Inference in Repeated Games with Random States*

Arek Szydłowski

University of Leicester

February 17, 2017

Abstract

We provide methods for empirical analysis of repeated games with i.i.d. shocks. The number of possible equilibria in these games is large and, usually, theory is silent about which equilibrium will be chosen in practice. Thus, our method remains agnostic about selection among these multiple equilibria, which leads to partial identification of the parameters of the game. We propose profiled likelihood ratio test for building confidence sets for these parameters and use multiplier bootstrap to obtain the critical value for the test.

1 Introduction

The theory of repeated strategic interactions has been thoroughly studied by economists. However, the empirical analysis of repeated games is scarce. Repeated games usually feature many equilibria making their analysis difficult. Abito & Szydlowski (2017) are the first to investigate identification in stochastic repeated games under the assumption of no state dependence and without making, necessarily, arbitrary assumptions on the choice mechanism between these equilibria. They develop convenient but not sharp characterization of the identified set in these games. In this paper we use a sharp maximum likelihood characterization and develop inferential methods for estimating confidence sets for the parameters of interest using a multiplier bootstrap profiled likelihood ratio test. We derive the asymptotic distribution of the test statistic.

*We would like to thank Mike Abito, Wojciech Olszewski, Subir Bose and Andre Stenzel for useful discussion and suggestions.

Our method can be applied in many contexts. For example, Ciliberto & Tamer (2009) analyse entry and exit into the US airline markets using a static game allowing for multiple equilibria. They use data from the second quarter of 2001 and view it as a long-run outcome of strategic interaction between the carriers, thus arguing that the static game approach is appropriate. However, this interaction can also be seen as repeated game, where players repeatedly make decisions about the presence in particular market. In fact, as Table 1 shows, there is a lot of entry and exit in the US airline market. For example, American Airlines dropped out of 25% of the markets they served in 1995, even more strikingly, in 2015 they came back to 12% of their 1995 markets after not serving them in 2005. Thus, the network of served markets is far from stable and carriers often react to changing market conditions by exit and entry.

Table 1: Share of markets by presence of major carriers over time (in %)

Presence in Q2 1995 - Q2 2005 - Q2 2015	American Airlines	Delta	United	US Airways
in - in - in	57.91	64.02	83.77	42.83
in - out - in	12.39	4.73	4.05	14.36
in - out - out	24.77	20.28	6.67	33.71

Note: Markets are defined as directional routes between origin and destination airports (irrespective of the number of stops on the way). "In" means that a carrier served at least one flight on the route.

Similarly, our method can be employed in the analysis of entry and exit into procurement auctions. These auctions are usually repeated over time, often with the same parties involved. Bajari, Hong & Ryan (2010) analysed entry decisions in these auctions using a static game model, whereas our method treats such auction as a repeated game where the unobserved component of the payoff is i.i.d. over time.

Our article is related to the literature on estimation of dynamic games. However, it differs from the approach in Bajari, Benkard & Levin (2007) (abbreviated by BBL) in important aspects. We focus on a simple form of a dynamic game without state dependence and with complete information. In this subset of dynamic games we do not have to restricts ourselves to Markov perfect equilibria, allowing for a rich set of equilibrium strategies potentially dependent on the whole history of play (in fact, Markov perfect equilibria in this game correspond to stage game equilibria, thus Markov restriction makes the dynamic dimension irrelevant).

In principle, our method can be extended to general dynamic games with non-trivial state dependence, however implementing such an extension poses a difficult computational challenge so we leave it for further research.

2 Repeated Game with Random States

A defining feature of our model is that the agents, unlike in a standard repeated game, play a different game every period. The payoffs in the stage game are stochastic and distributed i.i.d. over time (conditional on observables). Payoffs are fully observed by the players but the econometrician observes only the history of play and has only limited information about the non-stochastic part of the payoff. In particular, she does not observe the realization of the i.i.d. shocks.[1]

We focus on binary games. Let $A = \{0,1\}^2$ denote the set of actions with a typical element a. There are two players identified by $i = 1, 2$ playing an infinitely repeated game, $t = 0, 1, 2 \ldots, \infty$. Every period a random payoff relevant vector $\varepsilon_t(a) = (\varepsilon_{1t}(a), \varepsilon_{2t}(a))$ is drawn from the distribution $F_\varepsilon : \mathcal{E} \to [0,1]$ with the following properties:

Assumption 1. *$\varepsilon_t(a)$ are i.i.d. across time and $E(\varepsilon_t(a)) = \mathbf{0}$ for all a and $\varepsilon_t(\mathbf{0}) = \mathbf{0}$ (normalization);*

Denote player i's payoff in period t by $\tilde{u}_i(a_t, \varepsilon_{it}(a_t); \alpha_i)$, where α_i is a finite dimensional parameter. Define $\alpha = (\alpha_1, \alpha_2)$. In practice some observed characteristics $X \in \mathcal{X} \subset R^K$ will usually enter the payoff function. For now we suppress them from notation for the ease of exposition. We assume that the payoffs are additive in the random shock:

$$\tilde{u}_i(a_t, \varepsilon_{it}(a_t); \alpha_i) = u_i(a_t; \alpha_i) - \varepsilon_{it}(a_t)$$

This assumption is usually imposed in the empirical dynamic games literature (see Assumption DC in BBL). Payoffs in future periods are discounted by a common discount factor $\delta \in (0, 1)$.

Let $\mathcal{H}^t = (\mathcal{E} \times A)^t$ denote the set of period t ex ante histories with a typical element $h^t = \{\varepsilon_s(a_s), a_s\}_{s=0}^{t-1}$. Let $\tilde{\mathcal{H}}^t = (\mathcal{E} \times A)^t \times \mathcal{E}$ denote the set of period t ex post histories with a typical element $\tilde{h}^t = (h^t, \varepsilon_t)$. A pure strategy profile, σ, is a pair of mappings from $\tilde{\mathcal{H}}^t$ to A. Player i's

[1] For the rest of the paper we will simply refer to our repeated game with random states as a "repeated game".

expected lifetime payoff from playing the game is:

$$U_i(\sigma;\alpha_i) = E^{\sigma}\left[\sum_{t=0}^{\infty}\delta^t \tilde{u}_i(a_t, \varepsilon_{it}(a_t);\alpha_i)\right]$$

where the expectation is taken over the histories induced by σ.

