



Munich Personal RePEc Archive

How much can we identify from repeated games?

Abito, Jose Miguel

University of Pennsylvania

31 August 2015

Online at <https://mpra.ub.uni-muenchen.de/66378/>

MPRA Paper No. 66378, posted 02 Sep 2015 07:04 UTC

How much can we identify from repeated games?

Jose Miguel Abito*

August 31, 2015

[PRELIMINARY DRAFT. COMMENTS ARE WELCOME]

Abstract

I propose a strategy to identify structural parameters in infinitely repeated games without relying on equilibrium selection assumptions. Although Folk theorems tell us that almost any individually rational payoff can be an equilibrium payoff for sufficiently patient players, Folk theorems also provide tools to explicitly characterize this set of payoffs. I exploit the extreme points of this set to bound unobserved equilibrium continuation payoffs and then use these to generate informative bounds on structural parameters. I illustrate the identification strategy using (1) an infinitely repeated Prisoner's dilemma to get bounds on a utility parameter, and (2) an infinitely repeated quantity-setting game to get bounds on marginal cost and provide a robust test of firm conduct.

*University of Pennsylvania, Business Economics and Public Policy, abito@wharton.upenn.edu. Thanks...

1 Introduction

The application of game theoretical models and insights to empirical settings has been critical in capturing strategic interaction among economic agents in empirical work.¹ By combining assumptions on agents' rationality and an equilibrium concept, empirical researchers are able to link observed behavior with underlying model primitives such as agents preferences and constraints. Rationality allows one to put structure on how primitives translate to individual incentives for a given action. Because of strategic interaction, an equilibrium concept is used to link different agents' individual incentives and beliefs into outcomes in a consistent way.

Multiple equilibria typically precludes a unique mapping from primitives to observed behavior, making standard inference problematic without additional, often strong, assumptions. An important part of the literature on the econometrics of games have focused on the multiple equilibria problem. Substantial progress has been made in the analysis of static games (Tamer (2003), Ciliberto and Tamer (2009), Beresteanu, Molchanov, and Molinari (2011), Galichon and Henry (2011)). However, not much progress has been made for dynamic games. One important class of dynamic games where identification results are virtually unexplored are repeated games.

Repeated games provide a useful framework to model long-run relationships which lead to incentives and outcomes that are not otherwise captured by one-shot interactions. The framework provides insights on how agents can manage to cooperate in a non-cooperative environment without formal or explicit contracts, and how to build a reputation. Despite the richness of repeated games, there is some apprehension in its use due to the framework's general inability to make sharp predictions (Dal Bó and Frèchette, 2011). This has hampered the use of repeated games both in applied theory and empirical research. The multiplicity problem in repeated games is so perverse that a significant part of the theoretical literature has been devoted to deriving Folk Theorems that basically say that almost any individually rational payoff can be sustained as an equilibrium payoff. Thus, an open question in applying repeated games to empirical work is whether we can

¹De Paula (2013) provides a recent survey of the literature.

even learn anything from the data despite the Folk Theorem.

I develop an empirical strategy to identify structural parameters of infinitely repeated games without equilibrium selection assumptions. This strategy allows one to determine how much we can learn from the data just from the game's basic structure. Moreover, it provides a strategy that can be used in empirical research involving repeated games.

The identification problem and the consequent strategy is as follows. Actions taken by players today depend on what players expect to happen in the future. The one-stage deviation principle (OSDP) allows us to rationalize a given player's chosen action (conditional on the rival's) as an inequality comparing the sum of stage game payoffs and *equilibrium* continuation payoffs for the chosen action versus the best single-stage deviation. While OSDP simplifies the analysis by limiting to equilibrium continuation payoffs, analysis remains complicated since (almost) any individually rational payoff can be sustained as an equilibrium continuation payoff as long as agents are sufficiently patient. Hence, there are potentially infinitely many ways to rationalize observed actions and so one cannot directly invert the mapping from observed actions to primitives unless we know or assume what equilibrium is being played.

Folk Theorems explicitly characterize the set of equilibrium (average) continuation payoffs. I show how to use extreme points of this set to generate informative bounds on the structural parameter given data on frequency of stage game action profiles. While inference is limited to sets instead of singleton estimates of primitives (partial identification), the proposed identification strategy reveals that the basic structure of repeated games is quite informative about agents' primitives despite the Folk Theorem. Moreover, I show that set estimates of primitives can be used to provide a useful test of behavior (firm conduct).

In section 2, I start off with an analysis of a standard infinitely-repeated Prisoner's Dilemma to illustrate the identification strategy in a simple and familiar setting. I show how bounds on a utility parameter can be constructed using the incentive constraints from OSDP, and the upper and lower bounds of the set of equilibrium average payoffs from the Folk Theorem. Moreover, I study how the size of the identified set (i.e. how informative it is) depends on parameters of the

model.

The analysis of the Prisoner’s Dilemma yields general insights on identification. First, there are two levels of equilibrium selection assumptions—a static equilibrium selection assumption and a dynamic one. Dynamic equilibrium selection assumptions refer to assumptions on how players will play the game in the future which maps to what equilibrium continuation payoffs are. However, despite knowing what equilibrium continuation payoffs are, the model still turns out to be *incomplete*, in the sense similar to the multiple equilibrium problem in static games (Tamer, 2003). Finally, even if both static and dynamic equilibrium selection assumptions are made, point identification may still fail if there is no ample source of variation that the econometrician has some knowledge of (i.e. either this variation is observed or has some known distribution).

In section 3, I discuss the identification strategy in a more general setting. The same key ideas of using OSDP and the Folk Theorem drives the identification strategy. I then apply the identification strategy to a quantity-setting game in section 4. This application differs from the Prisoner’s Dilemma in an important way. In the Prisoner’s Dilemma, the econometrician explicitly observes stage game actions “Cooperate” or “Deviate” which contains some information on players’ long-run strategies. In the quantity-setting game, I assume that the econometrician only observes the actual values of the quantities chosen, and *not* whether these quantities map to “Cooperate” (e.g. joint monopoly quantity) or “Deviate” (e.g. Cournot quantity). I derive bounds for marginal cost and show how parameters of the model affect informativeness of these bounds. Finally I show how one can use the set estimates of marginal cost to test for firm conduct.

