Identification and Estimation of a Discrete Game by Observing its Correlated Equilibria

Wang, Yafeng and Graham, Brett

Institute of Social Science Survey, Peking University, 100871 Beijing, China, The Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, 361005 Xiamen, Fujian, China

8 October 2010

Identification and Estimation of a Discrete Game by Observing its Correlated Equilibria

Yafeng Wang* and Brett Graham†

Abstract

This paper studies the problem of identifying and estimating the normal-form payoff parameters of a simultaneous, discrete game of complete information where the equilibrium concept employed is correlated equilibrium rather than Nash equilibrium. We show that once we extend the equilibrium concept from Nash equilibrium to correlated equilibrium, the identification and estimation of game-theoretic econometric models becomes simpler, since this extension avoids the usual requirement of computing all the equilibria of a given game. To deal with the presence of multiple equilibria, unlike most other work on empirical games, we make use of the moment inequality restrictions induced by the underlying game-theoretic econometric models without the need to make any equilibrium selection assumptions. The resulting identified features of the model are sets of parameters such that the choice probabilities predicted by the econometric model are consistent with the empirical choice probabilities estimated from the data. The importance sampling technique is used to reduce computational burden and overcome the non-smoothness problems. We also show that the model selection tests for moment inequality models can be used to test equilibrium concepts such as correlated equilibrium versus Nash equilibrium.

Keywords: Game-Theoretic Econometric Models; Correlated Equilibrium; Partial Identification; Moment Inequality Restrictions; Importance Sampling.

JEL Classification Numbers: C35, C51, C72.

*Corresponding author. Institute of Social Science Survey (ISSS) , Peking University, 100871 Beijing, China. E-mail: econwang@pku.edu.cn.

†The Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, 361005 Xiamen, Fujian, China. E-mail: bgraham@xmu.edu.cn.

1 Introduction

Game theory is one of the cornerstones of modern economic theory, and much progress has been made in clarifying the nature of strategic interaction in economic models. It is the benchmark theoretical model for analyzing strategic interactions among a few players. Given the importance of gaming in economic theory, the empirical analysis of games has been the focus of much recent literature in econometrics and industrial organization. Since the seminal work of Bresnahan and Reiss (1990, 1991), it is common to assume, as in a standard discrete choice model, that each player's utility or payoff is a linear function of covariates and a random preference shock. However, unlike a discrete choice model, utility also depends on the actions of other agents.

Although there are numerous studies on the empirical estimation of a wide range game-theoretic econometric models, the most widely studied is the class of incomplete information games (both static and dynamic[1]. Complete information games have received less attention due to their computational complexity, since estimation involves multidimensional integrals. Moreover, complete information will generally induce the presence of multiple equilibria (Morris and Shihn (2003)). Dealing with multiple equilibria is a difficult task because a particular realization of observables and a particular set of payoffs may be consistent with different model outcomes. To address the problem presented by the requirement to compute multidimensional integrals, Bajari, Hong, and Ryan (2010) and Ciliberto and Tamer (2009) provide simulation-based estimators for static complete information games. Bajari, Hong, and Ryan (2010) outline three main alternatives to address the presence of multiple equilibria. The first approach is to introduce an equilibrium selection mechanism that determines which equilibrium will be played among several equilibria. Bajari, Hong, and Ryan (2010) and Jia (2008) are examples of such an approach.

[1] Studies of incomplete information static games include Sweeting (2005), Seim (2006), Aradillas-Lopez (2007, 2010) and Bajari, Hong, Krainer, and Nekipelov (2010), while the studies of dynamic game includes Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008) among others.

While an equilibrium selection mechanism allows for identification the underlying game, in general, we have limited knowledge about the equilibrium selection mechanism, and any misspecification about it will lead to inconsistent estimation. The second approach which was first used by Bresnahan and Reiss (1990) is to map the sets of equilibrium action profiles associated with a particular set of payoff profile to some other variable that is constant over each set. In Bresnahan and Reiss (1990), for example, the variable used is the number of entrants in the market. Of course, this is a useful method only as long as such a variable can be found. The last approach, proposed by Tamer (2003), is to partially identify parameters and thus, eliminate the need to make assumptions about any underlying equilibrium selection mechanism. Berry and Tamer (2007) and Ciliberto and Tamer (2009) are examples of this approach. Although each of these approaches can make inferences in the presence of a multiplicity of equilibria, a common practical issue is that all of them require computation of all the Nash equilibria of the underlying game. This heavy computational burden will make estimation extremely difficult, if not impossible, when dealing with large games[2].

Here, we depart from the commonly used equilibrium concept – Nash equilibrium, and assume that the outcome of the game is generated by a broader rationality rule proposed by Aumann (1974, 1987)– correlated equilibrium. A most interesting feature of this alternative equilibrium concept is that the identification and estimation of empirical games become simpler, even though it enlarges the corresponding equilibrium set[3]. Yang (2007) also uses the concept of correlated equilibrium to estimate simultaneous-move, discrete game of complete information[4]. In contrast to his paper, our error structure is more general, as will be discussed in section 2. As a result, the moment conditions used

[2] Mckelvey and Mclennan (1996) analyze the different computational methods for computing the set of Nash equilibria for general games and point out the difficulty associated with this issue.

[3] Chwe (2007) also studies the identification of games based on correlated equilibrium in a deterministic environment.

[4] We were unaware of this research until our own work was completed, and are grateful to Zhou Yang for bringing it to our attention.

for estimation differ between the two papers. The advantages of correlated equilibrium in the context of identification comes from its convexity, that is, any convex combination of correlated equilibria is also a correlated equilibrium, a property not held by the set of Nash equilibria. This property reduces the computation burden associated with estimation. We also adopt the partial identification approach to deal with the presence of multiple equilibria following Berry and Tamer (2007) and Ciliberto and Tamer (2009), where the identified set is characterized by moment inequality restrictions. However, our approach does not require the computation of all the equilibria (either correlated or Nash), but only needs to compute some "extreme" equilibria–equilibria that realize a particular outcome least or most, which can be obtained from simple linear programming. This does not mean that computing the required set of correlated equilibria is simple[5], the key feature is that it does not need the whole set of equilibria. The importance sampling technique is used to approximate the multi-dimensional integrals associated with these extreme equilibria. Given the existing research on empirical games based on Nash equilibrium, and the results established in this paper based on correlated equilibrium, we also provide a framework for testing equilibrium concepts based on the moment inequality model selection test developed by Shi (2010). The nested relationship between Nash equilibrium and correlated equilibrium makes this test similar to the famous Hausman test (Hausman, 1978).

The paper is organized as follows. In section 2 we outline the general discrete, static, complete information game to be estimated and formulate the implied equilibrium conditions associated with the concept of correlated equilibrium. Several important properties of correlated equilibrium are also presented. In section 3 we discuss the problem of partial identification of the model. Section 4 describes the procedure for estimating the identified set which is formulated in section 3, and also describes how importance sampling is used

[5] Papadimitriou and Roughgarden (2008) develop a polynomial-time algorithm for finding correlated equilibria and also discuss the difficulty in computing the complete set of correlated equilibria.

to approximate the evaluation of multiple integrals and the computation of "extreme" correlated equilibria. Section 5 introducing a test procedure for testing the behavioral hypothesis of correlated equilibrium versus Nash equilibrium. A simple Monte Carlo experiment is conducted in section 6. Section 7 concludes the paper.

2 The Model

We use the strategic environment of Bajari, Hong, and Ryan (2010) to develop our estimation method. There are T independent repetitions of a simultaneous-move (normal form) discrete game of complete information. In each game there are $i = 1, ..., N$ players. In each repetition of this game, each player i chooses an action a_i from the finite set of actions A_i simultaneously. Define $A^N = \times_i A_i$ and let $a = (a_1, ..., a_N)$ denote a generic element of A. Player i's von Neumann-Morgenstern (vNM) utility is a mapping $u_i : A^N \to R$, where R is the real line. We will follow the convention of Bajari, Hong, and Ryan (2010) and sometimes drop the subscript t for simplicity when no ambiguity would arise.