A normalized continuation payoff of player i after a history $\tilde{h}^t$ is given by:

$$V_i^{\sigma}(\tilde{h}^t;\alpha_i) = (1-\delta)u_i(\sigma(\tilde{h}^t);\alpha_i) - (1-\delta)\varepsilon_{it}(\sigma(\tilde{h}^t)) + \delta\int V_i^{\sigma}(\tilde{h}^{t+1};\alpha_i)dF_{\varepsilon}$$

3 Identification

Throughout our analysis we will assume that

Assumption 2. ***(a)*** *δ and F_{ε} are known,* ***(b)*** *$u(a_i, a_{-i};\alpha) = 0$ for some $a_i \in A$ and all $a_{-i} \in A$*

This assumption is frequently made in dynamic discrete choice and dynamic games literature. In the case of dynamic discrete choice models, Magnac & Thesmar (2002) showed that the payoff functions are not identified without knowledge of the discount factor and distribution of shocks or without normalizing the payoffs for one of the alternatives.

We assume that we observe finite history of play, h^t, in at least one game (market) and the observed play is generated by some subgame-perfect equilibrium. In principle, characterizing the set of equilibria in our game is complicated since there are numerous strategies that can generate a given equilibrium outcome. Fortunately, we do not have to work with **strategies** directly but rather characterize sufficient and necessary condition(s) for observed **choices** to be made in equilibrium. Let $p_t(a)$ denote the probability of a being played in equilibrium at some t. We will restrict this probability to be stationary.

Assumption 3. *Players play a stationary subgame-perfect equilibrium, i.e. $p_t(a) = p(a)$ for all t.*

In principle, we could work with non-stationary equilibria. However, this would require that we observe the same game being played in many different markets such that we can calculate $p_t(a)$ for every t by looking at probability of a being played across markets. Such data is hard to come by since usually payoffs differ between markets due to different market characteristics. In some games

Assumption 3 rules out some important equilibria, e.g. some efficient symmetric equilibria (see Remark 5.7.1 in Mailath & Samuelson (2006)), thus one needs to be careful if it's not too strong in a particular game of interest.

Proposition 1. *Let* $\mathcal{V}_S(\alpha)$ *denote the set of pairs of expected lifetime payoffs that can be reached in a stationary equilibrium, i.e. for any stationary subgame-perfect equilibrium strategy* σ *we have* $(U_1(\sigma;\alpha_1), U_2(\sigma;\alpha_2)) \in \mathcal{V}_S(\alpha)$*, and* $v(a) = (v_1(a), v_2(a)) \in \mathcal{V}_S(\alpha)$ *denote expected continuation payoffs associated with playing* a *in the current period. Define* $V = (v(a), v(a_1', a_2), v(a_1, a_2'), v(a_1', a_2'))$ *with* $a_1 \neq a_1', a_2 \neq a_2'$*. Then, for some distribution* $F_V : \mathcal{V}_S(\alpha)^4 \to [0,1]$*:*

$$p(a) = \int_{\mathcal{V}_S(\alpha)^4} P(a \text{ is a Nash equilibrium in the normal form game with payoffs } g_i(a)|V = v) dF_V(v) \tag{1}$$

where:

$$g_i(a) = (1-\delta)(u_i(a;\alpha_i) - \varepsilon_i(a)) + \delta v_i(a),$$

and $(\varepsilon_1(a), \varepsilon_2(a))$ *are drawn from* F_ε*.*

Proof. Fix V. Proposition 5.7.3 in Mailath & Samuelson (2006) implies that action profile a is played in period t as part of a subgame-perfect equilibrium if and only if a is a Nash equilibrium of a normal form game with payoffs:

$$g_i(a) = (1-\delta)(u_i(a;\alpha_i) - \varepsilon_i(a)) + \delta v_i(a)$$

Now it remains to integrate with respect to V using the equilibrium selection distribution F_V to obtain the unconditional probability of observing a being played in equilibrium. □

The distribution F_V can be interpreted as an equilibrium selection function since every $v \in \mathcal{V}_S(\alpha)$ corresponds to a different subgame-perfect equilibrium in our stochastic game. If data comes only from a single game and mixing between different equilibria is not allowed, then F_V is degenerate on some particular V. However, if data comes from different games (markets) and we allow different equilibria to be played in different markets, then F_V may put non-trivial mass on several V's corresponding to different markets.

Proposition 1 is useful for two reasons. First, it allows us to describe the identified set in the model without dealing with strategy functions, which belong to a complicated functional space. Instead, in order to verify if a given parameter value α can be reconciled with observed probabilities, it is enough to check if this probability can be generated in equilibrium by some (combination of) continuation payoffs in $\mathcal{V}_S(\alpha)$[4]. Second, calculating the sets $\mathcal{V}_S(\alpha)$ is relatively easy for a game with discretely supported independent shocks and we can approximate these sets for a game with continuous F_ε by increasing the number of support points of ε. We discuss methods for obtaining these sets in the next section.

As a result of Proposition 1, we have the following corollary:

Corollary 1. *Suppose that the econometrician observes the probabilities $p(a)$ for each $a \in A$. Define the identified set Θ^S_{01} as the set of α's for which observed probabilities are consistent with some stationary subgame-prefect equilibrium. Then:*

$$\Theta^S_{01} = \{\alpha : \exists F_V \quad \forall_{a \in A} \text{ condition (1) holds}\}$$

This set is sharp.