The identification strategy can be adapted to other dynamic games, specifically to stochastic games. Stochastic games with Markovian strategies are widely used in fields such as industrial organization and labor, although actual estimation of these games has been more limited. Moreover, feasibility of current methods require an equilibrium selection assumption. The advantage of my identification strategy is that it does not require equilibrium selection assumptions. Although the strategy requires knowing something about equilibrium continuation payoffs, there is no need to compute the whole set. All is needed is to compute the very extreme points of this set, which may

be easier to compute. I offer more discussion at the end in section 5.

2 Example: Prisoner's Dilemma

Figure 1 is a Prisoner's dilemma with unknown utility parameter $\alpha \in (0, 5)$. The goal of the econometrician is to estimate α given data on outcomes. Suppose this game is repeated infinitely many times with discount factor δ . Assume that the econometrician observes a sample of actions, i.e. stage game outcomes, and also the discount factor. For now let us assume that the econometrician observes a time series of stage game outcomes for a single market (or prison).

| | | | |
|----|----------|------------------|------------------|
| | | P2 | |
| | | <i>C</i> | <i>D</i> |
| P1 | <i>C</i> | α, α | $\alpha - 5, 10$ |
| | <i>D</i> | $10, \alpha - 5$ | $0, 0$ |

Figure 1: Prisoner's dilemma stage game payoffs

To estimate α , we want to link observed frequencies of stage outcomes with predicted frequencies coming from a theoretical model. Consider observing the frequency of (C, C) in the data, i.e. $\Pr(C, C)$. The one-stage deviation principle (OSDP) allows us to characterize a player's incentive to choose C when the rival chooses C as an inequality

$$(1 - \delta)\alpha + \delta v_{CC|h} \geq (1 - \delta)10 + \delta v_{CD|h}$$

where $v_{kk'|h}$'s are equilibrium average continuation payoffs proceeding immediately the outcome (k, k') and history h . Let $\Delta_{C|h} \equiv v_{CC|h} - v_{DC|h}$ and $\Delta_{D|h} \equiv v_{DD|h} - v_{DC|h}$. Using the OSDP, the subgame perfect Nash equilibrium of the infinitely repeated game are just Nash equilibria of the

normal form game in figure 2 for a given history h . This gives necessary and sufficient conditions on when we will observe the stage game outcome $(C, C$ and hence the frequency of this outcome occurring.

| | | | |
|----|-----|--|--|
| | | P2 | |
| | | C | D |
| P1 | C | $(1 - \delta)(\alpha - 10) + \delta\Delta_{C h}, (1 - \delta)(\alpha - 10) + \delta\Delta_{C h}$ | $(1 - \delta)(\alpha - 5) + \delta\Delta_{D h}, 0$ |
| | D | $0, (1 - \delta)(\alpha - 5) + \delta\Delta_{D h}$ | $0, 0$ |

Figure 2: Normal form game using one-stage deviation principle: Prisoner's Dilemma

A necessary condition for observing (C, C) given history h is that it is a Nash equilibrium of the normal form game. If we restrict to pure strategies, this necessary condition is given by²

$$(1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} > 0.$$

Thus,

$$\Pr(C, C|h) \leq \Pr((1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} > 0|h). \quad (1)$$

Similarly, a sufficient condition for observing (C, C) given history h is that it is the unique Nash equilibrium of the normal form game:

$$(1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} > 0 \cap (1 - \delta)(\alpha - 5) + \delta\Delta_{D|h} > 0$$

and so

$$\Pr((1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} > 0 \cap (1 - \delta)(\alpha - 5) + \delta\Delta_{D|h} > 0|h) \leq \Pr(C, C|h). \quad (2)$$

²If we allow for (uncorrelated) mixed strategies, then there is a mixed strategy equilibrium involving each player choosing C with probability $\rho = \frac{(1 - \delta)(\alpha - 5) + \delta\Delta_{D|h}}{(1 - \delta)5 + \delta(\Delta_{D|h} - \Delta_{C|h})}$. In order for $\rho \in (0, 1)$, we need $(1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} < 0$ and $(1 - \delta)(\alpha - 5) + \delta\Delta_{D|h} > 0$. Thus, $\Pr(C, C|h) \leq \Pr((1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} > 0|h) + \rho \cdot \Pr[(1 - \delta)(\alpha - 10) + \delta\Delta_{C|h} < 0 \cap (1 - \delta)(\alpha - 5) + \delta\Delta_{D|h} > 0|h]$.

The above analysis shows that despite knowing the equilibrium continuation payoffs $(\Delta_{C|h}, \Delta_{D|h})$, the model is still *incomplete* in that we cannot write the likelihood for (C, C) (given h) (Tamer, 2003). In this sense, we have a static equilibrium selection problem on top of the dynamic equilibrium selection problem, i.e. knowing the equilibrium continuation payoffs. The identification strategy deals with both the static and dynamic equilibrium selection problems.

2.1 Deriving the identified set

Let F_{Δ_h} as the joint distribution of $(\Delta_{C|h}, \Delta_{D|h})$. This distribution captures the dynamic equilibrium selection mechanism. Dynamic equilibrium selection assumptions are essentially assumptions on F_{Δ_h} . For example, if we assume that players are implementing a Grim-Trigger strategy that sustains (C, C) forever using a punishment of (D, D) , then F_{Δ_h} puts all its mass on $\Delta_{C|h} = \alpha - 0 = \alpha$ and $\Delta_{D|h} = 0 - 0 = 0$ for histories involving cooperation in the past, and $\Delta_{C|h} = \Delta_{D|h} = 0$ for histories where some deviation in the past occurred.




