step estimator to be greater than those obtained for the "infeasible" estimator (which was only possible to do in this case since we have used "artificial data" where we know the true θ^* and can therefore compute the true $p_a(x_a,x_b,\theta^*)$ and $p_b(x_a,x_b,\theta^*)$ for any (x_a,x_b).

As to be expected, the monte carlo study reveals that the FIML-NFXP estimator is substantially more efficient than the semiparametric two step estimator. Thus, there is a price to be paid for avoiding the numerical burden of doing FIML-NFXP. However from figure 8 below, we see that the infeasible "full information" version of the semiparametric two step estimator comes close to FIML in terms of efficiency. In all cases we see that the finite sample distribution of the estimators (computed using a kernel density estimator for 500 realized monte carlo values of the parameters) is close to the asymptotic normal approximation to the sampling distributions of the estimator.

Thus, we have verified that both the FIML-NFXP and semi parametric two step estimators "work" in the sense of yielding parameter estimates that are centered about the true values. Furthermore we also see that asymptotic theory "works" in the sense that it provides a relatively good aapproximation to the finite sample distributions of the estimators, estimated using the results of our monte carlo study with 500 replications.

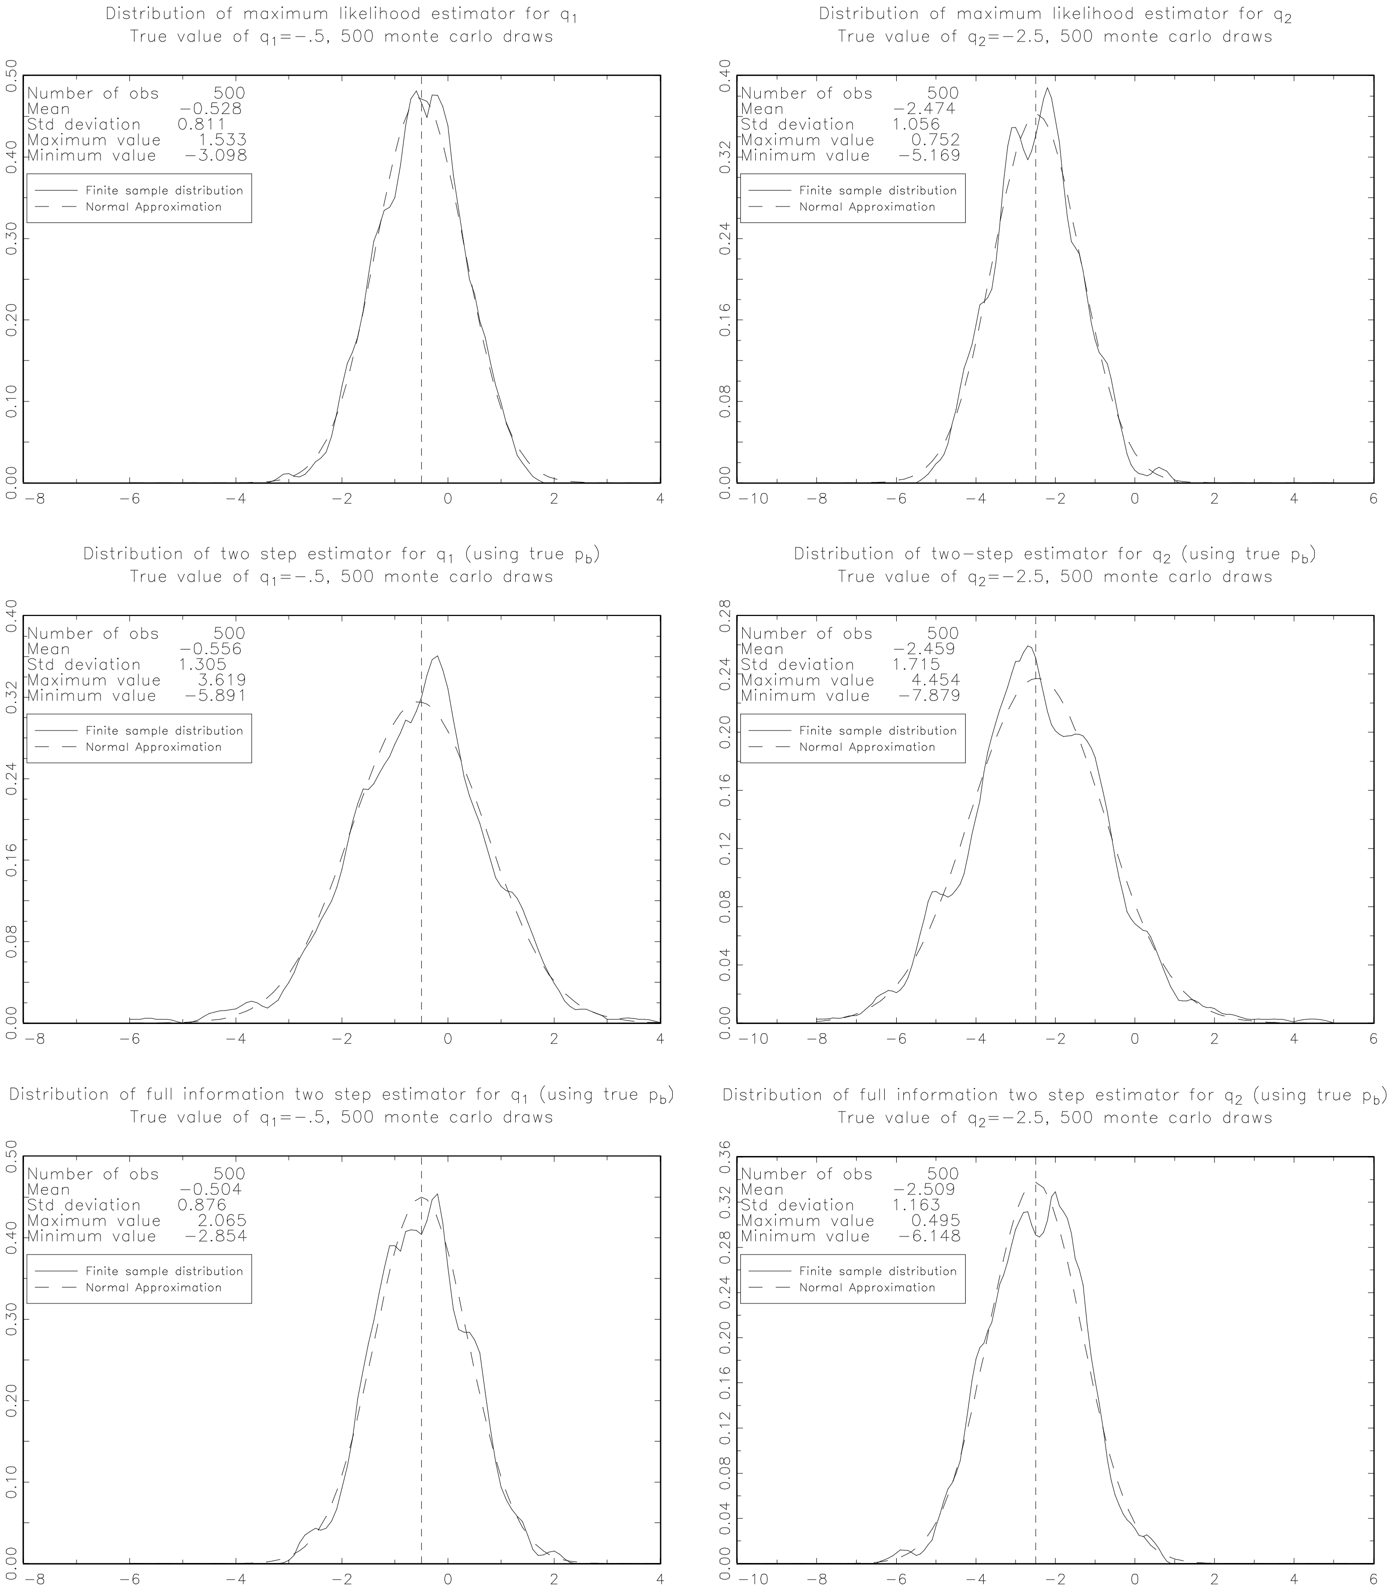

Figure 7: Expected Payoffs to Confessing as a Function of the Probability the Other Prisoner Confesses