The vNM utility of player i is assumed to be:

$$u_i(a, x_i, \epsilon_i; \theta_1) = \Pi_i(x_i, a; \theta_1) + \epsilon_i(a) \tag{1}$$

where $a \in A^N$. In Equation (1), player i's vNM utility from outcome a is the sum of two terms. The first term $\Pi_i(x_i, a; \theta_1)$ is a function which depends on the vector a of actions taken by all of the players, the covariates x, which are observed by the econometrician, and parameters θ_1. The second term $\epsilon_i(a)$, is a random preference shock which reflects all the information about utility that is common knowledge to the players but not observed by the econometrician. Unlike most other study on empirical games, here the preference shocks depend on the entire vector of actions a, not just the actions taken by player i. As argued by Bajari, Hong, and Ryan (2010), this is a more general error structure than normally assumed in the literature, for example Tamer (2003). As a simple example, consider a simple two-firm entry game. This structure allows for the information unobserved by the econometrician about a firm's payoffs to depend not only on his choice of whether or not to enter a market but also on the choice of the other firm. Let ϵ_i denote the vector of the individual shocks $\epsilon_i(a)$ and ϵ denote the vector of all preference shocks, $\epsilon_i(a)$ are assumed to be independent with a density $g_i(\epsilon_i(a)|\theta_2)$ and joint distribution $g(\epsilon|\theta_2) = \prod_i \prod_{a \in A} g_i(\epsilon_i(a)|\theta_2)$, where θ_2 denotes the parameters of the distribution.

For each repetition $t \in \{1, \ldots, T\}$ of the game with the above structure, the researcher observes covariates x_t and the outcome a_t (a vector of the actions chosen by the players in period t). Unlike most other studies of empirical games (e.g., Bajari, Hong, and Ryan (2010) and Ciliberto and Tamer (2009)), we assume that observed outcomes are consistent with the concept of correlated equilibrium rather than Nash equilibrium. As a generalization of Nash equilibrium, Aumann (1974, 1987) shows that correlated equilibrium is the appropriate solution concept if players do not know the beliefs of other players but that in every state of the world each player's rationality is common knowledge. The most notable feature of correlated equilibrium is that it does not require explicit randomization on the part of the players. Rather the equilibrium can be interpreted as a set of (possibly) correlated signals that players receive that determine a unique optimal choice.

Formally, assume the game structure defined above, and let (Ω, π) be a probability space, $\mathcal{P}_i$ be a partition of Ω, $i = 1, ..., N$, and let

$$Q_i = \{q_i : \Omega \to A_i | q_i \text{ is } \mathcal{P}_i \text{ measurable.}\} \tag{2}$$

If we refer to the partition as $\mathcal{P}_i = \{P_i(\omega)\}_{\omega \in \Omega}$, where $P_i(\omega)$ is the element of the partition containing ω, then correlated equilibrium can be defined as[6]:

[6]This following definitions and discussion of correlated equilibrium follow Bergin (2005). Similar

Definition 2.1 (Correlated Equilibrium) *The collection* $(\Omega, \pi, \{\mathcal{P}_i\}_{i=1}^N, \{q_i\}_{i=1}^N)$ *is a correlated equilibrium if* $\forall i$,

$$\sum_{\omega} u_i(q_{-i}(\omega), q_i(\omega))\pi(\omega) \geq \sum_{\omega} u_i(q_{-i}(\omega), \tau_i(\omega))\pi(\omega),\ \forall \tau_i \in Q_i \tag{3}$$

where for each i, q_i *is constant on each member of* $\mathcal{P}_i$.

Intuitively, given the information received through the partitioning of Ω and the realized state ω, players choose actions to maximize their expected utility. Thus, (3) define sufficient conditions for this utility maximization.

The formulation of a correlated equilibrium in Definition 2.1 allows for complete flexibility in defining the elements of the state space of a correlated equilibrium and, thus, leads to a broad range of interpretations of a correlated equilibrium (e.g. sunspot equilibria). However, from a computational point of view, it is more useful to restrict attention to canonical correlated equilibria, correlated equilibria where the state space is identified with the space of pure strategies, that is $\Omega = A^N$ and, for each player, the partition $\mathcal{P}_i$ of Ω is generated by A_i.

The following proposition states the strategic equivalence between correlated equilibrium and canonical correlated equilibrium[7].

Proposition 2.1 *Let* $(\Omega, \pi, \{P\}_{i=1}^N, \{q_i\}_{i=1}^N)$ *be a correlated equilibrium. Then there is a canonical correlated equilibrium yielding the same distribution on actions and the same expected payoff to each player.*

Based on this strategic equivalence, we will restrict attention to the concept of canonical correlated equilibrium and, without ambiguity, refer to it as correlated equilibrium for the remainder of the paper.

treatments can be found in Osborne and Rubinstein (1994), Forges (2009) and Sorin (1997).

[7] Proofs of all of the following propositions regarding correlated equilibrium can be found in Bergin (2005) or Osborne and Rubinstein (1994).

The following properties of correlated equilibrium will also be useful in identifying and estimating the underlying parameters of our model:

Proposition 2.2 *The set of Nash equilibrium payoffs is a subset of the set of correlated equilibrium payoffs.*

Proposition 2.3 *The set of correlated equilibrium payoffs is a convex set.*

Proposition (2.2) shows that the set of probability distributions over outcomes induced by the set of Nash equilibria is equivalent to the set of probability distributions over correlated equilibrium that are the product of independent probability distributions over each player's actions. The convexity of the set of correlated equilibrium payoffs will facilitate the computation of correlated equilibrium, the resulting identified region will allow us to restrict attention to "extreme" correlated equilibria–for each outcome, the equilibria that attach the least and the most probability to that outcome.

The nested relationship between correlated equilibrium and Nash equilibrium makes our test of equilibrium concepts similar to the famous Hausman test (Hausman, 1978). If the outcomes of the underlying game are consistent with Nash equilibrium, but the researcher estimates the parameters of the game based on the concept of correlated equilibrium, then the estimate are consistent but inefficient, while if the outcomes are consistent with correlated equilibrium, but the researcher estimates the game based on Nash equilibrium, then the estimates are inconsistent.

Given the structure of the discrete normal form game described above, assume that the observed outcomes of such a game are consistent with the concept of correlated equilibrium, i.e., there exists a distribution π over the set of outcomes A^N such that:

$$\sum_{a_{-i}} u_i(a_{-i}, a_i, x_i, \epsilon_i; \theta_1)\pi(a_{-i}|a_i) \geq \sum_{a_{-i}} u_i(a_{-i}, a_i', x_i, \epsilon_i; \theta_1)\pi(a_{-i}|a_i),\ \forall a_i' \in A_i, i = 1, ..., N \tag{4}$$

Our task is to estimate and draw an inference about the parameters of the payoff functions, θ_1, and the parameters of the distribution of random preference shocks, θ_2, from the observed outcomes a_t^o, and some exogenous covariates which affect the payoffs, x_t. Note that the actual payoff levels are unobserved, i.e., they are latent variables. To aid the discussion, let $S_\pi(u(x, \epsilon; \theta_1))$ denote the collection of distributions over outcomes which satisfies the equilibrium condition (4), i.e., any distribution over outcomes associated with a correlated equilibrium of the underlying game, given the payoffs associated with the outcomes of the game. These payoffs, in turn, are determined by the set of covariates, x, the set of random shocks, ϵ and the parameters of the utility function θ_1. Let $\pi(u(x, \epsilon; \theta_1))$ denote a generic elements of the set $S_\pi(u(x, \epsilon; \theta_1))$. For the purposes of exposition, we will sometimes simply refer to it as π, or $\pi \in S_\pi(u(x, \epsilon; \theta_1))$.

3 Identification

The general strategy to identify the structural parameters of a game-theoretic econometric model is to match the choice probabilities predicted by the model with the empirical choice probabilities observed from the data (see Bajari, Hong, Krainer, and Nekipelov (2010), Bajari, Hong, and Ryan (2010) and Ciliberto and Tamer (2009)). The empirical choice probabilities can usually be obtained from the data nonparametrically. However, the multiplicity of equlibria associated with a solution concept generally, and correlated equilibrium in particular, in addition to the absence of an observed equilibrium selection mechanism makes obtaining the choice probabilities predicted by the game structure problematic. To solve this problem, we follow Ciliberto and Tamer (2009), who use the restrictions on the distribution of the selection mechanism over the set of Nash equilibria implied by the laws of probability to partially identify the model parameters. We use these same restrictions over the set of correlated equilibria. This is in contrast to the approach of Bajari, Hong, and Ryan (2010) who introduce an explicit equilibrium selection mechanism over the set of Nash equilibria to achieve point identification of structural parameters.

