

STOCHASTIC GAMES IN ECONOMICS AND RELATED FIELDS: AN OVERVIEW

RABAH AMIR

*Université Catholique de Louvain
Louvain-la-Neuve, Belgium*

Abstract. This survey provides an extensive account of research in economics based on the stochastic games paradigm. Its area-by-area coverage is in the form of an overview, and includes applications in resource economics, industrial organization, macroeconomics, market games, and experimental and empirical economics. As to methodologically defined frameworks, the coverage is somewhat more detailed (to the extent that the material is not covered elsewhere in this volume), and includes the open-loop concept, the linear-quadratic model, myopic equilibrium, games of perfect information, and stochastic games with a continuum of players. It is hoped that the survey might be useful as a general guide both to economists and to game theorists.

1. Introduction

This section provides a general idea of the contents and organization of this survey of a large body of research in economics and related fields loosely defined by the adoption of the common methodology of stochastic games. We note at the outset that this class of games has also been referred to in various contexts as dynamic games, difference games, state-space games, sequential games and Markov games.¹ Given the breadth of this task, some omission is inevitable. We begin by describing the intended goals and limi-

¹Shapley coined the term “stochastic games” by analogy to “stochastic processes,” thus implicitly capturing the presence of dynamics. Since most applications actually involve models with deterministic transitions, this may appear somewhat misleading here, and “dynamic game” seems more appropriate. This is particularly true of studies considering open-loop equilibria, an essentially meaningless concept for games with chance moves.

tations of this survey, the confines and special features of stochastic games in economics, and the general organization of this survey.

Studies with general reward and transition functions and with monotonicity of equilibrium strategies and value functions as central features of the problem at hand are presented in the previous chapter, with a summary of the associated methodology of lattice programming.

1.1. PURPOSE AND SCOPE OF THE SURVEY

This chapter provides a general survey of applications of stochastic games in economics and related fields. We identify clusters of studies according to methodological considerations (e.g., reliance on open-loop equilibria or perfect information or computational simplicity), or to relevant subfields (e.g., industrial organization or resource economics). The primary concern has been to come up with convenient and natural categories that are consistent with the general purpose of this volume while appealing to a diverse readership. For strands of literature defined by a common methodological framework, a summary of the main results is provided. Otherwise, a list of references, along with some general descriptive comments, is given.

This survey will not encompass the continuous-time case, or differential games,² except in some cases where the results have direct qualitative analogs in discrete-time, or are otherwise of relevance to issues raised here. Likewise, although some links exist with the repeated games literature and with the standard two-stage game framework, these will not be dealt with here.

The reader is referred to the previous chapters for all general definitions and results associated with stochastic games, including the classification of the different types of strategy.

1.2. SPECIAL FEATURES OF ECONOMIC APPLICATIONS OF STOCHASTIC GAMES

In relating the present survey to the rest of this volume, one must keep in mind that a number of motivations and widely held beliefs among economists have in large part shaped the nature and the focus of the studies invoking the theory of stochastic games in economics. A brief account of these beliefs is now given. (For further discussion of various aspects, including the appropriateness of the different types of strategies, see [7], [35], [36], [64].)

1. *Discounted payoffs.* As situations where finite-time events are irrelevant are unnatural in economics, only models with discounted payoffs have been considered. The presence of a positive rate of interest is ubiquitous

²See [7], [22] and references therein.

in economic life. Thus, for economic models, the undiscounted case can be relevant only as a robustness check on a model with discounting.

2. *Pure strategies.* Due to a lack of compelling universal interpretation and to their inherent ex-post regret property, mixed strategies have enjoyed limited acceptance in economics in general, and this area is no exception. Mixed strategies have been considered only in a limited number of cases, when pure-strategy equilibria fail to exist.

3. *Uncountable state and action sets.* Owing mostly to the prevalence of calculus-based methods, there is a continuing tradition in economics of working with uncountable spaces, although the theory of stochastic games is much more complete for the case of finite state and action spaces, and reality sometimes conforms more to the latter case (e.g., discrete units for prices). Some recent models (listed below) do utilize finite spaces, though.

4. *Simplicity.* To avoid fixed-point arguments in function spaces and complex systems of functional equations, several studies rely on specific functional forms that allow closed-form equilibrium strategies, such as the linear-quadratic and the myopic models. A key advantage of this approach, in addition to the obvious computational appeal, is that it allows for clear-cut comparative statics conclusions, otherwise a rare luxury in dynamic games. Another simplicity-inspired choice is the nature of the strategies allowed, with many models being limited to open-loop behavior often without compelling contextual economic justification.