To illustrate the identification argument consider a simple game with the following stage game payoffs (for further reference, we will call this game Σ^S):

		P 2: 1	P 2: 0
P 1	1	$\alpha_1 - \varepsilon_1, \alpha_2 - \varepsilon_2$	$-\varepsilon_1, 0$
	0	$0, -\varepsilon_2$	$0, 0$

where ε_1 and ε_2 are identically and independently distributed (with some abuse of notation, let F_ε denote their distribution). The corresponding normal form of the repeated game is given by:

		P 2: 1	P 2: 0
P 1	1	$(1-\delta)(\alpha_1 - \varepsilon_1) + \delta v_1(1,1), (1-\delta)(\alpha_2 - \varepsilon_2) + \delta v_2(1,1)$	$-(1-\delta)\varepsilon_1 + \delta v_1(1,0), \delta v_2(1,0)$
	0	$\delta v_1(0,1), -(1-\delta)\varepsilon_2 + \delta v_2(0,1)$	$\delta v_1(0,0), \delta v_2(0,0)$

Let $\alpha_1 = \alpha_2 = \alpha > 0$. By analogy to static binary games (see Tamer (2003)), we may have multiple Nash equilibria in this normal form. For example, assuming that players play a static Nash equilibrium in every period, i.e. $v_i(a) = 0$ for all i and a, both (1,1) and (0,0) can arise in equilibrium when ε_1 and ε_2 are between $-\alpha$ and 0.

Figure 1: Multiple equilibria in the normal form

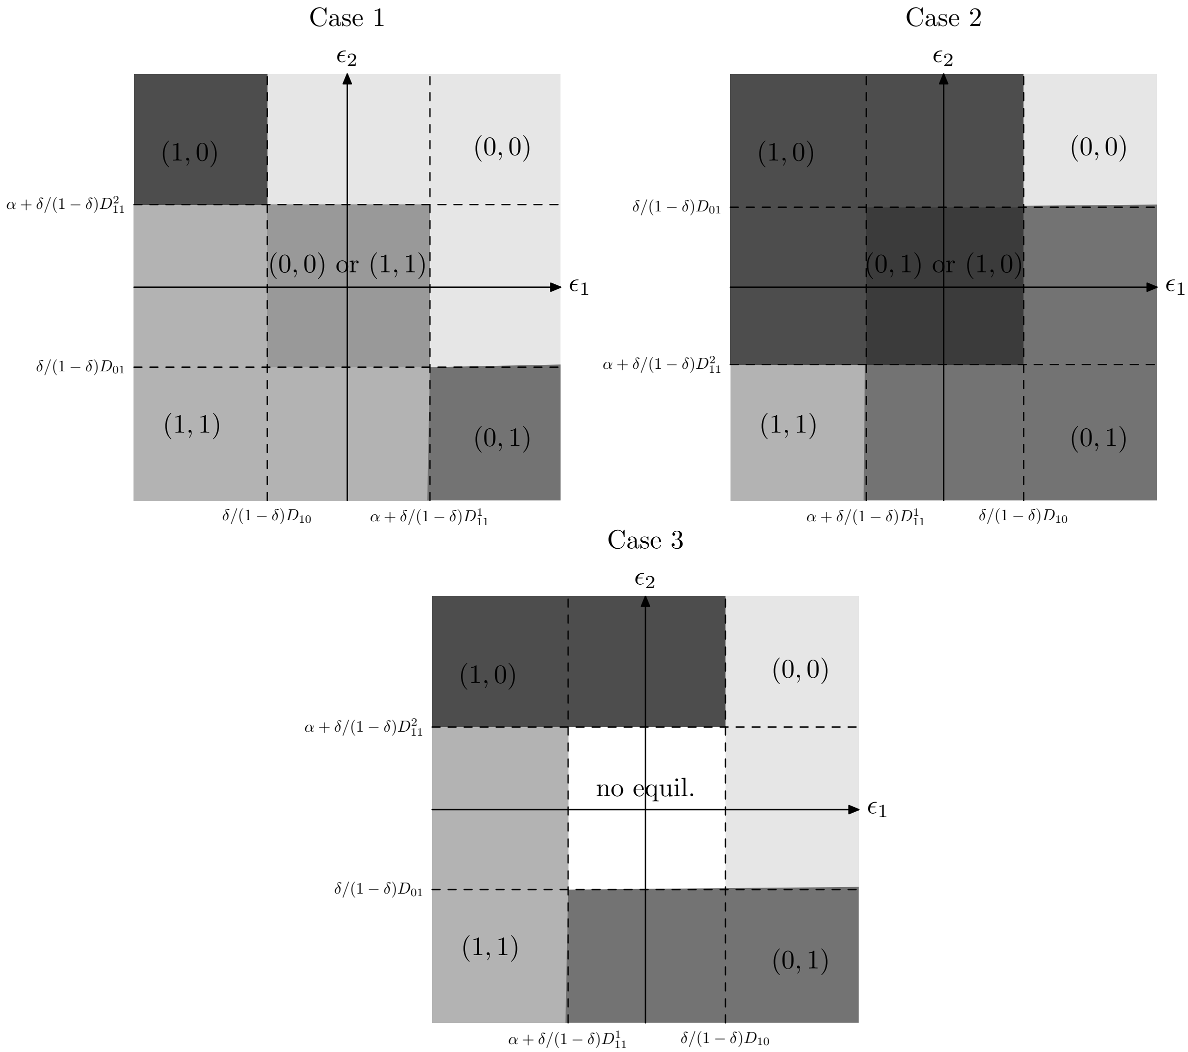

In general, the region of ε for which we will have multiple equilibria depends on V, and more specifically value contrasts, D_a^i:

$$D_{10} = v_1(1,0) - v_1(0,0), \quad D_{01} = v_2(0,1) - v_2(0,0)$$

$$D_{11}^1 = v_1(1,1) - v_1(0,1), \quad D_{11}^2 = v_2(1,1) - v_2(1,0)$$

Figure 1 shows regions of values of $(\varepsilon_1, \varepsilon_2)$ for which different equilibria occur for different values of D^i_a's. There are several regions where the game has two equilibria. Let us focus on Case 1 and the probability that (1,1) is played in subgame-perfect equilibrium. For simplicity assume that only symmetric equilibria are played in this game. Let s denote the probability that (0,0) is played in equilibrium in the region where both (0,0) and (1,1) can be played. Then, we can write the probability of (1,1) being played in this repeated game, conditional on the continuation values V as:

$$p(1,1|V) = \left[F_\varepsilon\left(\alpha + \frac{\delta}{1-\delta}D_{11}\right)\right]^2 - s\left[F_\varepsilon\left(\alpha + \frac{\delta}{1-\delta}D_{11}\right) - F_\varepsilon\left(\frac{\delta}{1-\delta}D_{01}\right)\right]^2$$

Thus, conditional on V the problem boils down to a standard identification problem in static games, where we may have multiple equilibria and we do not observe the equilibrium selection probability s. However, in our repeated game lack of identification is magnified by the fact that we observe neither s nor V. In the numerical example in Section 5 we show that this does not preclude us from obtaining meaningful bounds on the model parameters.