Figure 3: Set of individually rational payoffs in in our Prisoner's Dilemma ($\alpha \in (0, 5)$)

The Folk Theorem in fact characterizes the set of equilibrium average payoffs for sufficiently patient players. Specifically, almost any individually rational stage payoff can be sustained as a subgame perfect Nash equilibrium for a high enough discount factor. Figure 3 gives the set \mathcal{F}

of individually rational payoffs for our Prisoner's Dilemma. Let $\underline{v}(\alpha)$ and $\bar{v}(\alpha)$ be the lower and upper bounds of this set, which are functions of the unknown parameter α . For any discount factor less than or equal to one, the support of F_{Δ_h} for any history h must be a subset³ of the set of differences in equilibrium payoffs in \mathcal{F} , i.e. $\text{supp}(F_{\Delta_h}) \subseteq [\underline{\Delta}(\alpha), \bar{\Delta}(\alpha)]$, where $\underline{\Delta}(\alpha) = \underline{v}(\alpha) - \bar{v}(\alpha)$ and $\bar{\Delta}(\alpha) = \bar{v}(\alpha) - \underline{v}(\alpha)$. The following proposition shows how one can use these bounds to derive an identified set for the unknown parameter α without any equilibrium selection assumptions.

Proposition 1. Assume players are restricted to pure strategies. Moreover, suppose we only observe the frequency $\Pr(C, C)$ and the discount factor δ . The identified set for α is given by

$$\mathcal{H}(\alpha) = \{\alpha : \Pr(C, C) \in [\Pr(LB(\alpha) > 0), \Pr(UB(\alpha) > 0)]\}$$

where $LB(\alpha) = (1 - \delta)(\alpha - 10) + \delta\underline{\Delta}(\alpha)$, $UB(\alpha) = (1 - \delta)(\alpha - 10) + \delta\bar{\Delta}(\alpha)$, $\bar{\Delta}(\alpha) = \frac{50 - (\alpha - 5)(\alpha - 10)}{5}$ and $\underline{\Delta}(\alpha) = -\bar{\Delta}(\alpha)$.

Proof. Since $\Delta_{C|h} \leq \bar{\Delta}(\alpha)$ for all histories h , the necessary condition given by inequality 1 implies

$$\Pr(C, C|h) \leq \Pr((1 - \delta)(\alpha - 10) + \delta\bar{\Delta}(\alpha) > 0).$$

Integrating this inequality across histories yields

$$\Pr(C, C) \leq UB(\alpha)$$

where $UB(\alpha) = (1 - \delta)(\alpha - 10) + \delta\bar{\Delta}(\alpha)$.

Next, since $\Delta_{C|h} \geq \underline{\Delta}(\alpha)$ and $\Delta_{D|h} \geq \underline{\Delta}(\alpha)$ for all histories h , the sufficient condition given by inequality 2 implies

$$\Pr((1 - \delta)(\alpha - 10) + \delta\underline{\Delta}(\alpha) > 0 \cap (1 - \delta)(\alpha - 5) + \delta\underline{\Delta}(\alpha) > 0) \leq \Pr(C, C|h).$$

Integrating this across histories gives

$$\Pr((1 - \delta)(\alpha - 10) + \delta\underline{\Delta}(\alpha) > 0 \cap (1 - \delta)(\alpha - 5) + \delta\underline{\Delta}(\alpha) > 0) \leq \Pr(C, C).$$

³Moreover, since we restrict to pure strategies, the actual set of equilibrium average payoffs is actually a proper subset of \mathcal{F} even for arbitrarily large discount factors.

Since $(1 - \delta)(\alpha - 10) + \delta\underline{\Delta}(\alpha) > 0$ implies $(1 - \delta)(\alpha - 5) + \delta\underline{\Delta}(\alpha) > 0$, the lefthand side of the above inequality is just equal to

$$\Pr((1 - \delta)(\alpha - 10) + \delta\underline{\Delta}(\alpha) > 0).$$

Thus,

$$LB(\alpha) \leq \Pr(C, C)$$

where $LB(\alpha) = (1 - \delta)(\alpha - 10) + \delta\underline{\Delta}(\alpha)$.

To derive $\overline{\Delta}(\alpha)$ and $\underline{\Delta}(\alpha)$, recall that $\overline{\Delta}(\alpha) = \overline{v}(\alpha) - \underline{v}(\alpha)$ and $\underline{\Delta}(\alpha) = \overline{\Delta}(\alpha)$. From figure 3, we have $\underline{v}(\alpha) = 0$ while $\overline{v}(\alpha)$ is characterized by the point $(0, \overline{v}(\alpha))$ that lies on the line segment connecting $(\alpha - 5, 10)$ and (α, α) . (or equivalently, the point $(\overline{v}(\alpha), 0)$ that lies on the line segment connecting (α, α) and $(10, \alpha - 5)$). After some algebra, one can show that $\overline{v}(\alpha) = \frac{50 - (\alpha - 5)(\alpha - 10)}{5}$. \square

To explicitly derive the identified set for α , we need to derive all values of α such that the observed frequency $\Pr(C, C)$ is inside the interval $[\Pr(LB(\alpha), \Pr(UB(\alpha))]$. Since both $LB(\alpha)$ and $UB(\alpha)$ are deterministic functions of α , $\Pr(LB(\alpha)) = \mathbb{I}\{LB(\alpha)\}$ and $\Pr(UB(\alpha)) = \mathbb{I}\{UB(\alpha)\}$ where $\mathbb{I}\{\cdot\}$ is the indicator function. Moreover, since $LB(\alpha)$ and $UB(\alpha)$ are continuous functions of α , the set of α 's that satisfy $\Pr(C, C) \in [\mathbb{I}\{LB(\alpha)\}, \mathbb{I}\{UB(\alpha)\}]$ are just a finite union of intervals over the set of possible α 's.

There are three possible identified sets depending on $\Pr(C, C)$. Figure 4 illustrates these identified sets graphically. If $\Pr(C, C) = 0$, then it must be that $LB(\alpha) \leq 0$ which gives some bound on α , i.e. $\alpha \in [\alpha_1, \alpha_3]$ in figure 4. Notice that there is no restriction provided by $UB(\alpha)$ since $\mathbb{I}\{UB(\alpha)\}$ can be zero or one. Intuitively, $\Pr(C, C) = 0$ requires that (C, C) is *not* a unique Nash equilibrium of the induced normal form game in figure 2. However, $\Pr(C, C) = 0$ does not rule out that (C, C) can be one of multiple Nash equilibria.