5. *Predictive power of models.* Since applications are typically motivated by the search for clear-cut conclusions, only highly structured and relatively aggregated models of stochastic games (typically with scalar state and action sets) have been studied. This is also due to the relatively complex nature of this class of games. Also, history-dependent behavior and folk theorem-type outcomes have generally been avoided in applications, with some exceptions.

1.3. ORGANIZATION OF THE SURVEY

In keeping with the overview nature of this survey, the only items covered in some detail are those that are not discussed elsewhere in this volume, and yet are important from the point of view of applications. Section 2 provides a summary of the properties of open-loop equilibria and a list of references by area. Section 3 deals with dynamic games of resource extraction. Section 4 considers the class of linear-quadratic games, ubiquitous in economics and systems theory. Section 5 presents the class of stochastic games with myopic equilibria, with applications. Section 6 collects various applications in the area of oligopoly theory that are not covered above. Section 7 mentions the case of perfect information games. Section 8 deals with stochastic games

with a continuum of players and some applications. Finally, Section 9 lists some work in empirical and in experimental economics.

2. Open-Loop Equilibrium in Deterministic Dynamic Games

Open-loop strategies are widely used in deterministic dynamic games (i.e., those with no chance moves). This section provides an overview of the main properties of open-loop equilibrium and then lists some of the studies in economics relying on such behavior.

2.1. DEFINITION AND PROPERTIES OF OPEN-LOOP STRATEGIES

Throughout this section, we consider a Markov dynamic game with deterministic transitions. An open-loop strategy is defined as a sequence of actions depending only on the initial state and on the date (or period). An open-loop strategy is thus a sequence of length $(T+1)$, where T is the last period in the (possibly infinite) horizon. Open-loop behavior rests on the premise that the players simultaneously commit at the beginning of the game to a completely specified list of actions to be played without any possibility of update or revision during the entire course of the game. Hence, no contingency planning of any sort is possible. An alternative way of thinking about open-loop strategies is as Markovian strategies³ where at each stage players use only constant functions of the current state. With open-loop strategies, a game may thus be viewed as a static game with sequences of length $(T+1)$ as strategy spaces.

Several important properties of open-loop equilibria are discussed next. To begin with, in *deterministic* Markov one-person dynamic optimization, there always exists an optimal open-loop strategy, so restricting oneself to open-loop policies results in no loss of value compared to using more sophisticated behavior. This fact is certainly intuitive, as is its failure in the presence of chance moves or stochastic transitions.

The game-theoretic analog of the above fact is perhaps less intuitive: in deterministic dynamic games, an open-loop equilibrium remains an equilibrium when the strategy spaces are expanded to include Markovian or history-dependent strategies. The reason is that if all of a given player's rivals are using open-loop strategies, the player cannot achieve a higher payoff by using more sophisticated strategies than open-loop. This follows directly by invoking the above fact for the player's best-response problem which, given the open-loop strategies of the rivals, is a *deterministic* Markov dy-

³In the systems theory and macroeconomics literature [7], Markovian strategies are usually referred to as feedback strategies, and sometimes as closed-loop (no-memory) strategies.

dynamic program.⁴ No general results are known about the comparison of equilibrium payoffs under open-loop vs. Markovian behavior.

Open-loop equilibria are generally not subgame-perfect.⁵ By contrast, open-loop optima in one-player deterministic problems clearly satisfy the principle of optimality, since the optimal Markovian and open-loop policies lead to the same actions and states at every period.

Open-loop equilibria are typically much simpler to analyze than Markovian equilibria. In particular, the usually difficult question of existence of pure-strategy equilibrium is most often straightforward in the open-loop case, where it amounts to using Brouwer's fixed-point theorem with the action set viewed as a subset of R^∞ (with the product topology), under standard conditions on the primitives. This relative simplicity is at the heart of the widespread use of open-loop strategies in the early stages of the adoption of stochastic games, despite the broad consensus that the commitment to a completely specified course of action over the indefinite future is not a realistic behavioral postulate in most cases of interest.⁶ The simultaneous presence of explicit long-term dynamics and of restricted static-like behavior seems contradictory. Furthermore, subgame perfection is broadly viewed as a desirable property of equilibrium behavior. Consequently, focus has markedly shifted towards Markovian strategies.

⁴This argument is similar to the better-known argument that Markovian (resp. Markov-stationary) equilibria of a Markov (resp. Markov-stationary infinite-horizon) stochastic game remain equilibria when history-dependent strategies are allowed. This also follows from the fact that with all rivals playing Markovian (resp. Markov-stationary) strategies, a player's best-response problem is a Markov (resp. Markov-stationary) dynamic program, for which there exists a Markov (resp. Markov-stationary) optimal policy. This argument is equally valid in the presence of chance moves (i.e., stochastic transitions). These important justifying arguments, as well as the so-called one-shot deviation principle, follow directly from the theory of dynamic programming.

