International Journal of Industrial Organization 31 (2013) 92–101

Contents lists available at SciVerse ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Equilibrium analysis of dynamic models of imperfect competition

Juan F. Escobar

Center for Applied Economics, Department of Industrial Engineering, University of Chile, Republica 701, Santiago, Chile

ARTICLE INFO

Article history:
Received 18 August 2010
Received in revised form 25 October 2012
Accepted 26 October 2012
Available online 5 November 2012

JEL classification:
C73
C61
C62

Keywords:
Industry dynamics
Dynamic stochastic games
Markov perfect equilibrium

ABSTRACT

Motivated by recent developments in applied dynamic analysis, this paper presents new sufficient conditions for the existence of a Markov perfect equilibrium in dynamic stochastic games. The main results imply the existence of a Markov perfect equilibrium provided the sets of actions are compact, the set of states is countable, the period payoff functions are upper semi-continuous in action profiles and lower semi-continuous in actions taken by rival firms, and the transition function depends continuously on actions. Moreover, if for each firm a static best-reply set is convex, the equilibrium can be taken in pure strategies. We present and discuss sufficient conditions for the convexity of the best replies. In particular, we introduce new sufficient conditions that ensure the dynamic programming problem each firm faces has a convex solution set, and deduce the existence of a Markov perfect equilibrium for this class of games. Our results expand and unify the available modeling alternatives and apply to several models of interest in industrial organization, including models of industry dynamics.

1. Introduction

This paper considers infinite horizon games in which at each period, after observing a payoff-relevant state variable, players choose actions simultaneously. The state of the game evolves stochastically parameterized by past history in a stationary Markov fashion. The setting includes a broad class of models, including Ericson and Pakes' (1995) model, as well as more general dynamic models of imperfect competition.

We present a general existence theorem for dynamic stochastic games and offer several applications to industrial organization. A strict implication from our main result, Theorem 1, is the following. A dynamic stochastic game possesses a behavior strategy Markov perfect equilibrium if the sets of actions are compact, the set of states is countable, the period payoff functions are upper semi-continuous in action profiles and lower semi-continuous in rivals' actions, and the probability distribution of the next state depends continuously on the actions chosen. Moreover, if for each player a static best-reply set is convex, the equilibrium can be taken in pure strategies.

As in previous work (Doraszelski and Satterthwaite, 2010; Horst, 2005), to obtain existence in pure strategies, we need to impose convexity restrictions on the dynamic game. Our result requires the game to have convex best replies, meaning that for all rivals' actions and all (bounded) continuation functions, each firm's static best-reply set is convex. This condition resembles (and indeed reduces to) the standard convexity restriction imposed on the payoff functions in strategic-form games to ensure the existence of Nash equilibrium. We state independent, sufficient conditions that ensure the convexity of the best replies. Our first sufficient condition is the uniqueness of the set of best replies, a condition requiring best-reply sets to be single-valued. This condition reduces to the convexity condition introduced by Doraszelski and Satterthwaite (2010) in an industry dynamics model. The second sufficient condition, satisfied by the so-called games with concave reduced payoffs, ensures each player's maximization problem is concave and so best replies are convex-valued. Although these two conditions do not cover all the games that have convex best replies, they significantly broaden the modeling alternatives that existing results offer.

Our main results have several applications; Section 4 provides a few. We analyze an industry dynamics model similar to that introduced by Ericson and Pakes (1995). Doraszelski and Satterthwaite (2010) have recently studied a version of the Ericson–Pakes model and introduced a condition, the *unique investment choice* (UIC) condition, to guarantee equilibrium existence. Under the UIC condition, best replies are single-valued and thus our convexity restrictions are met. Moreover, we provide a new alternative condition for the existence in the Ericson–Pakes model and discuss how this new condition permits modeling alternatives uncovered by Doraszelski and Satterthwaite's (2010) analysis. In particular, our results allow for multidimensional investment decisions and complementarities among firms' investments.

We also study a Markov Cournot game — in which firms compete in quantities, and at each round, a decision-controlled demand shock is realized. We provide sufficient conditions ensuring equilibrium existence. We show how restrictions on how rivals' actions affect payoffs and on how the transition function depends on action profiles make current results unsatisfactory (Amir, 1996; Horst, 2005; Nowak, 2007). Notably, to ensure equilibrium existence, we do not need to restrict the number

E-mail address: jescobar@dii.uchile.cl.

0167-7187/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijindorg.2012.10.005

of firms nor do we need to assume the transition function is linear in action profiles. We also consider a version of the Markov Cournot game in which firms have fixed costs, and show results ensuring the existence of behavior strategy equilibrium.

Finally, we also apply our results to an incomplete information dynamic model extensively studied and applied recently (e.g. Bajari et al., 2007; Doraszelski and Escobar, 2010).

Dating back to Shapley (1953), several authors have studied the problem of equilibrium existence in dynamic stochastic games. Among these, Mertens and Parthasarathy (1987), Nowak and Raghavan (1992), and Duffie et al. (1994) constitute important contributions that neither generalize nor are generalized by our results. Two strands of the literature are more closely related to this work. First, Horst (2005), Doraszelski and Satterthwaite (2010), and Nowak (2007) deal with the pure-strategy equilibrium existence problem. Some of these results cover state spaces uncovered by our results and prove not only existence but also uniqueness. Although our main result is formally unrelated to these authors', this paper identifies convexity conditions that expand and unify available modeling alternatives. Indeed, a game satisfying any of the convexity conditions those authors impose has convex best replies as required by our main result. Moreover, games such as those that Horst (2005) and Nowak (2007) consider are games with concave reduced payoffs and so, according to Proposition 3, have convex best replies.[1] This work contributes to this literature by identifying convexity restrictions that are significantly weaker than the conditions so far available.[2]

These results also contribute to the literature on dynamic games with countable state spaces. Federgruen (1978) and Whitt (1980) provide existence results that are corollaries to our main behavior strategy result, Corollary 2, in that they do not permit payoffs to be discontinuous. In particular, they do not deal with the problem of pure strategy existence, nor do they answer whether a nontrivial class of models could satisfy a convexity condition as the one we impose.