First enumerate the elements of A from $a = \{1, ..., \#A\}$. A is the set of pure strategy profiles and $a \in A$. Given a correlated equilibrium $\pi \in S_\pi(u(x, \epsilon; \theta_1))$ is a distribution over A, we have

$$\pi = (\pi(1), ..., \pi(a), ..., \pi(\#A))' \tag{5}$$

and

$$\sum_{a=1}^{\#A} \pi(a) = 1; 0 \leq \pi(a) \leq 1; \forall a \in A \tag{6}$$

Let $\mathcal{Y}$ be the set of *potentially observable outcomes*. Since we assume that the observable outcome of the game is the equilibrium actions chosen by all the players, then $\mathcal{Y} = A$. Let $\Pr(y = a|x; \theta)$ denote the the probability that action profile a be the equilibrium action profile predicted by the model, where $\theta = (\theta_1, \theta_2)$, and let $\Pr(y = a|x)$ be the empirical choice probability identified from the data which is independent of the values of the structural parameters.

Identification requires the following assumptions:

Assumption 1 *The parameter space Θ is compact.* Assumption

Assumption 2 *The payoffs of one action for each player are fixed at a known constant.*

Assumption 3 *The joint distribution of $\epsilon = (\epsilon_i(a))$, $G(\epsilon|\theta_2)$ is independent, independent of x, and known to all agents and the econometrician, and let $g(\epsilon|\theta_2)$* be the corresponding density.

Assumption 4 (Identification of $Pr(y|x)$) *The econometrician observes data that identifies* $\Pr(y = a|x)$, $\forall a \in A$.

Assumption 1, compactness of the parameter space is critical for the construction the large sample property of our estimator (Chernozhukov, Hong, and Tamer (2007)). Assumptions 2 and 3 are common in the literature (see Berry (1992) and Ciliberto and Tamer

(2009)). Similar to the model environment, the phrasing we use here is taken from Bajari, Hong, and Ryan (2010). One can clearly see from the equilibrium condition (4) that any affine transformation of all deterministic payoffs does not change the set of equilibria. Thus, the need for the location normalization of Assumption 2. The scale normalization is included as part of Assumption 3. Assumption 4 requires that the empirical choice probabilities can be identified from the data. Clearly, this is necessary since identification relies on matching this probability with the choice probability predicted by the model.

As discussed before, the set of correlated equilibria, $S_\pi(u(\theta, x, \epsilon))$ will usually be an uncountably infinite set. If $S_\pi(u(\theta, x, \epsilon))$ is non-singleton, in order to derive the choice probability predicted by the model, $\Pr(y = a|x; \theta)$, we need to introduce an equilibrium selection mechanism:

$$\psi(\cdot|x, \epsilon) : S_\pi(u(x, \epsilon; \theta_1)) \to [0, 1]^{d[S_\pi(u(x,\epsilon;\theta_1))]} \tag{7}$$

such that

$$\psi(\cdot|x, \epsilon) \geq 0 \text{ and} \tag{8}$$

$$\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x, \epsilon) = 1 \tag{9}$$

where $d[S_\pi(u(x, \epsilon; \theta_1))]$ is the dimension of $S_\pi(u(x, \epsilon; \theta_1))$. This equilibrium selection mechanism specifies the probability, $\psi(\pi|x, \epsilon)$, that any correlated equilibrium $\pi \in S_\pi(u(x, \epsilon; \theta_1))$ be the chosen equilibrium. Since the $d[S_\pi(u(x, \epsilon; \theta_1))]$ is, in general, infinite, we should use a continuous distribution to express this equilibrium selection mechanism, but for purposes of exposition, we use the discrete distribution.

Given the equilibrium selection mechanism (7), the choice probabilities implied by the model can be written as:

$$\Pr(y = a|x;\theta) = \int \left(\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a) \right) dG(\epsilon|\theta_2) \tag{10}$$

where $\pi(a)$ is the probability that action profile a is selected (or realized) if the correlated equilibrium is π, and $\psi(\pi|x,\epsilon)$ is the probability that π is the selected equilibrium. Thus, $\psi(\pi|x,\epsilon)\pi(a)$ is the joint probability that action profile a and correlated equilibrium π are selected. Clearly, action profile a may be associated with other correlated equilibria, thus the summation of these probabilities, $\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a)$, is the probability that action profile a is the selected equilibrium action profile. Based on the choice probabilities implied by equation (10), we can define the sharp identified set for the parameter $\theta = (\theta_1, \theta_2)$.

Definition 3.1 (Sharp Identified Set) *The sharp identified set for the parameter vector $\theta \in \Theta$ is given by:*

$$\Theta_I = \left\{ \begin{array}{ll} & \exists \psi,\ \forall a \in \mathcal{Y} \\ \theta \in \Theta : & \textit{such that: } E[\Pr(y = a|x)] = E[\Pr(y = a|x;\theta)] \\ & \quad = E\left[\int \left(\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a) \right) dG(\epsilon|\theta_2) \right] \end{array} \right\} \tag{11}$$

Inference on the set Θ_I based on (11) is not practically feasible since one needs to deal with the infinite dimensional nuisance parameters $\psi(\cdot|x,\epsilon)$ that result from the multiplicity of equilibria. Note further that the equilibrium selection mechanism also depends on the unobserved random preference shock ϵ. It is possible to follow the approach of Bajari, Hong, and Ryan (2010) here and specify a parametric equilibrium selection mechanism that is characterized by a finite number of parameters. In general, however, we do not have sufficient information to specify a particular equilibrium selection mechanism, and any misspecification of this mechanism will induce inconsistent estimation. For particular model settings one may instead use a more refined equilibrium concept, such as perfect correlated equilibrium (Dhillon and Mertens, 1996) or maximum entropy correlated equilibrium (Ortiz, Schapire, and Kakade, 2007). However, for general models such refinements do not guarantee a unique equilibrium. In the spirit of Ciliberto and Tamer (2009), we leave the equilibrium selection mechanism unspecified but, instead, exploit the fact that the equilibrium selection mechanism $\psi(\pi|x,\epsilon)$ is a probability and hence bounded between zero and one to derive a outer identified set for the structural parameters.

Since the equilibrium selection mechanism $\psi(\pi|x,\epsilon)$ is a probability distribution, then

$$0 \leq \psi(\pi|x,\epsilon) \leq 1,\ \forall \pi \in S_\pi(u(x,\epsilon;\theta_1)) \tag{12}$$

Based on this natural property of probability, we can derive an outer identified set for the parameter θ. Formally, let $H_1^a(\theta,X)$ denote the lower bound of the choice probability of action profile a implied by the model, $\Pr(y=a|x;\theta)$, and $H_2^a(\theta,X)$ the upper bound, then:

$$H_1^a(\theta,X) = \min \int \left[\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a) \right] dG(\epsilon|\theta_2) \tag{13}$$

$$H_2^a(\theta,X) = \max \int \left[\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a) \right] dG(\epsilon|\theta_2) \tag{14}$$

Given the exogenous covariates, X, and payoff parameter θ_1, define $R_1^\pi(\theta_1,X)$ as the set of random preference shocks ϵ such that the game admits π as the unique equilibrium

$R_2^\pi(\theta_1, X)$ as the complement of $R_1^\pi(\theta_1, X)$ Thus we have:

$$\begin{aligned}
&H_1^a(\theta, X) \\
&= \min \int \left[\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a) \right] dG(\epsilon|\theta_2) \\
&= \underbrace{\int_{R_1^\pi(\theta,X)} \pi(a) dG(\epsilon|\theta_2)}_{(1)} + \underbrace{\int_{R_2^\pi(\theta,X)} \min\{\pi(a) : \pi \in S_\pi(u(x,\epsilon;\theta_1))\} dG(\epsilon|\theta_2)}_{(2)}.
\end{aligned} \tag{15}$$