⁵For the issue to be well defined, it is clear that one needs to assume that the action sets are essentially independent of the state. On the other hand, an equilibrium in Markovian strategies is always subgame-perfect in a strong sense: uniformly in the starting state.

⁶A very common framework of analysis adopted in industrial economics consists of modelling competing firms as making two decisions each, e.g., R&D levels and prices (or outputs). This can be done in a one-shot framework (with two decisions per firm), or in a two-stage game where R&D levels are chosen in the first stage, and prices are then chosen in the second stage, conditional on the observed R&D decisions. These two different timing structures can be viewed as relying on open-loop and closed-loop strategies, respectively. For instance, Brander and Spencer [13] provide a comparison of the two cases in a study of oligopolistic R&D. In such models, the use of open-loop strategies is much easier to justify as approximating real behavior, as it simply amounts to assuming that a firm does not get to observe its rivals' new technology before choosing its output level.

2.2. OPEN-LOOP EQUILIBRIUM IN ECONOMIC MODELS

Open-loop equilibrium originated and has been extensively analyzed in systems theory. See [7] for a detailed account. For problems with a linear-quadratic structure (covered in Section 4), open-loop equilibria are easily computed and characterized.

A class of applications that is of interest both from an economic and from a methodological point of view deals with continuous-time patent races. This class includes work by Loury [57], Lee and Wilde [53], and Reinganum [62], [63], among others. These papers a priori postulate differential games with stochastic duration corresponding to the occurrence of a success in an R&D project that would lead to a patent. The probability of a success for a firm follows an exponential distribution with parameter depending on the R&D expenditure of the firm. Due to the special structure of the model, in particular to the memoryless property of the exponential distribution, using Markovian strategies leads to an open-loop equilibrium, so that these games actually boil down to simple static games.

In the economics of natural resource exploitation and sustainability, studies that rely on the open-loop information structure tend to be older. They include, among many others, [69], [55], [21].

Various intrinsically dynamic problems in industrial organization were also considered under open-loop strategies early on. Spence [75] deals with investment in a new market, Spence [76] and Fudenberg and Tirole [34] propose models of the learning curve, Flaherty [32] studies dynamic limit pricing, and Flaherty [33] and Spence [77] are among the early attempts to model the effects of strategic process R&D.

The above list is far from complete, but can provide the reader with a flavor of the various approaches to, and results in, dynamic strategic competition relying on open-loop interaction.

3. Strategic Resource Extraction or Capital Accumulation

This is one of the areas of economics that has witnessed a high level of research activity involving stochastic/dynamic games as the key methodological approach. The seminal paper of Levhari and Mirman [54] considers two agents noncooperatively exploiting a natural resource. With z_t , a_t^i denoting the resource stock and Agent i 's consumption at time t , and β_i his discount factor, his payoff and the stock equation are given respectively by⁷

$$\sum_{t=0}^{\infty} \beta_i^t \log a_t^i \quad \text{and} \quad z_{t+1} = (z_t - a_t^1 - a_t^2)^\alpha, \quad \text{with } 0 < \alpha, \beta_i < 1.$$

⁷The one-player version of this problem is an example of the Solow-Cass-Koopmans growth model.

Using standard induction, Levhari and Mirman showed for this “Great Fish War” that (i) for every finite horizon with end-period T , there is a unique Markovian equilibrium with linear consumption strategies and logarithmic value functions,⁸ (ii) the limits of these strategies as $T \rightarrow \infty$, given for (say) Agent 1 by $\frac{\alpha\beta_2(1-\alpha\beta_1)z}{1-(1-\alpha\beta_1)(1-\alpha\beta_2)}$, constitute a Markov-stationary equilibrium of the infinite-horizon game,⁹ (iii) a tragedy of the commons prevails in both cases, in that the given equilibria are not Pareto-optimal and lead to over-consumption of the resource stock (relative to a Pareto-optimal path), and (iv) the equilibrium resource stock converges to a unique globally stable steady-state level given by $\bar{z} = \left\{ \frac{1}{\alpha\beta_1} + \frac{1}{\alpha\beta_2} - 1 \right\}^{\frac{\alpha}{\alpha-1}}$.

Cave [15] termed the “Cold Fish War” the situation where the two agents, observing the entire history of play, employ trigger strategies. Specifically, agents coordinate on cooperative extraction paths secured by the threat of reversion to the Markov-stationary strategies in case of defection. Assuming that $\beta_1 = \beta_2$, Cave characterizes the resulting open set of equilibria, which are clearly subgame-perfect. A simple necessary and sufficient condition is given for this set to include a Pareto-optimal extraction path.