The paper is organized as follows. Section 2 presents the model. Section 3 presents and discusses the main theorem. Section 4 provides a number of applications of our results. Section 5 concludes. All proofs are in the appendix, except where the proof provides important intuition.

2. Setup

In this section we introduce our dynamic game model and define our equilibrium notion. Similar to many studies in industrial organization, we consider a dynamic stochastic game played by a finite set of firms. In each round of play, there is a payoff-relevant state variable (e.g., the identity of the incumbent firms). The state variable evolves stochastically, and firms can influence its evolution through actions (e.g., by entering or exiting the market). The goal of each firm is to maximize the expected present value of its stream of payoffs.

2.1. Model

There is a finite set of firms denoted by I. At the outset of period $t=1$, firms are informed about the initial state of the game, s_1. Then they simultaneously pick their actions $a_1=(a_1^i)_{i\in I}$. At the outset of period $t=2$, firms are informed of the new state of the game s_2 and then simultaneously pick their actions $a_2=(a_2^i)_{i\in I}$. And so on for $t\geq 3$.

The state space is S. For each firm i, the set of actions is A^i. In most applications, we will assume A^i is contained in $\mathbb{R}^{L^i}$, where L^i is a natural number, but allowing some more generality will be useful when studying models of imperfect competition in which firms have private information (see Section 4.3).

When firms make decisions at round t, they know the whole sequence of realized states $s_1,\ldots,s_t$, and past actions $a_1,\ldots,a_{t-1}$. The evolution of the state variable is Markovian in the sense that (a_t,s_t) fully determines the distribution over the state in the next round s^{t+1}. The Markovian transition function takes the form $\mathbb{P}[s_{t+1}\in B\,|(a_t,s_t)]=Q(B;a_t,s_t)$, where $B\subseteq S$. Given realized sequences of actions $(a_t)_{t\geq 1}$ and states $(s_t)_{t\geq 1}$, the total payoff for firm i is the discounted sum of period payoffs

$$\sum_{t=1}^{\infty}\left(\delta^i\right)^{t-1}\pi^i(a_t,s_t),$$

where $\delta^i\in[0,1[$ is the discount factor, and $\pi^i(a,s)$ is the per period payoff function.

This dynamic stochastic game model is flexible and, indeed, several models widely used in the literature fit into this framework. We will discuss applications and examples in detail in Section 4.

Throughout the paper, we will maintain the following assumptions.

- A1 S is a countable set.
- A2 For all i, A^i is compact and contained in a linear metric space.[3]
- A3 For all i, π^i is a bounded function.
- A4 The transition function Q is *setwise continuous* in $a\in A$ (Royden, 1968, Chapter 11.4): for every $B\subseteq S$ and $s\in S$, $Q(B;a,s)$ is continuous in $a\in A$.

In applications, Assumption (A1) is perhaps the most demanding one. While this assumption is usually made in industry dynamics models (Doraszelski and Satterthwaite, 2010; Ericson and Pakes, 1995), it rules out dynamic stochastic games in which the state variable is continuous. From Assumption (A3), we can define π_l^i and π_u^i as, respectively, the lower and upper bounds for the function π^i, and denote $\|\pi^i\|_\infty=\sup_{a\in A,s\in S}|\pi^i(a,s)|$.

2.2. Markov perfect equilibria

We now present the equilibrium notion with which we work. One may study subgame perfect equilibria of our dynamic model, but recent research has focused on Markov perfect equilibria. Markov perfect equilibria are a class of subgame perfect equilibrium strategies in which players condition their play only on payoff-relevant information.[4] The idea is that, in a given round, firms choose actions depending on the current state, with the purpose of maximizing the sum of current and future expected discounted payoffs.

A *Markov strategy* for firm i is a function $\bar{a}^i:S\to A^i$ mapping current states to actions. Thus, a Markov strategy defines a dynamic game strategy in which in each round t, firm i chooses action $\bar{a}^i(s_t)$, where s_t is the state realized in round t. A tuple of Markov strategies $\left(\bar{a}^i\right)_{i\in I}$ is a *Markov perfect equilibrium* if it is a subgame perfect equilibrium of the dynamic game. In a Markov perfect equilibrium, although firms condition their play only on the current state, they may deviate to arbitrary strategies conditioning on the whole transpired history. We will also consider *behavior Markov perfect equilibria*, defined as subgame perfect equilibria in which each firm i uses a strategy $\bar{a}^i:S\to\Delta(A^i)$ that maps current states to a distribution over actions.

[1] Bernheim and Ray (1989) and Dutta and Sundaram (1992) derive pure strategy results formally unrelated to ours. For a class of dynamic models, they restrict the strategy sets so that best replies are single valued and the games therefore satisfy the convexity restrictions required by our analysis.

[2] While Amir (1996) and Curtat (1996) restrict their attention to supermodular stochastic games, they do need to impose convexity conditions that, as we explain in Section 3.2, cannot be deemed as less stringent than ours.

[3] A linear metric space is a vector space endowed with a metric. For example, A^i could be a compact subset of $\mathbb{R}^{L_i}$ for some L_i.

[4] Several arguments in favor of this restriction can be given; see Maskin and Tirole (2001) for a particularly insightful discussion.