Similarly, if $\Pr(C, C) = 1$, then it must be that $UB(\alpha) > 0$ but there is no restriction provided by $LB(\alpha)$. Finally, when $\Pr(C, C) \in (0, 1)$, it must be that $UB(\alpha) > 0$ and $LB(\alpha) \leq 0$ since (C, C) must be a Nash equilibrium, but is not unique. The identified set when $\Pr(C, C) \in (0, 1)$ is tighter

than the other two cases since it provides more restrictions on α . However, different values for $\Pr(C, C) \in (0, 1)$ yields the same identified set so what only matters is that the frequency is neither zero or one, and no additional information is provided. The reason is the following. Variation in outcomes, say of (C, C) , are driven by the distribution F_{Δ_h} . The shape of this distribution provides useful information about how outcomes would vary in the sample. However, at this point, we only know the support of this distribution, specifically, its endpoints, which is not enough to explain what determines the “non-degenerate” variation leading to specific values for $\Pr(C, C) \in (0, 1)$. Thus the researcher needs additional known sources of variation, either through an observed variable or an unobserved variable with some known distribution that moves around payoffs and outcomes.




Figure 4: Identified set

Note: Blue and red curves are graphs of $UB(\alpha)$ and $LB(\alpha)$ as defined in Propostion 1, respectively. The dotted blue and red curves give $\Pr(UB(\alpha) > 0)$ and $\Pr(LB(\alpha) > 0)$, respectively. Consider frequency of observing (C, C) . If $\Pr(C, C) = 0$, then $\mathcal{H}(\alpha) = [\alpha_1, \alpha_3]$. If $\Pr(C, C) = 1$, then $\mathcal{H}(\alpha) = (\alpha_2, \alpha_4)$. Finally, if $\Pr(C, C) \in (0, 1)$, then $\mathcal{H}(\alpha) = (\alpha_2, \alpha_3]$.

2.2 Point identification

As it turns out, knowing how equilibrium is selected (both dynamic and static equilibrium selection) is not sufficient to get point identification of α . To show this, suppose players are playing Grim-Trigger strategies and that δ is fixed at some value greater than 0.6, say $\bar{\delta} = 0.7$ (known to the researcher). Then the researcher observes $\Pr(C, C) = 1$ and that $(1 - \delta)(\alpha - 10) + \delta\alpha > 0$ and $(1 - \delta)(\alpha - 5) < 0$. These imply $\alpha > 10(1 - \bar{\delta}) = 3$ and $\alpha < 5$, which only yields $\alpha \in (3, 5)$.

The above example illustrates that knowledge of equilibrium selection is not sufficient to pin down the value of α . Additional sources of variation that moves payoffs (and outcomes) are needed to get point identification. For example, suppose we observe different independent prisons indexed by m , where all players are playing Grim-Trigger strategies. Suppose the discount factor of prisoners differ across prisons, say δ_m , and that we have enough variation in δ_m . Specifically, we observe that for prisons with $\delta_m > 0.6$, $\Pr(C, C) = 1$ while for $\delta_m < 0.6$, $\Pr(C, C) = 0$. Then we have $(1 - \delta)(\alpha - 10) + \delta\alpha = 0$ for $\delta = 0.6$ which implies $\alpha = 4$.

3 Identification strategy

The identification strategy consists of two key ideas. First, we can use the one-stage deviation principle (OSDP) to derive a set of incentive (inequality) constraints that describe conditions such that a given outcome is part of an equilibrium. The incentive constraints are functions of stage game and equilibrium continuation payoffs, and if both are known up to the unknown parameter, then one can use these constraints to identify the parameter (although we might still only achieve partial identification due to the static equilibrium selection problem). However, the econometrician does not typically know what equilibrium strategies are being played hence what equilibrium continuation payoffs go into these incentive constraints. In fact, there are potentially infinitely many equilibrium payoffs that are consistent with observed outcomes—a result from Folk Theorems. The second idea is that we can harness the power of Folk Theorems to provide lower and upper bounds on these equilibrium continuation payoffs. As long as the incentive constraints

are monotonic in these payoffs, we can then use the conditions implied by these worse case bounds to derive an identified set for the unknown parameter.

3.1 Set-up

There are two players choosing actions simultaneously in each period indefinitely with common discount factor δ . Stage game actions for each player i are elements of the compact set A_i . At this point, it is not necessary to specify whether actions are discrete or continuous. Moreover, actions need not have specific meaning as in the Prisoner's Dilemma, and can just be any real number.

Player i 's payoff is denoted by $\pi_i(x_i, x_j; \alpha)$ where x_i is the chosen action of player i , x_j is the chosen action of i 's rival j , and α is some parameter. I assume that the econometrician knows the functional form of $\pi_i(x_i, x_j; \alpha)$ but does not know α . Additionally, the econometrician observes the chosen stage game actions $\{\hat{x}_i, \hat{x}_j\}$, and also the frequency that these occur, i.e. $\Pr(\hat{x}_i, \hat{x}_j)$. Finally, let $x_i^*(x_j)$ be the *stage game* best response of i to j 's action x_j , i.e. $x_i^*(x_j) = \arg \max_{x \in A_i} \pi_i(x, x_j; \alpha)$. This best response can be a function of α , as long as it can be explicitly written.

3.2 One-stage deviation principle (OSDP)

OSDP allows us to focus on single-stage deviations to fully characterize incentives of players for choosing specific actions. That is, player i chooses \hat{x}_i given j plays \hat{x}_j if and only if

$$(1 - \delta)\pi_i(\hat{x}_i, \hat{x}_j; \alpha) + \delta\hat{v}_i \geq (1 - \delta)\pi_i(x_i^*(\hat{x}_j), \hat{x}_j; \alpha) + \delta\tilde{v}_i$$

for some $\hat{v}_i, \tilde{v}_i \in \mathcal{V}$ where \mathcal{V} is the set of *equilibrium* average payoffs. When constructing a deviation (righthandside of the above inequality), we allow i to deviate only for one period by choosing its best response $x_i^*(\hat{x}_j)$ but require i to go back to equilibrium play afterwards. Equilibrium play then leads to player i receiving equilibrium continuation payoffs (written as an average).