4 Equilibrium continuation payoff sets

In this section we illustrate how one can obtain the sets $\mathcal{V}_S(\alpha)$. In fact, we will show how to obtain an outer approximation to this set, $\mathcal{V}(\alpha)$, since we will allow $\mathcal{V}(\alpha)$ to contain continuation payoffs both from stationary and non-stationary equilibria. Having calculated $\mathcal{V}(\alpha)$, we can characterize an outer set of the identified set by

$$\Theta_{01} = \{\alpha : \exists F_V \quad \forall_{a\in A} \text{ condition (1) holds with } \mathcal{V}_S(\alpha) \text{ replaced with } \mathcal{V}(\alpha)\}$$

We will show using numerical examples that, even though $\Theta^S_{01} \subset \Theta_{01}$ the outer set obtained in this way may still be quite narrow for some games of interest.

With finite discrete support of shocks we can view games with different draws of ε as separate non-stochastic repeated games and use the algorithms introduced in Abreu, Pearce & Stacchetti (1990) and Abreu & Sannikov (2014) to find the sets of equilibrium payoffs in these games, $\mathcal{V}_{\varepsilon^m}(\alpha), m = 1, \ldots, M$. Then, the set of equilibrium payoffs in the full game can be calculated as a

Minkowski sum:

$$\mathcal{V}(\alpha) = \{V : V = V_1 \cdot P(\varepsilon = \varepsilon^1) + \ldots + V_M \cdot P(\varepsilon = \varepsilon^M), V_1 \in \mathcal{V}_{\varepsilon^1}(\alpha), \ldots, V_M \in \mathcal{V}_{\varepsilon^M}(\alpha)\}$$

(see Remark 5.7.1 in Mailath & Samuelson (2006)). We will sometimes refer to this set of equilibrium payoffs as the '$\mathcal{V}$ set'.

Figure 2: Set of equilibrium payoffs, $\mathcal{V}$

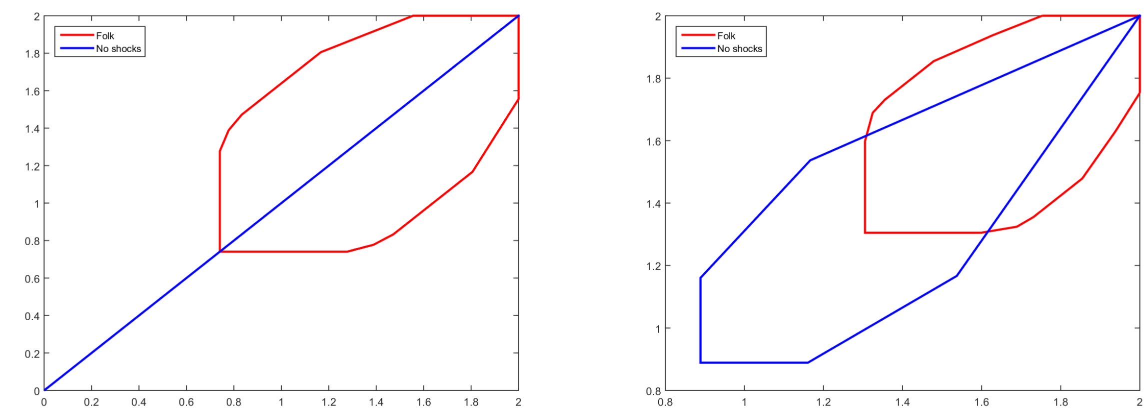

Note: $\alpha = 1, \delta = 0.75$. Blue line demarcates the set of payoffs in the non-stochastic version of the game with $\varepsilon_1 = \varepsilon_2 = 0$. Red line delineates the set in the stochastic game.

We will illustrate the construction of this set using the game Σ^S and a simplified version of an entry game in which entrants engage in Cournot competition (see Tamer (2003), henceforth, referred to as "Cournot entry game"), both with discrete support of shocks over $\mathcal{E} = \{-2, 0, 2\}^2$. The stage game payoffs in the latter game are given by:

		P 2: 1	P 2: 0
P 1	1	$\alpha - \varepsilon_1, \alpha - \varepsilon_2$	$\frac{9}{4}\alpha - \varepsilon_1, 0$
	0	$0, \frac{9}{4}\alpha - \varepsilon_2$	$0, 0$

Figure 2 compares the equilibrium payoff sets of the non-stochastic and stochastic version of game Σ^S (left panel) and the Cournot entry game (right panel). Note that due to symmetry the $\mathcal{V}$ set in Σ^S without shocks is just a straight line connecting stage game payoffs in the two Nash equilibria. However, once we add the stochastic shocks the game is not necessarily symmetric and the $\mathcal{V}$ set will have a non-empty interior for some values of ε. As a result, the $\mathcal{V}$ set in the stochastic game has a non-empty interior as well. The $\mathcal{V}$ sets in the stochastic game exclude some low payoff pairs from the non-stochastic version since now players can condition their strategies on the realizations of the shocks and, as a result, achieve better outcomes.

We can view a game with finite support of shocks as an approximation of a game with continuous F_ε where the approximation becomes more precise as the number of support points, M, grows. For the purpose of inference we can make $M \to \infty$ as the sample size $n \to \infty$.