The last equality of equation (15) separates the calculation of the lower bound of the choice probability of a into an integral over the support of the preference shocks that admits a unique equilibrium and an integral over the support of the preference shocks that admits multiple equilibria. The integrand of the first integral $\pi(a)$ is the choice probability of the unique correlated equilibrium π implied by the model over this support. The integrand of the second integral is the probability of a associated with the correlated equilibrium that realizes a with the lowest probability. The true equilibrium selection mechanism must select outcome profile a with at least this probability. Thus, we identify the lower bound of $\Pr(y = a|x; \theta)$. Similarly, the upper bound $H_2^a(\theta, X)$ can be derived as:

$$\begin{aligned}
&H_2^a(\theta, X) \\
&= \max \int \left[\sum_{\pi \in S_\pi(u(x,\epsilon;\theta_1))} \psi(\pi|x,\epsilon)\pi(a) \right] dG(\epsilon|\theta_2) \\
&= \underbrace{\int_{R_1^\pi(\theta,X)} \pi(a) dG(\epsilon|\theta_2)}_{(1)} + \underbrace{\int_{R_2^\pi(\theta,X)} \max\{\pi(a) : \pi \in S_\pi(u(x,\epsilon;\theta_1))\} dG(\epsilon|\theta_2)}_{(2)}
\end{aligned} \tag{16}$$

In an equivalent fashion to equation (15), the integrand of the second integral on the last line of (16) is the probability of a associated with the correlated equilibrium that realizes a with the highest probability. The true equilibrium selection mechanism can select outcome profile a with no more than this probability. Thus, we identify the upper bound of $\Pr(y = a|x; \theta)$.

Based on the lower bound and upper bound of the choice probabilities implied by the model, we have:

$$H_1^a(\theta, X) \leq \Pr(y = a|x; \theta) \leq H_2^a(\theta, X) \tag{17}$$

And when $\theta \in \Theta_I$

$$E[\Pr(y = a|x)] = E[\Pr(y = a|x; \theta)] \tag{18}$$

Thus we can define the outer identified set for the model parameter θ as:

Definition 3.2 (Outer Identified Region) *The outer identified set for model parameter $\theta = (\theta_1, \theta_2) \in \Theta$ is*

$$\Theta_O = \left\{ \begin{array}{cc} & \forall a \in \mathcal{Y} \\ \theta \in \Theta : & \textit{such that:} \\ & E[H_1^a(\theta, X)] \leq E[\Pr(y = a|x)] \leq E[H_2^a(\theta, X)] \end{array} \right\} \tag{19}$$

By introducing the following definitions:

$$\boldsymbol{H}_1(\theta, X) = (H_1^1(\theta, X), ..., H_1^a(\theta, X), ..., H_1^{\#A}(\theta, X))'$$

$$\boldsymbol{H}_2(\theta, X) = (H_2^1(\theta, X), ..., H_2^a(\theta, X), ..., H_2^{\#A}(\theta, X))'$$

and

$$\Pr\left(\boldsymbol{y}|\boldsymbol{x}\right) = (\Pr(y = 1|x), ..., \Pr(y = a|x), ..., \Pr(y = \#A|x))',$$

conditions that define the outer identified set can be stated as:

$$E[\boldsymbol{H}_1(\theta, X)] \leq E[\Pr(\boldsymbol{y}|\boldsymbol{x})] \leq E[\boldsymbol{H}_2(\theta, X)] \tag{20}$$

Note that the outer identified set Θ_O is broader than the sharp identified set Θ_I. Given that we do not have enough information about the equilibrium selection mechanism, the outer identified set Θ_O is the most we can learn about parameter θ from the underlying game and observed data. In general, the set is not a singleton, as it is characterized by the moment inequality restrictions. Such a model is called a partially identified econometric model, in contrast to the usual point identified case.

4 Estimation

The estimation problem is based on the moment inequality (20)

$$E[\mathbf{H}_1(\theta, \mathbf{X})] \leq E[\Pr(\mathbf{y}|\mathbf{x})] \leq E[\mathbf{H}_2(\theta, \mathbf{X})]. \tag{21}$$

We follow Chernozhukov, Hong, and Tamer (2007) which provide a general framework for moment inequality models to build a consistent estimator for the outer identified set Θ_O. Since the upper and lower bounds in moment conditions (21) contain multi-dimensional integrals, we first provide a simulation procedure to approximate these integrals. Due to the discreteness problem associated with simple Monte Carlo integration, we make use instead of importance sampling Monte Carlo integration, in the spirit of Ackerberg (2009) and Bajari, Hong, and Ryan (2010)[8].

4.1 Importance Sampling Approximation

Importance sampling is most noted for its ability to reduce simulation error and computational burden, and was first used in game-theoretic models by Bajari, Hong, and Ryan (2010). From the derivation (15) of $H_1^a(\theta, X)$ and (16) of $H_2^a(\theta, X)$ it is easily seen that $\pi(a|u)$ and the set of correlated equilibria $S_\pi(u(x, \epsilon; \theta_1))$ are both determined by the payoff

[8] McFadden (1989) noted the ability to use importance sampling to smooth simulations which is extended by Ackerberg (2009).

level u, and are influenced by θ_1 only through its effect on u. Thus, we can change the variable of integration in (15) and (16) from ϵ to u. Let $h(u|X,\theta)$ denote the density of u, conditional on x and θ. Based on the utility function $u_i(a, x_i, \epsilon_i; \theta_1) = \Pi_i(x_i, a; \theta_1) + \epsilon_i(a)$ and the density for ϵ, $g(\epsilon|\theta_2)$, $h(u|X,\theta)$ can be derived as:

$$h(u|X,\theta) = \prod_i \prod_{a \in A} g(u_i(a) - \Pi_i(x_i, a; \theta_1)|\theta_2) \tag{22}$$

Thus,

$$\begin{aligned}
&H_1^a(\theta, X) \\
&= \int\limits_{R_1^\pi(\theta,X)} \pi(a|u) dG(\epsilon|\theta_2) + \int\limits_{R_2^\pi(\theta,X)} \min\{\pi(a|u) : \pi \in S_\pi(u(x,\epsilon;\theta_1))\} dG(\epsilon|\theta_2) \qquad (23) \\
&= \int\limits_{R_1'} \pi(a|u) h(u|X,\theta) du + \int\limits_{R_2'} \min\{\pi(a|u) : \pi \in S_\pi(u)\} h(u|X,\theta) du
\end{aligned}$$

and

$$\begin{aligned}
&H_2^a(\theta, X) \\
&= \int\limits_{R_1^\pi(\theta,X)} \pi(a|u) dG(\epsilon|\theta_2) + \int\limits_{R_2^\pi(\theta,X)} \max\{\pi(a|u) : \pi \in S_\pi(u(x,\epsilon;\theta_1))\} dG(\epsilon|\theta_2) \qquad (24) \\
&= \int\limits_{R_1'} \pi(a|u) h(u|X,\theta) du + \int\limits_{R_2'} \max\{\pi(a|u) : \pi \in S_\pi(u)\} h(u|X,\theta) du
\end{aligned}$$

where R_1' is the set of u such that the game admits a unique equilibrium, and R_2' is the set of u such that the game admits multiple equilibria. By introducing an importance density $q(u)$, we can rewrite (23) and (24) as:

$$
\begin{aligned}
& H_1^a(\theta, X) \\
& = \int\limits_{R_1'} \pi(a|u) \frac{h(u|X,\theta)}{q(u)} q(u) du + \int\limits_{R_2'} \min\{\pi(a|u) : \pi \in S_\pi(u)\} \frac{h(u|X,\theta)}{q(u)} q(u) du
\end{aligned} \tag{25}
$$

and

$$
\begin{aligned}
& H_2^a(\theta, X) \\
& = \int\limits_{R_1'} \pi(a|u) \frac{h(u|X,\theta)}{q(u)} q(u) du + \int\limits_{R_2'} \max\{\pi(a|u) : \pi \in S_\pi(u)\} \frac{h(u|X,\theta)}{q(u)} q(u) du
\end{aligned} \tag{26}
$$