The subgame perfect Nash equilibrium (SPNE) of the infinitely repeated game is a Nash equilibrium of every possible history. For each (relevant) history, we can define a normal form game based on the incentive constraints derived using OSDP. The SPNE of the repeated game is the set

of average discounted payoffs such that observed stage game outcomes are Nash Equilibria of these normal form games.

The next step is to translate the incentive constraints as bounds on the probability of observing a given stage outcome. Take for example the stage game outcome (\hat{x}_i, \hat{x}_j) chosen under history h . For this stage game outcome to be a Nash equilibrium in pure strategies of the induced normal form game, it must be that

$$(1 - \delta)\pi_i(\hat{x}_i, \hat{x}_j; \alpha) + \delta\hat{v}_{i|h} \geq (1 - \delta)\pi_i(x_i^*(\hat{x}_j), \hat{x}_j; \alpha) + \delta\tilde{v}_{i|h}$$

and

$$(1 - \delta)\pi_j(\hat{x}_j, \hat{x}_i; \alpha) + \delta\hat{v}_{j|h} \geq (1 - \delta)\pi_j(x_j^*(\hat{x}_i), \hat{x}_i; \alpha) + \delta\tilde{v}_{j|h}$$

With mixed strategies, the set of conditions will be more complicated though are still functions of equilibrium average payoffs. We can define these *necessary* conditions as a set $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ such that $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ is increasing⁴ in Δ_h where Δ_h is an appropriately defined difference of equilibrium average payoffs. Thus for each h , $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha) \subseteq \Psi(\hat{x}_i, \hat{x}_j, \bar{\Delta}); \alpha)$, where $\bar{\Delta}$ is an upperbound that does not depend on h .

Next, we can derive a set of conditions such that (\hat{x}_i, \hat{x}_j) is the *unique* Nash equilibrium of the induced normal form game.⁵ Denote the set of *sufficient* conditions as $\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ for each history h , where again $\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ is increasing in Δ_h . Thus for each h , $\Xi(\hat{x}_i, \hat{x}_j, \underline{\Delta}; \alpha) \subseteq \Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$, where $\underline{\Delta}$ is a lowerbound that does not depend on h .

3.3 Main theorem

The following theorem characterizes the identified set for the unknown parameter α .

Theorem 1. Construct $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ and $\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ such that both are increasing in Δ_h .

⁴That is, $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha) \subseteq \Psi(\hat{x}_i, \hat{x}_j, \Delta'_h; \alpha)$ for $\Delta'_h > \Delta_h$.

⁵It is possible for this set to be empty.

The identified set for α is given by

$$\mathcal{H}(\alpha) = \bigcap_{(\hat{x}_i, \hat{x}_j)} \{ \alpha : \Pr(\hat{x}_i, \hat{x}_j) \in [\Pr(\Xi(\hat{x}_i, \hat{x}_j, \underline{\Delta}); \alpha), \Pr(\Psi(\hat{x}_i, \hat{x}_j, \overline{\Delta}); \alpha)] \}$$

Proof. Consider some arbitrary outcome $(\hat{x}_i, \hat{x}_j,)$. Since $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ is necessary for one to observe $(\hat{x}_i, \hat{x}_j,)$ under some history h in an equilibrium, we have

$$\Pr(\hat{x}_i, \hat{x}_j | h) \leq \Pr(\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha) | h).$$

Because $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha) \subseteq \Psi(\hat{x}_i, \hat{x}_j, \overline{\Delta}; \alpha)$, we have

$$\Pr(\hat{x}_i, \hat{x}_j | h) \leq \Psi(\hat{x}_i, \hat{x}_j, \overline{\Delta}; \alpha)$$

for all histories. Thus,

$$\Pr(\hat{x}_i, \hat{x}_j) \leq \Psi(\hat{x}_i, \hat{x}_j, \overline{\Delta}; \alpha).$$

Next, a sufficient condition for observing $(\hat{x}_i, \hat{x}_j,)$ under some history h in an equilibrium is captured by the set $\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$. Hence,

$$\Pr(\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha) | h) \leq \Pr(\hat{x}_i, \hat{x}_j | h).$$

Since $\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ is increasing in Δ_h , and after integrating across histories, we have

$$\Pr(\Xi(\hat{x}_i, \hat{x}_j, \underline{\Delta}); \alpha) \leq \Pr(\hat{x}_i, \hat{x}_j).$$

Thus for each outcome (\hat{x}_i, \hat{x}_j) we can derive an identified set as the set of α 's that satisfy $\Pr(\hat{x}_i, \hat{x}_j) \in [\Pr(\Xi(\hat{x}_i, \hat{x}_j, \underline{\Delta}); \alpha), \Pr(\Psi(\hat{x}_i, \hat{x}_j, \overline{\Delta}); \alpha)]$. The identified set that uses all data, $\mathcal{H}(\alpha)$ is just the intersection of these sets. \square

An important part of the identification strategy is to construct the set of necessary and sufficient conditions $\Psi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ and $\Xi(\hat{x}_i, \hat{x}_j, \Delta_h; \alpha)$ such that both are increasing in Δ_h . For each data point $(\hat{x}_i, \hat{x}_j,)$, we can construct bounds on the probability of observing this data point using the necessary and sufficient conditions evaluated at the extreme points $\underline{\Delta}$ and $\overline{\Delta}$. These lower and upper envelopes can then be used to identify upper and lower bounds for the unknown parameter. See figure 5 for a graphical intuition.




Figure 5: Graphical illustration of how to derive the identified set

4 Application: Quantity-setting game

I apply the identification strategy to a simple quantity-setting game. There are two firms, firms 1 and 2, competing by choosing how much to produce in each period. The firms compete for an indefinite amount of periods with a common discount factor δ . For ease of exposition, I restrict quantities to two levels: low quantity q_L or high quantity q_H , where $q_L < q_H$.⁶ I assume firms are symmetric and have a constant marginal cost of c . Finally, inverse demand is given by $P = 1 - Q$, where $Q = q_1 + q_2$.