Finally, it's worth noting that, although the probabilities $p(a)$ depend on only on the the continuation value contrasts, e.g. $D^1_{11} = v_1(1,1) - v_1(1,0), D^2_{11} = v_2(1,1) - v_2(0,1)$, working with the value contrasts as a primitive of the game is not that convenient because the restriction $V \in \mathcal{V}(\alpha)^4$ may put some non-trivial restrictions on the set of equilibrium D's. For example, setting $D^1_{11} = \max_{(v_1,v_2)\in\mathcal{V}} v_1 - \min_{(v_1,v_2)\in\mathcal{V}} v_1$ and $D^2_{11} = \max_{(v_1,v_2)\in\mathcal{V}} v_2 - \min_{(v_1,v_2)\in\mathcal{V}} v_2$ may not be feasible since the points $(\max_{(v_1,v_2)\in\mathcal{V}} v_1, \max_{(v_1,v_2)\in\mathcal{V}} v_2)$ and $(\min_{(v_1,v_2)\in\mathcal{V}} v_1, \min_{(v_1,v_2)\in\mathcal{V}} v_2)$ may not belong to the $\mathcal{V}$ set. Thus, we work with the equilibrium payoffs V rather than value contrasts because this allows us to handle the constraints more easily.

5 Numerical examples

In order to gain some insight about the size of the identified sets obtained using our method, we perform a small computational experiment using the two examples in the previous section. We set $\alpha = 1$ and assume that players play static Nash equilibrium every period. With our three point specification of the shocks the standard deviation of ϵ_i is equal to 1.6α so the noise component of payoffs is significant. In the simple game Σ^S we assume that in the region of multiplicity equilibrium (1,1) is played with probability zero, in the Cournot entry game (0,1) and (1,0) are played with equal probabilities in cases where both equilibria occur. Together, this implies that we have $p(1,0) = p(0,1)$ in both games.

Our numerical exercise consists of finding values of α for which we can match the probabilities under static Nash assumption:

	Σ^S	Cournot entry
p(1,1)	4/9	4/9
p(0,1)	2/9	22/81

and the theoretical probabilities given by (1) where we assume that F_V is degenerate. Note that, since the games are symmetric, it is enough to focus on probability of $(1,1)$ and $(0,1)$ being played (if we match these two probabilities it means we can also match p(1,0) because the problem there is the same as for p(0,1) and we can match $p(0,0) = 1 - 2p(0,1) - p(1,1)$).

Table 2 contains the resulting bounds. We also include the intersection of the bounds obtained from matching $p(1,1)$ and $p(0,1)$, which contains the identified set in the model (for a specific α we may need different values of V for matching probability of (1,1) and (0,1) thus the intersection is an outer set of the identified set).

Table 2: Numerical examples - bounds on α

	game Σ^S		
	p(1,1)	p(0,1)	$p(1,1) \cap p(0,1)$
$\delta = 0.55$	[1, 6.498]	[1, 6.39]	[1, 6.39]
$\delta = 0.75$	[0.978, 14.28]	[0.947, 14.989]	[0.978, 14.28]
$\delta = 0.95$	[0.945, 14.72]	[0.86, 14.795]	[0.945, 14.72]
	Cournot entry game		
	p(1,1)	p(0,1)	$p(1,1) \cap p(0,1)$
$\delta = 0.55$	[0.892, 1.061]	[0.892 ,1.103]	[0.892 ,1.061]
$\delta = 0.75$	[0.903, 1.225]	[0.989 ,1.22]	[0.989 ,1.22]
$\delta = 0.95$	[0.995, 1.32]	[0.997 ,1.324]	[0.997 ,1.32]

Couple of interesting observations arise from our exercise. First, there seems to be no meaningful upper bound on α for large values of the discount factor in game Σ^S. On the other hand, the bounds in the Cournot entry example are pretty tight for all values of δ. This suggests that, despite the large number of equilibria, our approach is capable of delivering quite sharp predictions about parameters of interest. Note that the main difference between the two examples is that there is somehow more structure in the payoffs of the Cournot entry game – the parameter α affects both entry and exit payoffs. This imposes some restrictions on the equilibrium payoff set and helps to shrink the bounds.

6 Inference

We suggest using maximum likelihood for inference. In practical applications some components of the payoff vector are observed, X, and we are interested in estimating their coefficients, $\beta = (\beta_1, \beta_2)$. Thus, all the probabilities below should condition on X. For notational simplicity we suppress this conditioning.

We assume that the discount factor δ is known so the goal is to estimate a confidence set for $\theta_1 \equiv (\alpha, \beta), \theta_1 \in \Theta_1 \subset \mathcal{R}^{d_1}$. We allow for any equilibrium selection mechanism s but restrict F_V to be degenerate. Let $\theta_2 \equiv (V, s) \in \Theta_2(\theta_1)$ where $\Theta_2(\theta_1)$ is the correspondence mapping values of θ_1 to the corresponding continuation values V and selection probabilities s. Further, $\theta = (\theta_1, \theta_2) \in \Theta \subset \mathcal{R}^d, d = d_1 + d_2$. We suggest using a profiled likelihood ratio statistic for inference and building the confidence set as a collection of points not rejected by the LR ratio test.

We observe an i.i.d. sample of action pairs $\{Y_t\}_{t=1}^T$. The likelihood of observing a given Y_t can be written as:

$$p(Y_t, \theta) = p(1,1)^{Y_{1t}Y_{2t}} p(1,0)^{Y_{1t}(1-Y_{2t})} p(0,1)^{(1-Y_{1t})Y_{2t}} p(0,0)^{(1-Y_{1t})(1-Y_{2t})}$$

so the log-likelihood can be expressed as $L_T(\theta) = \sum_{t=1}^T \log p(Y_t, \theta)$.