We can then simulate $H_1^a(\theta, X)$ and $H_2^a(\theta, X)$ by drawing random variables $(u^1, ..., u^{ns}, ..., u^{NS})$ from the importance density $q(u)$. Note that here u^{ns} is a vector; a vector of utilities for all of the players of the underlying game. Based on these simulated utility values, the importance sampling simulators for $H_1^a(\theta, X)$ and $H_2^a(\theta, X)$ are $\tilde{H}_1^a(\theta, X)$ and $\tilde{H}_2^a(\theta, X)$, respectively.

$$
\begin{aligned}
\tilde{H}_1^a(\theta, X) \quad = \quad & \tfrac{1}{NS} \sum_{ns} I(u^{ns} \in R_1') \pi(a|u^{ns}) \tfrac{h(u^{ns}|X,\theta)}{q(u^{ns})} + \\
& \tfrac{1}{NS} \sum_{ns} I(u^{ns} \in R_2') \min\{\pi(a|u^{ns}) : \pi \in S_\pi(u^{ns})\} \tfrac{h(u^{ns}|X,\theta)}{q(u^{ns})}
\end{aligned} \tag{27}
$$

$$
\begin{aligned}
\tilde{H}_2^a(\theta, X) \quad = \quad & \tfrac{1}{NS} \sum_{ns} I(u^{ns} \in R_1') \pi(a|u^{ns}) \tfrac{h(u^{ns}|X,\theta)}{q(u^{ns})} + \\
& \tfrac{1}{NS} \sum_{ns} I(u^{ns} \in R_2') \max\{\pi(a|u^{ns}) : \pi \in S_\pi(u^{ns})\} \tfrac{h(u^{ns}|X,\theta)}{q(u^{ns})}
\end{aligned} \tag{28}
$$

From the theory of importance sampling, $\tilde{H}_1^a(\theta, X)$ and $\tilde{H}_2^a(\theta, X)$ are unbiased simulators for $H_1^a(\theta, X)$ and $H_2^a(\theta, X)$, respectively. Most importantly, these simulators will generally be continuous in the parameter θ since they only depend on θ through $h(u|x, \theta)$ which is continuous in θ given that $g(\epsilon|\theta_2)$ is continuous.

The theory of importance sampling proves $\tilde{H}_1^a(\theta, X)$ and $\tilde{H}_2^a(\theta, X)$ that are smooth and unbiased simulators for any choice of the importance density $q(u)$ which has sufficiently large support. However, as noted by Bajari, Hong, and Ryan (2010), as a practical matter, it is important to make sure that the tails of the importance density are not too thin in a neighborhood of the parameter which optimizes the objective function in our estimation procedure. We suggest using some pre-estimated $\mathring{\theta}$ to construct the importance density

$$q(u) = h(u|X, \mathring{\theta}), \tag{29}$$

which can be obtained from the estimates of the game with incomplete information developed by Bajari, Hong, Krainer, and Nekipelov (2010), or through the generalized maximum entropy estimator for static games of complete information (Golan, Karp, and Perloff, 2000). Note that these two studies on empirical games are both based on the concept of Nash equilibrium.

4.2 Estimation

Given the simulators obtained from the importance sampling, $\tilde{H}_1^a(\theta, X)$ and $\tilde{H}_2^a(\theta, X)$, for $H_1^a(\theta, X)$ and $H_2^a(\theta, X)$, respectively, define

$$\tilde{\boldsymbol{H}}_1(\theta, X) = (\tilde{H}_1^1(\theta, X), ..., \tilde{H}_1^a(\theta, X), ..., \tilde{H}_1^{\#A}(\theta, X))'$$

$$\tilde{\boldsymbol{H}}_2(\theta, X) = (\tilde{H}_2^1(\theta, X), ..., \tilde{H}_2^a(\theta, X), ..., \tilde{H}_2^{\#A}(\theta, X))'$$

From (21) we get the following simulated moment inequality restrictions:

$$E[\tilde{\boldsymbol{H}}_1(\theta, X_t)] \leq E[\Pr(\boldsymbol{y}|\boldsymbol{x}_t)] \leq E[\tilde{\boldsymbol{H}}_2(\theta, X_t)] \tag{30}$$

According to Chernozhukov, Hong, and Tamer (2007), our inferential procedure uses the objective function[9]:

$$\min_{\theta\in\Theta} Q(\theta) \equiv \int \left\|(\Pr(\boldsymbol{y}|\boldsymbol{x}) - \tilde{\boldsymbol{H}}_1(\theta, X))_-\right\|^2 + \left\|(\Pr(\boldsymbol{y}|\boldsymbol{x}) - \tilde{\boldsymbol{H}}_2(\theta, X))_+\right\|^2 dF_x \tag{31}$$

to estimate the unknown parameters associated with (30). If $\Pr(\boldsymbol{y}|\boldsymbol{x}) < \tilde{\boldsymbol{H}}_1(\theta, X)$, then $\left\|(\Pr(\boldsymbol{y}|\boldsymbol{x}) - \tilde{\boldsymbol{H}}_1(\theta, X))_-\right\|^2$ is strictly positive, and if $\Pr(\boldsymbol{y}|\boldsymbol{x}) > \tilde{\boldsymbol{H}}_2\ (\theta, X)$, then $\left\|(\Pr(\boldsymbol{y}|\boldsymbol{x}) - \tilde{\boldsymbol{H}}_2(\theta, X))_+\right\|^2$ is strictly positive. It is easy to see that $Q(\theta) \geq 0$ for all $\theta \in \Theta$ and that $Q(\theta) = 0$ if and only if $\theta \in \Theta_O$.

To estimate the outer identified set Θ_O, we need to take a sample analog of $Q(\theta)$. First replace $\Pr(\boldsymbol{y}|\boldsymbol{x})$ with a $\sqrt{T}$ consistent estimator $P_T(X)$[10]. The sample analog for $Q(\theta)$ is

$$Q_T(\theta) = \frac{1}{T^2}\sum_{t=1}^{T} \left[\left\|((P_T(X_t) - \tilde{\boldsymbol{H}}_1(\theta, X_t))_-\right\|^2 + \left\|P_T(X_t) - \tilde{\boldsymbol{H}}_2(\theta, X_t))_+\right\|^2\right]. \tag{32}$$

Our estimation for Θ_O is any solution that minimizing (32), which can be obtained from:

$$\hat{\Theta}_O = \{\theta \in \Theta : TQ_T(\theta) \leq v_T\} \tag{33}$$

where $v_T \to \infty$ and $\frac{v_T}{T} \to 0$, Chernozhukov, Hong, and Tamer (2007) propose a resampling method to obtain a suitable v_T.

Proposition 4.1 *Let Assumption 3 hold. Suppose that the regularity conditions of Theorem 3.1 in Chernozhukov, Hong, and Tamer (2007) hold. Then we have that $\hat{\Theta}_O$ is a Hausdorff consistent estimator for Θ_O, that is, $d_H(\hat{\Theta}_O, \Theta_O) = 0$ with probability one.*

[9]Let $\| x \|_+ = \| (x)_+ \|$ and $\| x \|_- = \| (x)_- \|$, where $(x)_+ := \max(x, 0)$, $(x)_- := \max(-x, 0)$ and $\| \cdot \|$ is the Euclidian norm.

[10]The convergence rate of nonparametric estimates for $P_T(X)$ are slower than $\sqrt{T}$ when there are continuous variables in x, a useful method is to discretize all the variables in x and use nonparametric frequency estimation.

The proof of Theorem 4.1 is the same as that for Theorem 3.1 in Chernozhukov, Hong, and Tamer (2007). To conduct inference about the above moment inequalities model, we use the methodology of Chernozhukov, Hong, and Tamer (2007) and Ciliberto and Tamer (2009), which requires construction of a set $\mathcal{C}_T$ for a prespecified $\alpha \in (0, 1)$ such that

$$\lim_{T\to\infty} (\theta_O \in \mathcal{C}_T) \geq \alpha \text{ for any } \theta_O \in \Theta_O. \tag{34}$$

Our construction is as follows. Let

$$\mathcal{C}_T(c) = \left\{\theta \in \Theta : T\left(Q_T(\theta) - \min_z Q_T(z)\right) \leq c\right\}. \tag{35}$$

We iterate once over the following steps:

1. Compute an initial estimate for Θ_O as $\mathcal{C}_T(c_0)$, for example $\mathcal{C}_T(c_0) = \mathcal{C}_T(0)$, then subsample the statistic $T\left(Q_T(\theta) - \min_z Q_T(z)\right)$ for $\theta \in \mathcal{C}_T(0)$ and obtain the estimate of its α-quantile, $c_1(\theta_0)$.