I assume that in each period, profits are perturbed by choice-specific shocks. This approach is similar to previous work that models oligopoly competition as repeated games (Green and Porter, 1984; Rotemberg and Saloner 1986; Fershtman and Pakes, 2000). In period t , total stage game profit of a given firm when it chooses i and the rival chooses j is the sum of its *base* profit π_{ij} and a choice-specific shock η_{ijt} : $\pi_{ij} + \eta_{ijt}$, for $i, j \in \{L, H\}$. I assume shocks are iid across periods (and

⁶One can easily allow for more *discrete* choices. However, since I will include choice-specific shocks in the model, assuming actions are continuous introduces some technical (measurability) issues.

| | | | |
|--------|-----|--|--|
| | | Firm 2 | |
| | | L | H |
| Firm 1 | L | $\pi_{LL} + \eta_{LLt}, \pi_{LL} + \eta_{LLt}$ | $\pi_{LH} + \eta_{LHt}, \pi_{HL} + \eta_{HLt}$ |
| | H | $\pi_{HL} + \eta_{HLt}, \pi_{LH} + \eta_{LHt}$ | $\pi_{HH} + \eta_{HHt}, \pi_{HH} + \eta_{HHt}$ |

Figure 6: Period t stage game payoffs for quantity-setting game

markets) have mean zero, are common to both players and are symmetric, i.e. $\eta_{ij} = \eta_{ji}$. Choice shocks η_{ijt} are revealed to both players only at the beginning of each period, and are common knowledge. Firms then decide on their quantities simulatenously. Figure 6 gives the payoff matrix for the stage game.

Just as in the Prisoner's Dilemma example in section 2, we can use the one-stage deviation principle to characterize the equilibrium of the infinitely repeated game as a Nash equilibrium in the normal form game illustrated in figure 7 for every relevant history h and time t . Here, $\epsilon_{Lt} = \eta_{LLt} - \eta_{HLt}$, $\epsilon_{Ht} = \eta_{HLt} - \eta_{HHt}$, $\Delta_{L|ht} = v_{LL|ht} - v_{HL|ht}$ and $\Delta_{H|ht} = v_{HL|ht} - v_{HH|ht}$ where $v_{ij|ht}$ are equilibrium expected (with respect to the distribution of η 's) average continuation payoffs.

| | | | |
|--------|-----|--|--|
| | | Firm 2 | |
| | | L | H |
| Firm 1 | L | $(1 - \delta)(\pi_{LL} - \pi_{HL}) + \delta\Delta_{L ht}, (1 - \delta)(\pi_{LL} - \pi_{HL}) + \delta\Delta_{L ht}$ | $(1 - \delta)(\pi_{LH} - \pi_{HH}) + \delta\Delta_{H ht}, \epsilon_{Lt}$ |
| | H | $\epsilon_{Lt}, (1 - \delta)(\pi_{LH} - \pi_{HH}) + \delta\Delta_{H ht}$ | $\epsilon_{Ht}, \epsilon_{Ht}$ |

Figure 7: Normal form game using one-stage deviation principle: Quantity-setting game




Figure 8: Nash equilibria of induced normal form

Since each of these normal form games depend on history h and time t , the Nash equilibria of each of these games also depend on the realization of $(\epsilon_{Lt}, \epsilon_{Ht})$.⁷ Assume firms are only playing pure strategies. Figure 8 summarizes the Nash equilibria of the normal form game in the $(\epsilon_{Lt}, \epsilon_{Ht})$ space. For example, if $\epsilon_{Lt} < (1 - \delta)(\pi_{LL} - \pi_{HL}) + \delta\Delta_{L|ht}$ and $\epsilon_{Ht} < (1 - \delta)(\pi_{LH} - \pi_{HH}) + \delta\Delta_{H|ht}$, then the stage outcome (q_L, q_L) at time t , after history h and upon observing these types of ϵ 's can be sustained by some equilibrium strategy.

4.1 Data generation

Let the true marginal cost be equal to $c_0 = 0.5$. I assume the the low and high quantities respectively correspond to monopoly and Cournot quantities given by

$$\begin{aligned} q_L = q^M &= \frac{1 - c_0}{4} = \frac{1}{8} \\ q_H = q^C &= \frac{1 - c_0}{3} = \frac{1}{6}. \end{aligned}$$

⁷Strategies themselves are functions of ϵ 's.

While firms are aware that $q_L = q^M$ and $q_H = q^C$, the econometrician does *not* have this information. All the econometrician observes are the actual numerical values of q_L and q_H .

Instead of defining strategies explicitly, I simulate data directly by assuming $\Delta_L = \pi^M - \pi^C$ and $\Delta_H = \pi^C - \pi^C = 0$ along the equilibrium path. Note that in a model with no shocks, i.e. $\eta_{ij} = 0$, these Δ 's are consistent with a Grim-Trigger strategy which would also imply $\Pr(q_L, q_L) = 1$. This need not be the case when there are shocks since outcomes such as (q_H, q_H) can occur along the equilibrium path for certain values of ϵ 's and given properly defined strategies that are conditional on ϵ 's.

I assume ϵ 's are independently distributed according to a logistic distribution with common variance σ^2 . Since ϵ 's have full support, multiple equilibria in the normal form game given in figure 7 can arise depending on the realization of ϵ 's (see figure 8). This corresponds to the equilibrium selection problem in static games (Tamer, 2003). I let s_1 and s_2 be static selection parameters where s_1 is equal to the probability that (q_L, q_L) is chosen when $\epsilon_{Lt} < (1 - \delta)(\pi_{LL} - \pi_{HL}) + \delta\Delta_{L|ht}$ and $\epsilon_{Ht} > (1 - \delta)(\pi_{LH} - \pi_{HH}) + \delta\Delta_{H|ht}$, while s_2 is equal to the probability that (q_L, q_H) is chosen when $\epsilon_{Lt} > (1 - \delta)(\pi_{LL} - \pi_{HL}) + \delta\Delta_{L|ht}$ and $\epsilon_{Ht} < (1 - \delta)(\pi_{LH} - \pi_{HH}) + \delta\Delta_{H|ht}$.