Define the identified sets for θ and θ_1 as:

$$\Theta_0 = \arg\sup_{\theta \in \Theta} E_0[\log p(Y_t, \theta)] \qquad \text{and} \qquad \Theta_{01} = \arg\sup_{\theta_1 \in \Theta_1} \sup_{\theta_2 \in \Theta_2(\theta_1)} E_0[\log p(Y_t, \theta)]$$

where the expectation E_0 is taken with respect to the true distribution of Y_t.

For a candidate value $\tilde{\theta}_1$ the profiled likelihood ratio statistic is defined as:

$$LR(\tilde{\theta}_1) = 2 \left[\sup_{\theta \in \Theta} L_T(\theta) - \sup_{\theta_2 \in \Theta_2(\tilde{\theta}_1)} L_T(\tilde{\theta}_1, \theta_2) \right] + o_P(1)$$

where the $o_P(1)$ term accommodates approximation error coming from using a discrete grid for ε for calculating $\mathcal{V}$ sets and optimization error. The analysis of our LR statistic is nonstandard since, without identification, the usual quadratic approximation in the Euclidean norm is not helpful. However, Liu & Shao (2003) showed that instead of Euclidean distance one can work with Pearson distance, χ, which equals zero for all the pairs of θ's in the identified set. This makes quadratic approximation in χ-distance a useful tool in establishing the asymptotic distribution of our statistic.

Denote $p_0(Y_t) = p(Y_t, \theta_0)$ for any $\theta_0 \in \Theta_0$. The Pearson distance between θ_1 and θ_2 is given by:

$$\chi^2(\theta_1, \theta_2) = E_0\left[\left(\frac{p(Y_t, \theta_1) - p(Y_t, \theta_2)}{p_0(Y_t)}\right)^2\right]$$

For $\epsilon > 0$ define:

$$\Theta_\epsilon = \{\theta : \chi(\theta, \theta_0) \leq \epsilon\}.$$

Assumption 4. *We have:*

(a) Θ_1 *is a compact, nonempty subset of* $\mathcal{R}^{d_1}$ *and* $\Theta_2(\theta_1)$ *is a compact, nonempty subset of* $\mathcal{R}^{d_2}$ *for every* $\theta_1 \in \Theta_1$.

(b) $E_0[\log p(y_t, \theta)]$ *is upper-semicontinuous on* Θ.

(c) The class of functions $\{f(\cdot, \theta) = \log p(\cdot, \theta) : \theta \in \Theta\}$ *is Glivenko-Cantelli.*

(d) $p(y, \theta)$ *is continuously differentiable in* $\theta \in \Theta$ *for every* y.

(e) For $G(Y_t, \theta) = \frac{\frac{\partial p(Y_t, \theta)}{\partial \theta}}{p_0(Y_t)}$ *we have* $\sup_{\theta \in \Theta_\epsilon} E_0[G(Y_t, \theta)' G(Y_t, \theta)] < M < \infty$ *for some* $\epsilon > 0$.

(f) The class of functions

$$\mathcal{D} = \left\{D(\cdot, \theta, \lambda) = \frac{G(\cdot, \theta)'\lambda}{\sqrt{E_0[(G(\cdot, \theta)'\lambda)^2]}} : \theta \in \Theta_\epsilon, \|\lambda\| = 1\right\}.$$

is Donsker.

The last assumption is crucial for deriving the asymptotic distribution of our statistic. Note that we require the Donsker property to hold in the neighbourhood of the identified set and **on** the set itself, which entails that the limits of $D(\cdot, \theta, \lambda)$ as $\chi(\theta, \theta_0) \to 0$ also have to be stochastically equicontinuous. Alternatively, one could assume that the Donsker condition holds only on the neighbourhood of Θ_0 and impose the completeness and continuous paths condition of Liu & Shao (2003) (p. 815) on the set of these limits but we prefer to state the condition in the current, more familiar form.

Define the following neighbourhoods of the identified set:

$$\mathcal{B}_T = \left\{\theta \in \Theta : 0 < \chi(\theta, \theta_0) \leq \sqrt{M/T}\right\}$$
$$\mathcal{B}_{1T} = \left\{\theta \in \Theta : 0 < \chi(\theta, \theta_0) \leq \sqrt{M/T} \quad \text{and } \theta_1 = \tilde{\theta}_1\right\}$$

and introduce the finite sample parameter spaces:

$$\Lambda_T = \{(\theta_0, \lambda) : \theta_0 \in \Theta_0, \lambda \in R^d : \|\lambda\| = 1, \theta_0 + \lambda/\sqrt{T} \in \mathcal{B}_T\}$$
$$\Lambda_T^2 = \{(\theta_0, \lambda) : \theta_0 \in \Theta_0, \lambda = [\mathbf{0}_{d_1} \quad \lambda_2]', \lambda_2 \in R^{d_2}, \|\lambda\| = 1, \theta_0 + \lambda/\sqrt{T} \in \mathcal{B}_{1T}\}$$

Let $\to^H$ denote convergence in Hausdorff distance and $\Rightarrow$ denote weak convergence in the sense described by Van der Vaart & Wellner (1996).

Theorem 1. *If Assumption 4 holds and $\Lambda_T \to^H \Lambda, \Lambda_T^2 \to^H \Lambda^2$, then for $\tilde{\theta}_1 \in \Theta_{01}$:*

$$LR(\tilde{\theta}_1) \Rightarrow \sup_{(\theta_0,\lambda)\in\Lambda} v_D^2(\theta_0, \lambda) - \sup_{(\theta_0,\lambda)\in\Lambda^2} v_D^2(\theta_0, \lambda)$$

where $v_D(\cdot, \cdot)$ is a tight Gaussian process with covariance kernel:

$$\Omega_D((\theta, \lambda), (\theta', \lambda')) = E_0\left[D(Y, \theta, \lambda)D(Y, \theta', \lambda')\right]$$

Theorem 1 derives the asymptotic distribution of our likelihood ratio test statistic. Note that the limiting distribution depends on the sets of cluster points Λ, Λ^2, which are difficult to estimate (λ's cannot be estimated consistently). We follow Chen, Tamer & Torgovitsky (2011) and suggest using multiplier bootstrap for obtaining the critical value for our statistic. Let $\{w_t\}_{t=1}^T$ be a sequence of positive random numbers with $E(w_t) = 1, Var(w_t) = 1$ and $\int_0^\infty \sqrt{(Pr(|w_t| > s)}ds < \infty$, drawn independently of $\{Y_t\}_{t=1}^T$. Then, the critical value at the κ level can be obtained as a $1-\kappa$ quantile across the bootstrap samples of:

$$LR^*(\tilde{\theta}_1) = 2\left[\sup_{\theta\in\Theta}\sum_{t=1}^T w_t \log p(Y_t, \theta) - \sup_{\theta_2\in\Theta_2(\hat{\theta}_1)}\sum_{t=1}^T w_t \log p(Y_t, (\hat{\theta}_1, \theta_2))\right]$$

where $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2) \in \arg\max_{\theta\in\Theta} L_T(\theta)$.