2. Update c through $c_1 = \sup_{\theta_0 \in \mathcal{C}_T(c_0)} c_1(\theta_0)$ and return to step 1, but replace c_0 with c_1.

Thus, $\mathcal{C}_T(c_2)$ is our confidence region for $\hat{\Theta}_O$. See Chernozhukov, Hong, and Tamer (2007) and Ciliberto and Tamer (2009) for more detail on this. Such a confidence region not only has the desired coverage property, but is also consistent in the sense of Theorem 4.1.

4.3 Computation of the Equilibria

The simulated lower and upper bounds, $\tilde{H}_1^a(\theta, X)$ and $\tilde{H}_2^a(\theta, X)$, contain the following equilibrium computations:

$$I(u \in R_1')\pi(a|u) \tag{36}$$

$$I(u \in R_2')\min\{\pi(a|u) : \pi \in S_\pi(u)\} \tag{37}$$

and

$$I(u \in R_2') \max\{\pi(a|u) : \pi \in S_\pi(u)\} \tag{38}$$

where u is a vector which contains the utility levels of all the players for each action profile. We first discuss the computation of (37) and (38), where the corresponding game admits multiple equilibria. First note that if once we identify the regions R_1' and R_2', then we only need to compute the correlated equilibrium which realizes action profile a with the lowest probability and compute the correlated equilibrium which realizes action profile a with the highest probability Both of these can be obtained through simple linear programming. The first correlated equilibrium solves

$$\begin{aligned} &\min_\pi \ \pi(a) \\ &s.t. \begin{cases} \displaystyle\sum_{a_{-i}} u_i(a_{-i}, a_i)\pi(a_{-i}, a_i) \geq \sum_{a_{-i}} u_i(a_{-i}, \tilde{a}_i)\pi(a_{-i}, a_i), \ \forall i, a_i \text{ and } \tilde{a}_i \neq a_i \\ \displaystyle\sum_{a \in A} \pi(a) = 1, \ \pi(a) \geq 0, \end{cases} \end{aligned} \tag{39}$$

and the second solves

$$\begin{aligned} &\max_\pi \ \pi(a) \\ &s.t. \begin{cases} \displaystyle\sum_{a_{-i}} u_i(a_{-i}, a_i)\pi(a_{-i}, a_i) \geq \sum_{a_{-i}} u_i(a_{-i}, \tilde{a}_i)\pi(a_{-i}, a_i), \ \forall i, a_i \text{ and } \tilde{a}_i \neq a_i \\ \displaystyle\sum_{a \in A} \pi(a) = 1, \ \pi(a) \geq 0. \end{cases} \end{aligned} \tag{40}$$

If the game has a unique equilibrium, then the solution for the system of linear inequalities:

$$\begin{aligned} &\sum_{a_{-i}} u_i(a_{-i}, a_i)\pi(a_{-i}, a_i) \geq \sum_{a_{-i}} u_i(a_{-i}, \tilde{a}_i)\pi(a_{-i}, a_i), \ \forall i, a_i \text{ and } \tilde{a}_i \neq a_i \\ &\sum_{a \in A} \pi(a) = 1, \ \pi(a) \geq 0 \end{aligned} \tag{41}$$

is unique. Thus, this unique equilibrium solves both linear program (39) and (40), which means that, in practice, we do not need to identify payoffs based on whether or not they admit multiple equilibria. The only computation required is linear program (39) or (40). Clearly, the computation of equilibria in our procedure is very simple. Studies which focus on empirical estimation of complete information games based on Nash equilibrium, need to compute all the Nash equilibrium of underlying game, which will, in general, induce a heavy computational burden. See for example, Berry and Tamer (2007), Bajari, Hong, and Ryan (2010) and Ciliberto and Tamer (2009).

5 Test of Equilibrium Concepts

In this paper, we use correlated equilibrium to empirically identify and estimate a static complete information game. Of course, it is an open question which equilibrium concept most appropriately model strategic choices. A non-exhaustive of solution concepts used in the literature include pure strategy Nash equilibrium, mixed strategy equilibrium, correlated equilibrium and evolutionary equilibrium. In this section we outline a formal empirical test of the suitablility of Nash equilibrium as a solution concept as compared to correlated equilibrium.

Formally, let $\mathcal{CE}$ denote the set of parameter estimates of a static complete information game using correlated equilibrium as the solution concept and $\mathcal{NE}$ the set of parameter estimates of the same game using Nash equilibrium as the solution concept. Since both solution concepts can be characterized by moment inequality restrictions, then

$$\mathcal{CE} = \bigcup_{\theta \in \Theta} \mathcal{CE}_\theta; \; \mathcal{NE} = \bigcup_{\beta \in B} \mathcal{NE}_\beta, \tag{42}$$

where

$$\mathcal{CE}_\theta = \{CE : E_{CE} m_j(X_i, \theta) \geq 0, \; j \in J_{CE}\} \tag{43}$$

and

$$\mathcal{NE}_\beta = \{NE : E_{NE}g_j(X_i, \beta) \geq 0,\, j \in J_{NE}\}. \tag{44}$$

If $\{X_i \in \mathcal{X}\}_{i=1}^n$ is the sample generated from distribution μ, and $m_j(X_i, \theta)$ and $g_j(X_i, \beta)$ are moment functions characterized by finite dimensional parameter θ and β, respectively, then $E_{CE}m_j(X_i, \theta) \geq 0$ is equivalent to the moment inequalities (20), while $E_{NE}g_j(X_i, \beta) \geq 0$ is equivalent to the moment conditions in Ciliberto and Tamer (2009)[11].

Given the above structure, we want to test which of the two distributions $\mathcal{CE}$ and $\mathcal{NE}$, is closer to the true distribution μ. Since both solution concepts are defined in terms of moment inequality restrictions, we can make use of the test for moment inequality models developed by Shi (2010). Consider the null hypothesis:

$$H_0 : d(\mathcal{CE}, \mu) = d(\mathcal{NE}, \mu) \tag{45}$$

where

$$d(\mathcal{CE}, \mu) = \inf_{CE \in \mathcal{CE}} d(CE, \mu);\; d(\mathcal{NE}, \mu) = \inf_{NE \in \mathcal{NE}} d(NE, \mu). \tag{46}$$

The distance $d(P, \mu)$ is defined as the Kullback-Leibler divergence measure:

$$d(P, \mu) = \int p_\mu \log p_\mu d\mu, \tag{47}$$

where p_μ is the density of P with respect to μ. We now construct the test statistics. For a data distribution μ, define the Lagrange multipliers:

$$\gamma_\mu^*(\theta) = \arg\min_\gamma \exp(\gamma' m(X_i, \theta)) \tag{48}$$

$$\lambda_\mu^*(\beta) = \arg\min_\lambda \exp(\lambda' g(X_i, \beta)) \tag{49}$$

[11] The moment conditions in Ciliberto and Tamer (2009) are based on pure strategy Nash equilibrium, to obtain the moment conditions for Nash equilibrium, one needs to extend that result. Berry and Tamer (2007) briefly discuss the problems that arise when allowing for mixed strategies.

and criterion functions:

$$\mathcal{M}_\mu(\gamma, \theta) = E_\mu \exp(\gamma' m(X_i, \theta)) \tag{50}$$

$$\mathcal{G}_\mu(\lambda, \beta) = E_\mu \exp(\lambda' g(X_i, \beta)) \tag{51}$$

Shi (2010) prove that the null hypothesis (45) can be stated as:

$$H_0 : \max_{\theta \in \Theta} \mathcal{M}_\mu(\gamma^*_\mu(\theta), \theta) = \max_{\beta \in B} \mathcal{G}_\mu(\lambda^*_\mu(\beta), \beta) \tag{52}$$