Define *individual choice probabilities* as

$$\begin{aligned} p(\Delta; c) &\equiv \Pr(\epsilon_L < (1 - \delta)(\pi_{LL} - \pi_{HL}) + \delta\Delta | \Delta) \\ q(\Delta; c) &\equiv \Pr(\epsilon_H < (1 - \delta)(\pi_{LH} - \pi_{HH}) + \delta\Delta | \Delta). \end{aligned}$$

The logistic distribution assumption implies *true* individual choice probabilities being equal to

$$\begin{aligned} p_0 &\equiv \frac{\exp\left[\frac{(1-\delta)(\pi_{LL}-\pi_{HL})+\delta\Delta_L}{\sigma}\right]}{1 + \exp\left[\frac{(1-\delta)(\pi_{LL}-\pi_{HL})+\delta\Delta_L}{\sigma}\right]} \\ q_0 &\equiv \frac{\exp\left[\frac{(1-\delta)(\pi_{LH}-\pi_{HH})+\delta\Delta_H}{\sigma}\right]}{1 + \exp\left[\frac{(1-\delta)(\pi_{LH}-\pi_{HH})+\delta\Delta_H}{\sigma}\right]} \end{aligned}$$

Finally observed frequencies of stage game outcomes are computed as:

$$\begin{aligned}
\Pr(q_L, q_L) &= p_0 q_0 + s_1 \cdot p_0 (1 - q_0) \\
\Pr(q_H, q_H) &= (1 - p_0)(1 - q_0) + (1 - s_1) \cdot p_0 (1 - q_0) \\
\Pr(q_L, q_H) &= s_2 \cdot (1 - p_0) q_0 \\
\Pr(q_H, q_L) &= 1 - \Pr(q_L, q_L) - \Pr(q_H, q_H) - \Pr(q_L, q_H).
\end{aligned}$$

I assume the econometrician observes (or knows or can estimate beforehand) the discount factor δ , the inverse demand function $P(Q)$, the distribution of the shocks including its variance σ^2 , the frequencies $\Pr(q_L, q_L)$, $\Pr(q_H, q_H)$, $\Pr(q_L, q_H)$ and $\Pr(q_H, q_L)$ and the actual values of q_L and q_H . The econometrician does not know the true marginal cost c_0 , the static selection parameters (s_1, s_2) , the equilibrium continuation payoffs that are implemented, and finally that q_L and q_H correspond to the joint monopoly and Cournot quantities, respectively.

4.2 Identified set

The goal of the econometric exercise is to compute bounds for marginal cost and test for firm conduct, i.e. $H_0 : q_L = q^M$ and $H_0 : q_L = q^C$. If both static and dynamic selection parameters were known (i.e. $s_1, s_2, \Delta_L, \Delta_H$), then one can compute theoretical stage game frequencies as functions of marginal cost, then find the value of marginal cost that best matches these with the observed stage frequencies.⁸ Without this information, one can instead use the method in section 3 to get an identified set for marginal cost.

The identification strategy requires knowing the extreme points of the set of equilibrium continuation payoffs. Instead of relying on a Folk Theorem for this modified repeated game, I impose the restriction⁹ $v_{kk'} \in [\pi^C(c), \pi^M(c)]$ where $\pi^M(c)$ and $\pi^C(c)$ are individual firm's (joint) Monopoly

⁸The ‘‘ample source of variation’’ in this case is the known distribution of shocks.

⁹Note that it is possible to have worse equilibrium punishments than the Cournot payoff using more complicated strategies (Abreu, 1986; Abreu, 1988). I will derive the set of individually rational payoffs and use in a future version of this paper. Nevertheless, restricting to Cournot payoff as being the worse punishment is still a useful benchmark

and Cournot profits:

$$\pi^M(c) = \frac{(1-c)^2}{8}, \pi^C(c) = \frac{(1-c)^2}{9}.$$

These imply $\bar{\Delta}(c) = \pi^M(c) - \pi^C(c) = (1-c)^2/72$ and $\underline{\Delta}(c) = -\bar{\Delta}(c)$. Finally, the identified set is the set of all values of c that simultaneously satisfy the following:

$$\begin{aligned} \Pr(q_L, q_L) &\in [p(\underline{\Delta}(c))q(\underline{\Delta}(c)), p(\bar{\Delta}(c))] \\ \Pr(q_H, q_H) &\in [(1-p(\bar{\Delta}(c)))(1-q(\bar{\Delta}(c))), 1-q(\underline{\Delta}(c))] \\ \Pr(q_L, q_H) &\in [0, (1-p(\underline{\Delta}(c)))q(\bar{\Delta}(c))] \\ \Pr(q_H, q_L) &\in [0, (1-p(\underline{\Delta}(c)))q(\bar{\Delta}(c))] \end{aligned}$$

4.3 Results

4.3.1 Bounds for marginal cost

I compute the bounds for different values of the standard deviation of the logit profit shocks σ ($=1$ and 0.001), discount factor δ ($=1$ and 0.001), and static selection parameters s_1 and s_2 ($s_1 = s_2 = 0, 0.5, 1$). Table 1 contains the results. Since the inverse demand function is $P = 1 - Q$, $[0, 1]$ is a natural bound on marginal cost. Thus, the computed bounds are only informative if it is a proper subset of $[0, 1]$.

The simulation parameter σ measures the standard deviation of the profit shocks. We can interpret the magnitude of σ by comparing it to monopoly profits under the true $c_0 = 0.5$, i.e. $\pi^M = 0.03125$. Clearly, a standard deviation equal to one is large relative to profits, hence representing an extreme case where decision-making will mostly be driven by these shocks. On the other hand, a standard deviation of 0.001 represents about a 3% movement in monopoly profits. The estimated bounds for marginal cost are much tighter when σ is small. This reflects the fact that the incentive compatibility constraints provide more information regarding outcomes when noise from exogenous profits shocks are smaller.

and does not trivialize the problem of multiple equilibria.

Estimated bounds are also tighter when the discount factor δ is low. In this case, the incentive compatibility constraints have more bite since collusion is harder to sustain with more impatient firms. The value $\delta = 0.5$ is the lowest discount factor that will make Grim-Trigger an SPNE in the game without profit shocks.