7 Conclusion

Though we focus on repeated games with random states, our identification and inference approach can be generalized to general stochastic games with state dependence. Recent advancements in the analysis of these games (Abreu, Brooks & Sannikov (2016)) provide readily available procedures for computing equilibrium payoff sets, which can be naturally embedded in our econometric approach. We investigate this extension in a separate paper.

A Mathematical Proofs

A.1 Useful lemmas

Lemma 1 (Liu & Shao (2003))**.** *For some* $\epsilon > 0$ *let* $\mathcal{F}_\epsilon = \left\{ S(Z,\theta) = \left(\frac{p(Y,\theta)}{p_0(Y)} - 1\right)/\chi(\theta,\theta_0) : \theta \in \Theta_\epsilon/\Theta_0 \right\}$ *form a Donsker class with a square integrable envelope function. Then, uniformly over* $\theta \in \Theta_\epsilon/\Theta_0$*:*

$$\max\left\{2\sum_{t=1}^{T}\log\frac{p(Y_t,\theta)}{p_0(Y_t)}, 0\right\} = \max\left\{2\sqrt{T}\chi(\theta,\theta_0)\frac{1}{\sqrt{T}}\sum_{t=1}^{T}S(Y_t,\theta) - T\chi^2(\theta,\theta_0), 0\right\} + o_p(1)$$

A.2 Proof of Theorem 1

Assumption 4(a)(b) implies that log-likelihood achieves its maximum. Let

$$\hat{\theta} \in \arg\max_{\theta\in\Theta} L_T(\theta) \quad \text{and} \quad \tilde{\theta}_2 \in \arg\max_{\theta_2\in\Theta_2(\tilde{\theta}_1)} L_T(\tilde{\theta}_1,\theta_2)$$

and $\tilde{\theta} = (\tilde{\theta}_1, \tilde{\theta}_2)$.

Step 1: Consistency

First we are going to prove consistency of the maximum likelihood estimator $\hat{\theta}$ in the χ-distance (consistency of $\tilde{\theta}_2$ follows similarly). Define the Kullback-Leibler divergence

$$KL(\theta_0,\theta) = E_0[\log(p_0(Y)/p(Y,\theta))].$$

Using $\log(x) \leq x - 1 - 1/2(x-1)^2$ we obtain:

$$KL(\theta_0,\theta) \geq \frac{1}{2}\chi^2(\theta,\theta_0).$$

Now note that $\chi^2(\hat{\theta}, \theta_0) > \epsilon$ for arbitrarily small $\epsilon > 0$ implies:

$$E_0 \log \frac{p_0(Y)}{p(Y, \hat{\theta})} \geq \inf_{\theta: \chi^2(\theta, \theta_0) > \epsilon} E_0 \log \frac{p_0(Y)}{p(Y, \theta)} \geq \frac{\epsilon}{2}. \quad (2)$$

Consider the following derivation:

$$|E_0 \log p_0(Y) - E_0 \log p(Y, \hat{\theta})| \leq \left| E_0 \log p_0(Y) - \frac{1}{T} \sum_{t=1}^{T} \log p_0(Y_i) \right| + \left| E_0 \log p(Y, \hat{\theta}) - \frac{1}{T} \sum_{t=1}^{T} \log p(Y_i, \hat{\theta}) \right| < \frac{\epsilon}{2}$$

for T large enough, where the last inequality follows from Assumption 4(c) . Together with (2) this implies $\lim_{T \to \infty} P(\chi^2(\hat{\theta}, \theta_0) > \epsilon) = 0$.

Step 2: Rate of convergence

Step 1 implies that $\hat{\theta} \in \Theta_\epsilon$ for T large enough. If $\hat{\theta} \in \Theta_0$, then $L_T(\hat{\theta}) - L_T(\theta_0) = 0$. Let's focus on the case $\hat{\theta} \in \Theta / \Theta_0$. Note that Assumption 4(e) implies that $\mathcal{D}$ is a subset of $\mathcal{L}^2(P_0)$. This and Assumption 4(f) allow us to use Lemma 1 to obtain:

$$L_T(\hat{\theta}) - L_T(\theta_0) = \sqrt{T} \chi(\hat{\theta}, \theta_0) \frac{1}{\sqrt{T}} \sum_{t=1}^{T} S(Y_t, \hat{\theta}) - \frac{1}{2} T \chi^2(\hat{\theta}, \theta_0) \geq 0 \quad (3)$$

Now note that using Assumption 4(d) we have $S(Y, \theta) = D(Y, \bar{\theta}, \lambda)$ for $\bar{\theta}$ between $\hat{\theta}$ and θ_0. Donsker property of $\mathcal{D}$ implies $\frac{1}{\sqrt{T}} \sum_{t=1}^{T} D(Y_t, \bar{\theta}, \lambda) = O_p(1)$. This and (3) implies $\chi(\hat{\theta}, \theta_0) = O_p(T^{-1/2})$.