The sample analog of $\mathcal{M}_\mu(\gamma^*_\mu(\theta), \theta)$ and $\mathcal{G}_\mu(\lambda^*_\mu(\beta), \beta)$ are:

$$\hat{\mathcal{M}}_n(\gamma, \theta) = \frac{1}{n}\sum_{i=1}^{n} \exp(\gamma' m(X_i, \theta));\ \hat{\mathcal{G}}_n(\lambda, \beta) = \frac{1}{n}\sum_{i=1}^{n} \exp(\lambda' g(X_i, \beta)) \tag{53}$$

where

$$\begin{aligned} \hat{\gamma}_n(\theta) &= \arg\min_{\gamma} \hat{\mathcal{M}}_n(\gamma, \theta), \quad & \hat{\lambda}_n(\beta) &= \arg\min_{\lambda} \hat{\mathcal{G}}_n(\lambda, \beta) \\ \hat{\Theta}_n &= \arg\max_{\theta \in \Theta} \hat{\mathcal{M}}_n(\hat{\gamma}_n(\theta), \theta) \quad & \hat{B}_n &= \arg\max_{\beta \in B} \hat{\mathcal{G}}_n(\hat{\lambda}_n(\beta), \beta) \end{aligned} \tag{54}$$

Then we can use the quasi-likelihood ratio statistic

$$QLR_n = \max_{\theta \in \Theta} \hat{\mathcal{M}}_n(\hat{\gamma}_n(\theta), \theta) - \max_{\beta \in B} \hat{\mathcal{G}}_n(\hat{\lambda}_n(\beta), \beta) \tag{55}$$

to test the null hypothesis, equation(52).

With several regularity conditions, Shi (2010) proves that under H_0:

$$QLR_n \overset{d}{\rightsquigarrow} N(0, \varpi^2_n) \tag{56}$$

where $\varpi^2_n = E_\mu[\exp(\gamma^*_\mu(\theta^*)' m(X_i, \theta^*)) - \exp(\lambda^*_n(\beta^*)' g(X_i, \beta^*))]^2$, $\theta^* \in \arg\max_{\theta \in \Theta} \mathcal{M}_\mu(\gamma^*_\mu(\theta), \theta)$, $\beta^* \in \arg\max_{\beta \in B} \mathcal{G}_\mu(\lambda^*_\mu(\beta), \beta)$. In practice, ϖ^2_n can be replaced with its sample analog $\hat{\varpi}_n^2$:

$$\hat{\varpi}_n^2 = \sup_{\theta \in \hat{\Theta}_n, \beta \in \hat{B}_n} \frac{1}{n} \sum_{\mu} [\exp(\hat{\gamma}_n(\theta)' m(X_i, \theta)) - \exp(\hat{\lambda}_n(\beta)' g(X_i, \beta))]^2. \tag{57}$$

The test criterion is

Test of Correlated Equilibrium versus Nash Equilibrium Let b_n be a sequence of positive numbers such that $b_n^{-1} + n^{-1} b_n \to 0$. Given the nominal size α and the $(1-\alpha/2)$ quantile of the standard normal distribution, $z_{\alpha/2}$.

(1) If $n\hat{\varpi}_n^2 > b_n$ and $n^{\frac{1}{2}} QLR_n / \hat{\varpi}_n > z_{\alpha/2}$, then reject H_0 in favor of the hypothesis that correlated equilibrium is the appropriate equilibrium concept.

(2) If $n\hat{\varpi}_n^2 > b_n$ and $n^{\frac{1}{2}} QLR_n / \hat{\varpi}_n < -z_{\alpha/2}$, then reject H_0 in favor of the hypothesis that Nash equilibrium is the appropriate equilibrium concept.

(3) If $\hat{\varpi}_n^2$ and $n^{\frac{1}{2}} QLR_n / \hat{\varpi}_n$ do not satisfy the condition in (1) and (2), then do not reject the null H_0.

This test criterion is based on the nested model selection test of Shi (2010). Recall Proposition 2.2, which states that the set of Nash equilibrium payoffs is a subset of the set of correlated equilibrium payoffs. An interesting case is when the test does not reject the null hypothesis. This implies that correlated equilibrium and Nash equilibrium are equally effective in explaining the observed data, or, put differently, the set of Nash equilibria is (approximately) equal to the set of correlated equilibria. Finally, the nesting of Nash equilibrium with correlated means that using correlated equilibrium to estimate empirical games is robust to the true state being that the appropriate solution concept is Nash equilibrium, but inefficient in this state. If instead, the true equilibrium concept is correlated equilibrium, then except in the special case that all correlated equilibria of the game are Nash equilibria , using Nash equilibrium to estimate the game will produce inconsistent estimates. This is similar to the choice between fixed effect and random effect in panel data models.

6 Monte Carlo Simulation

To demonstrate the performance of our estimates in finite samples, we conduct a Monte Carlo experiment using a simple static 2×2 entry game. In each of the T repetitions of the static complete information game, each has the following structure:

	0	1
0	$(0,0)$	$(0, \epsilon_2(0,1))$
1	$(\epsilon_1(1,0), 0)$	$(\theta_1 + \epsilon_1(1,1), \theta_2 + \epsilon_2(1,1))$

The action set of each player is $A_i = \{0, 1\}$, where 0 means no entry and 1 means entry. The utility function for player i is defined as:

$$u_i(a, \epsilon_i(a); \theta) = I(a_i = 1)(\theta_1 a_{-i} + \epsilon_i(a)) \tag{58}$$

As a simple experiment, we have not included any exogenous covariates x here. In accordance with the location and scale normalization requirement for identification, we set the utility of no entry equal to 0 and the variance of the random preference shock equal to 1. Thus, we only need to estimate the strategic effect parameters θ_1 and θ_2.

All random preference shocks $\epsilon_{1t}(1,0)$, $\epsilon_{2t}(0,1)$, $\epsilon_{1t}(1,1)$ and $\epsilon_{2t}(1,1)$ are independently drawn from a standard normal distribution. The parameter space Θ is set to $\Theta = [-5, 5]^2$, and the true values are

$$\theta_1 = -0.5; \ \theta_2 = -1 \tag{59}$$

Thus, entry by player i will decrease the payoff of player j given entry by player j. Given these random shocks and parameters, we generate the outcome of each game, i.e., the observed action profiles, by a simple maximum entropy equilibrium selection mechanism that solves

$$\begin{aligned} &\max_{\pi} -\sum_{a\in A}\pi(a)\ln\pi(a) \\ &s.t.\begin{cases} \sum_{a_{-i}} u_i(a_{-i},a_i)\pi(a_{-i},a_i) \geq \sum_{a_{-i}} u_i(a_{-i},\tilde{a}_i)\pi(a_{-i},a_i),\ \forall i, a_i \text{ and } \tilde{a}_i \neq a_i \\ \sum_{a\in A}\pi(a)=1,\ \pi(a)\geq 0. \end{cases} \end{aligned} \tag{60}$$

Obviously, this maximum entropy equilibrium selection mechanism will generate the most dispersive correlated equilibrium π^* among all correlated equilibria. We use simple random sampling on π^* to determine which action profile will be played. Based on the maximum entropy equilibrium selection mechanism, with a sample size of 500,

$$E[\Pr(\boldsymbol{y}_t)] = (0.291058, 0.274005, 0.35475, 0.080187). \tag{61}$$

We use the following procedure to estimate Equation (33). First, we use a simulated annealing algorithm to find a solution to the minimization of Equation (32), which we denote by $\tilde{\theta}$. Then we use rich directions[12] to grid search within the parameter space Θ until it condition (33) is satisfied.

We generate 1000 samples of size $T = 500, 1000$ to assess the finite sample properties of our estimator. We first use the importance sampling simulator to get simulated bounds of choice probabilities, then, based on the above numerical procedure compute the final estimates. The interval estimates are reported in Table 1, and the set estimators for $T = 500$ and $T = 1000$ are compared in Figure 1. Since we lack information regarding the true range of the outer identified set, we can not say much about the performance of our

[12] In this experiment, we choose 402 directions, which are randomly chosen according to a uniform distribution over $[0, 2\pi]$.