The static selection parameter s gives the probability that (q_L, q_L) and (q_L, q_H) are selected in the induced normal form game (see figure 8). The bounds are tighter when $s = 0.5$ and σ smaller. Interestingly however, the bounds become uninformative with $s = 0.5$ when σ is large.

Table 1: Bounds for marginal cost ($c_0 = 0.5$)

| | $\sigma = 1$ | $\sigma = 0.001$ |
|----------------|--------------|------------------|
| $\delta = 0.5$ | | |
| $s = 0$ | [0, 0.56] | [0.36, 0.56] |
| $s = 0.5$ | [0, 1] | [0.44, 0.59] |
| $s = 1$ | [0.49, 1] | [0.43, 0.61] |
| $\delta = 0.6$ | | |
| $s = 0$ | [0, 0.58] | [0.12, 0.58] |
| $s = 0.5$ | [0, 1] | [0.37, 0.62] |
| $s = 1$ | [0.35, 1] | [0.36, 0.67] |

4.3.2 Test of firm conduct

While the bounds on marginal cost are of interest themselves, they can also be used to provide a test for firm conduct. Note that if we had a point estimate for marginal cost, we can compute estimates of the monopoly and Cournot quantities \hat{q}^M and \hat{q}^C , respectively, predicted by theory. However, since we only have bounds for c and not point estimates, we can only get bounds for these predicted quantities. Nevertheless, these quantity bounds are still useful in determining whether observed quantities are consistent with collusion or Cournot competition.

Table 2: Bounds for theoretical quantities

| $\sigma = 0.001, \delta = 0.5$ | | | |
|--------------------------------|--------------|----------------|----------------|
| | c | $q^M = 1/8$ | $q^C = 1/6$ |
| $s = 0$ | [0.36, 0.56] | [0.11, 0.16] | [0.147, 0.213] |
| $s = 0.5$ | [0.44, 0.59] | [0.103, 0.140] | [0.137, 0.187] |
| $s = 1$ | [0.43, 0.61] | [0.098, 0.143] | [0.13, 0.19] |

Table 2 provides the bounds for marginal cost, and predicted monopoly and Cournot quantities (based on theory). Consider $s = 0$. Suppose we are interested in “testing”¹⁰ if the observed low q_L is consistent with q^M or q^C . In the data we observe $q_L = 0.125$. This value for the low quantity is inside the identified set for q^M and so we cannot reject that firms are colluding at the monopoly quantity. On the other hand, $q_L = 0.125$ is outside the identified set for q^C and thus we can reject that firms are competing in Cournot fashion. A similar test can be implemented for $q_H = q^M$ and $q_H = q^C$ (observed $q_H = 0.167$). Finally, the test has no power if identified set for c is too wide (not very informative).

5 Extensions

The examples I discuss in the paper are simple enough such that it is easy to characterize the set of equilibrium payoffs (or an non-trivial upper bound of it). In more complicated games, the set of equilibrium payoffs are harder to compute. Nevertheless, methods have been developed to compute equilibria in more complicated infinitely repeated games. For example, Judd, Yeltekin and Conklin (2003) develop a computationally feasible algorithm to approximate the set of equilibria in infinitely repeated games with perfect monitoring and public randomization, following the approach proposed by Abreu (1988), Abreu, Pearce and Stacchetti (1986, 1990), and Cronshaw and Luenberger (1994). In on-going work, I explore how their method can be extended to games

¹⁰This is an informal test since I assume we have infinite data and zero sampling error.

without public randomization (nor mixed-strategies) using Milman's converse to the Klein-Milman theorem (see for example, Ok (2007)).

Finally, the method proposed in the paper can be adapted to dynamic games with states (i.e. stochastic games), as long as the set of equilibrium continuation payoffs can be computed, or at least its boundary. Dutta (1995) is an earlier attempt to compute equilibria in these games while Yeltekin, Cai and Judd (2015) a more recent take on this problem.

References

- Abreu, D. (1988), "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56 (2), 383-396.
- Abreu, D., D. Pearce, and E. Stacchetti (1986), "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 251-269.
- Abreu, D., D. Pearce, and E. Stacchetti (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58 (5). 1041-1063.
- Beresteanu, A., I. Molchanov, and F. Molinari (2011), "Sharp Identification Regions in Models With Convex Moment Predictions," *Econometrica*, 79 (6), 1785-1821.
- Ciliberto, F. and E. Tamer (2009), "Market Structure and Multiple Equilibria in Airline Markets", *Econometrica*, 77 (6), 1791-1828.
- Dal Bó, P. and G. R. Fréchette (2011), "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence," *American Economic Review*, 101, 411-429.
- Dutta, P. K. (1995), "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66, 1-32.

- Cronshaw, M. and D. G. Luenberger (1994), "Strongly Symmetric Subgame Perfect Equilibrium in Infinitely Repeated Games with Perfect Monitoring and Discounting," *Games and Economic Behavior*, 6, 220-237.
- De Paula, A. (2013), "Econometric Analysis of Games with Multiple Equilibria", *Annual Review of Economics*, 5, pp. 107-131.
- Fershtman, C. and Pakes, A. (2000), "A Dynamic Game with Collusion and Price Wars," *RAND Journal of Economics*, 31 (2), pp. 207-236.
- Galichon, A. and Henry, M. (2011), "Set Identification in Models with Multiple Equilibria," *Review of Economic Studies*, 78 (4), pp. 1264-1298.
- Green, E. J. and Porter, R. H. (1984), "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, 52 (1), pp. 87-100.
- Judd, K. L., Yeltekin, S. and Conklin, J. (2003), "Computing Supergame Equilibria," *Econometrica*, 71 (4), pp. 1239-1254.
- Ok, E. A. (2007) *Real Analysis with Economic Applications*, Princeton: Princeton Univ Press.
- Rotemberg, J. J. and Saloner, G. (1986), "A Supergame-Theoretic Model of Price Wars During Booms," *American Economic Review*, 76, pp. 390-407.
- Tamer, E. (2003), "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria'," *Review of Economic Studies*, 70, pp. 147-165.
- Yeltekin, S., Cai, Y. and Judd, K.L. (2015), "Computing equilibria of dynamic games," working paper, Carnegie Mellon University.