Step 3: Asymptotic distribution

For $\theta_0 = (\tilde{\theta}_1, \theta_2) \in \Theta_0$ decompose:

$$L_T(\hat{\theta}) - L_T(\tilde{\theta}) = L_T(\hat{\theta}) - L_T(\theta_0) - [L_T(\tilde{\theta}) - L_T(\theta_0)]$$

Note that $\hat{\theta} \in \mathcal{B}_T \cup \Theta_0$ with probability approaching one. Lemma 1 and Assumption A1(d) imply:

$$
\begin{aligned}
L_T(\hat{\theta}) - L_T(\theta_0) &= \sup_{\theta \in \mathcal{B}_T \cup \Theta_0} L_T(\theta) - L_T(\theta_0) = \max\left\{ \sup_{\theta \in \mathcal{B}_T} L_T(\theta) - L_T(\theta_0), 0 \right\} \\
&= \max\left\{ \sup_{\theta \in \mathcal{B}_T} \sqrt{T}\chi(\theta,\theta_0) \frac{1}{\sqrt{T}} \sum_{t=1}^{T} S(Y_t,\theta) - \frac{1}{2} T \chi^2(\theta,\theta_0), 0 \right\} + o_p(1) \\
&\leq \max\left\{ \sup_{\theta \in \mathcal{B}_T} \sup_{b_T} b_T \frac{1}{\sqrt{T}} \sum_{t=1}^{T} S(Y_t,\theta) - \frac{1}{2} b_T^2 + o_p(1), 0 \right\} \\
&= \sup_{\theta \in \mathcal{B}_T} \frac{1}{2} \left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} S(Y_t,\theta) \right)^2 + o_p(1) \\
&= \sup_{\theta_0 \in \Theta_0, \lambda \in R^d : \|\lambda\|=1, \theta_0 + \lambda/\sqrt{T} \in \mathcal{B}_T} \frac{1}{2} \left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} D(Y_t,\bar{\theta},\lambda) \right)^2 + o_p(1)
\end{aligned}
$$

where $\bar{\theta}$ is between θ_0 and $\theta_0 + \lambda/\sqrt{T}$. Note that we used the Taylor expansion $\frac{p(Y,\theta_0+\lambda/\sqrt{T})}{p_0(Y)} - 1 = G(Y,\bar{\theta})'\lambda/\sqrt{T}$.

Now the same argument as in Liu & Shao (2003) (p. 817) implies that for $\theta \in \Theta_\epsilon/\Theta_0$ we can set $\sqrt{T}\chi(\theta,\theta_0) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T} S(Y_t,\theta) + o_p(1)$ so we get:

$$
2[L_T(\hat{\theta}) - L_T(\theta_0)] \geq \sup_{\theta_0 \in \Theta_0, \lambda \in R^d : \|\lambda\|=1, \theta_0 + \lambda/\sqrt{T} \in \mathcal{B}_T} \left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} D(Y_t,\bar{\theta},\lambda) \right)^2 + o_p(1)
$$

where $\bar{\theta}$ is between θ_0 and $\theta_0 + \lambda/\sqrt{T}$. Thus, we have:

$$
2[L_T(\hat{\theta}) - L_T(\theta_0)] = \sup_{(\theta_0,\lambda) \in \Lambda_T} \left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} D(Y_t,\bar{\theta},\lambda) \right)^2 + o_p(1) \Rightarrow \sup_{(\theta_0,\lambda) \in \Lambda} v_D^2(\theta_0,\lambda) \tag{4}
$$

as $\Lambda_T \rightarrow^H \Lambda$ where the convergence follows from Assumption 4(f), which implies asymptotic equicontinuity of the empirical process $1/\sqrt{T}\sum_{t=1}^{T} D(Y_t,\theta,\lambda)$.

By similar arguments as above we obtain:

$$
2[L_T(\tilde{\theta}) - L_T(\theta_0)] = \sup_{(\theta_0,\lambda) \in \Lambda_T^2} \left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} D(Y_t,\bar{\theta},\lambda) \right)^2 + o_p(1) \Rightarrow \sup_{(\theta_0,\lambda) \in \Lambda^2} v_D^2(\theta_0,\lambda) \tag{5}
$$

where Λ^2 is the Hausdorff limit of Λ_T^2. Putting (4) and (5) together we obtain the final result.

References

Abito, J. M. & Szydlowski, A. (2017), How much can we identify from repeated games? Working Paper.

Abreu, D., Brooks, B. & Sannikov, Y. (2016), A "Pencil-Sharpening" Algorithm for Two Player Stochastic Games with Perfect Monitoring. Working Paper.

Abreu, D., Pearce, D. & Stacchetti, E. (1990), 'Toward a theory of discounted repeated games with imperfect monitoring', *Econometrica* **58**(5), pp. 1041–1063.

Abreu, D. & Sannikov, Y. (2014), 'An algorithm for two-player repeated games with perfect monitoring', *Theoretical Economics* **9**(2), 313–338.

Bajari, P., Benkard, C. L. & Levin, J. (2007), 'Estimating dynamic models of imperfect competition', *Econometrica* **75**(5), 1331–1370.

Bajari, P., Hong, H. & Ryan, S. P. (2010), 'Identification and estimation of a discrete game of complete information', *Econometrica* **78**(5), 1529–1568.

Chen, X., Tamer, E. & Torgovitsky, A. (2011), Sensitivity analysis in semiparametric likelihood models. Working paper.

Ciliberto, F. & Tamer, E. (2009), 'Market structure and multiple equilibria in airline markets', *Econometrica* **77**(6), 1791–1828.

Liu, X. & Shao, Y. (2003), 'Asymptotics for likelihood ratio tests under loss of identifiability', *The Annals of Statistics* **31**(3), 807–832.

Magnac, T. & Thesmar, D. (2002), 'Identifying dynamic discrete decision processes', *Econometrica* **70**(2), pp. 801–816.

Mailath, G. & Samuelson, L. (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

Tamer, E. (2003), 'Incomplete simultaneous discrete response model with multiple equilibria', *The Review of Economic Studies* **70**(1), 147–165.

Van der Vaart, A. W. & Wellner, J. A. (1996), *Weak Convergence and Empirical Processes: with Applications to Statistics*, Springer-Verlag, New York.