Table 1: The Results of the Monte Carlo Simulations

	Initial Value For Set Estimation	Interval Estimates
	$T = 500$	
θ_1	-0.5677	$[-2.3142, 0.9580]$
θ_2	-1.0273	$[-2.8319, 0.4027]$
	$T = 1000$	
θ_1	-0.5659	$[-2.2794, 0.9370]$
θ_2	-1.0267	$[-2.8184, 0.3816]$
	Monte Carlo Repetitions: 1000	
	Importance Sampling Repetitions: 999	

estimator, except that the true value of the parameter lies in our estimated set. Moreover, from Figure 1 we see that when the sample size increase, the range of $\hat{\Theta}_O$ decreases, which is similar to convergence in the point identified case, but that the size of the set remains large at this sample size.

7 Conclusion

In this paper, we propose a framework for identifying and estimating the normal-form payoff parameters of a discrete, static, complete information game where the equilibrium concept employed is correlated equilibrium. Compared with existing studies based on Nash equilibrium, this extension of the equilibrium concepts simplifies the identification and estimation of game-theoretic econometric models, since our approach does not require the computation of the full set of equilibria, it only needs to compute some "extreme" equilibria which can be obtained through linear programming. We deal with the presence of multiple equilibria, by making use of the moment inequality restrictions induced by the underlying game-theoretic econometric models, rather than making any assumptions regarding equilibrium selection Thus we avoid the potential for misspecification of the equilibrium selection mechanism.This leads to the estimation of a partially identified model. Given the outer identified set characterized by moment restrictions, the set

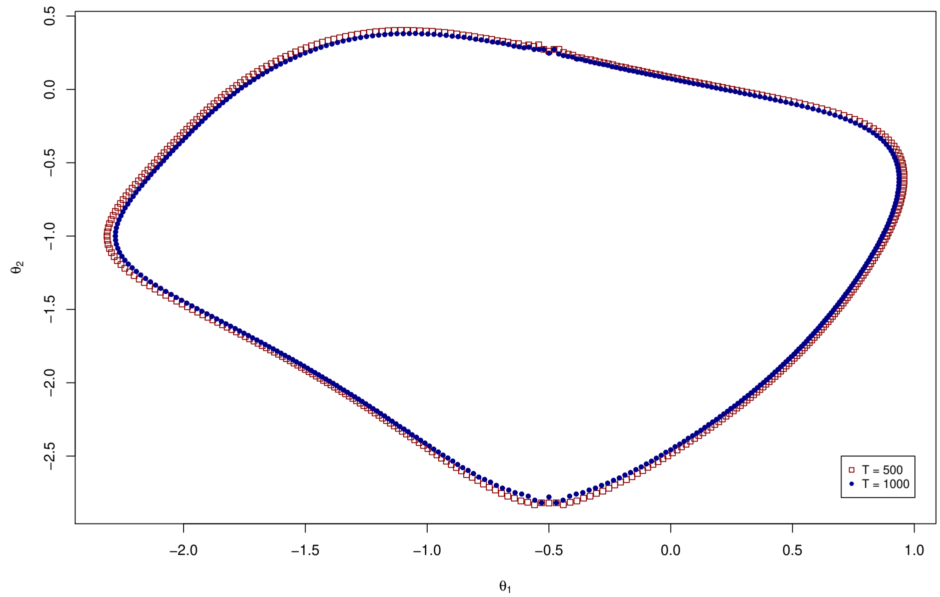

Figure 1: The estimated outer identified set $\hat{\Theta}_O$ for different sample sizes.

estimator developed by Chernozhukov, Hong, and Tamer (2007) is used to obtain its estimates. The importance sampling technique is used to reduce computational burden and overcome the non-smoothness problems. We also show that the model selection tests for moment inequality models developed by Shi (2010) can be used to test equilibrium concepts such as correlated equilibrium versus Nash equilibrium. The greatest limitation of our estimation method is that it requires the distribution of random preference shocks to be known by the researcher. Estimation of such models in ignorance of the the distribution of random preference shocks is an important topic for future research. Another possible extension is to update our estimator to one based on conditional moment restrictions.

References

Ackerberg, D. A. (2009): "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*, 7(4), 343 – 376.

Aguirregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75, 1–53.

Aradillas-Lopez, A. (2007): "Pairwise Difference Estimation of Incomplete Information Games," *Working Paper*.

——— (2010): "Semiparametric Estimation of a Simultaneous Game with Incomplete Information," *Journal of Econometrics*, 157(2), 409–431.

Aumann, R. J. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1(1), 67–96.

——— (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55(1), 1–18.

Bajari, P., H. Hong, J. Krainer, and D. Nekipelov (2010): "Estimating Static Models of Strategic Interactions," *Journal of Business and Economic Statistics*, 28(4), 469–482.

Bajari, P., H. Hong, and S. Ryan (2010): "Identification and Estimation of a Discrete Game of Complete Information," *Econometrica*, 78(5), 1529–1568.

Bergin, J. (2005): *Microeconomic Theory: A Concise Course*. Oxford University Press.

Berry, S., and E. Tamer (2007): "Identification in Models of Oligopoly Entry," in *Advanced in Economics and Econometrics: Theory and Application*, vol. II, chap. 2, pp. 46–85. Cambridge University Press.

Berry, S. T. (1992): "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 60(4), 889–917.

Bresnahan, T. F., and P. C. Reiss (1990): "Entry in Monopoly Markets," *Review of Economic Studies*, 57, 531–553.

——— (1991): "Empirical models of Discrete Games," *Journal of Econometrics*, 48, 57–81.

Chernozhukov, V., H. Hong, and E. Tamer (2007): "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, 75(5), 1243–1284.

Chwe, M. S. (2007): "Identifying a Game by Observing its Correlated Equilibria," *Working Paper*.

Ciliberto, F., and E. Tamer (2009): "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 77(6), 1791–1828.

Dhillon, A., and J. F. Mertens (1996): "Perfect Correlated Equilibria," *Journal of Economic Theory*, 68(2), 279–302.

Forges, F. (2009): "Correlated equilibrium and communication in games," in *Encyclopedia of Complexity and Systems Science*, ed. by R. Meyers, pp. 1587–1596. Springer.

Golan, A., L. S. Karp, and J. M. Perloff (2000): "Estimating Coke's and Pepsi's Price and Advertising Strategies," *Journal of Business and Economic Statistics*, 18(4), 398–409.

Hausman, J. A. (1978): "Specification Tests in Econometrics," *Econometrica*, 46(6), 1251–71.

Jia, P. (2008): "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 76, 1263–1316.

McFadden, D. (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica*, 57(5), 995–1026.

McKelvey, R., and A. McLennan (1996): "The computation of equilibrium in finite games," in *The Handbook of Computational Economics*, ed. by H. Amman, D. A. Kendrick, and J. Rust, vol. 1996. Elsevier Science, The Netherlands.

Morris, S., and H. Shihn (2003): "Global Games: Theory and Applications," in *Advances in Economics and Econometrics*, ed. by M. Dewatripont, L. Hansen, and S. Turnovsky, pp. 56–114. Cambridge University Press.

Ortiz, L. E., R. E. Schapire, and S. M. Kakade (2007): "Maximum Entropy Correlated Equilibria," *Eleventh International Conference on Artificial Intelligence and Statistics*.

Osborne, M. J., and A. Rubinstein (1994): *A Course in Game Theory*. The MIT Press.

Papadimitriou, C., and T. Roughgarden (2008): "Computing Correlated Equilibria in Multi-Player Games," *The Proceedings of the 37th Annual Symposium on the Theory of Computing*.

Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, 75, 901–928.

Seim, K. (2006): "An Empirical Model of Firm Entry with Endogenous Product-type Choices," *RAND Journal of Economics*, 37(3), 619–640.

Shi, X. (2010): "Model Selection Tests for Nonnested Moment Inequality Models," *Job Market Paper of Yale University*.

Sorin, S. (1997): "COMMUNICATION, CORRELATION AND COOPERATION," in *Cooperation: Game-theoretic Approaches*, ed. by S. Hart, and A. Mas-Colell, pp. 199–218. Springer.

Sweeting, A. (2005): "Coordination Games, Multiple Equilibria and the Timing of Radio Commercials," *Working Paper.*

Tamer, E. T. (2003): "Incomplete Bivariate Discrete Response Model with Multiple Equilibria," *Review of Economic Studies*, 70, 147–167.

Yang, Z. (2007): "Essays in Empirical Industrial Organization," Ph.D. thesis, University of Michigan.