The Levhari-Mirman analysis has been extended to more complex resource dynamics, including interactive fish species by Fisher and Mirman [30], [31]. Furthermore, related work dealing with general utility and growth functions includes Dutta and Sundaram [24], [25], [26]. Other general analysis papers, as well as studies of strategic bequests are covered in the previous chapter dealing with the lattice-theoretic approach [1].¹⁰

4. The Class of Linear-Quadratic Games

In a general linear-quadratic game, player i 's objective functional and the state equation are, for $t = 1, 2, \dots, T$,

$$\max \sum_{t=1}^T \frac{1}{2} \left\{ z'_{t+1} Q_{t+1}^i z_{t+1} + \sum_{j \in N} a_t^{j'} R_t^{ij} a_t^j \right\} \text{ and } z_{t+1} = A_t z_t + \sum_{j \in N} B_t^j a_t^j,$$

⁸Interestingly, the complementary choices of functional forms for the utility and biological growth functions in this model produce the same convenient qualitative results as in the linear-quadratic case. Here, due to the linearity of the equilibrium strategies, the value functions inherit the log nature of the utility function.

⁹Uniqueness of equilibrium in the infinite-horizon game remains an open question to date. There may exist infinite-horizon equilibria that are not necessarily related to the finite-horizon equilibrium strategies.

¹⁰There is also an extensive literature on strategic resource extraction in continuous time, which we will not cover here. The results therein are often quite different from their natural discrete-time counterparts.

where¹¹ z_t and a_t^j denote respectively the state vector (an element of $\mathbb{R}^n$) and player j 's action vector (an element of $\mathbb{R}^{l_j}$), at time t ; A_t , B_t^j , Q_{t+1}^i , and R_t^{ij} are matrices with appropriate dimensions, R_t^{ij} is negative definite, and Q_{t+1}^i is symmetric and negative semi-definite.

A Markov equilibrium can be given in closed form as follows. Let P_t^i be matrices satisfying, for $i = 1, 2, \dots, N$; $t = 1, 2, \dots, T$,

$$[R_t^{ii} + B_t^{i'} Z_{t+1}^i B_t^i] P_t^i + B_t^{i'} Z_{t+1}^i \sum_{j \neq i} B_t^j P_t^j = B_t^{i'} Z_{t+1}^i A_t, \quad (4.1a)$$

where the Z_t^i are defined recursively by

$$Z_t^i = F_t' Z_{t+1}^i F_t + \sum_{j \in N} P_t^{j'} R_t^{ij} P_t^j + Q_t^i, \text{ with } Z_{T+1}^i = Q_{T+1}^i, \quad (4.1b)$$

$$\text{and } F_t \triangleq A_t - \sum_{i \in N} B_t^i P_t^i.$$

There is a unique Markov equilibrium if and only if (4.1) has a unique solution set $\{P_t^{j*}\}$, with equilibrium strategies (specifying player i 's action vector at time t in a T -period horizon problem) and value function for player i from stage t onwards given by, for $i = 1, 2, \dots, n$; $t = 1, 2, \dots, T$,

$$\gamma_t^{i*} = -P_t^{i*} z_t, \text{ and } V_t^i(z_t) = \frac{1}{2} z_t' (Z_t^i - Q_t^i) z_t.$$

Some extensions of this class of games are now noted: (i) the state equation or the payoff functions may include additional linear terms (affine-quadratic games; the resulting equilibrium strategies are then affine functions of the state); (ii) exact conditions for P_t^{j*} to exist and be unique can be given in terms of invertibility of a composite matrix formed from the primitives of the problem; (iii) uncertainty in the form of an additive Gaussian vector (i.i.d. across time) in the state equation is easily incorporated, resulting in no qualitative changes in the solution; and (iv) the open-loop equilibrium is also easily computed.

Next, consider the infinite-horizon *undiscounted* stationary version of the game, obtained by letting $T = \infty$ and A , B^i , Q , R^{ij} be time-invariant. Sufficient conditions on the primitives that guarantee existence are not

¹¹Matrices are denoted by capital letters, vectors by lower-case letters and the transpose operation by a "prime" sign. Further details may be found in [7].

known at this point.¹² Nonetheless, the following partial answer (involving assumptions on derived objects) is known. Consider the following matrix equations, which are clearly limits of (4.1a-b):

$$[R^{ii} + B^{i'} \bar{Z}^i B^i] \bar{P}^i + B^{i'} \bar{Z}^i \sum_{j \neq i} B^j \bar{P}^j = B^{i'} \bar{Z}^i A, i = 1, 2, \dots, N, \quad (4.2a)$$

where $\bar{Z}^i$ is defined by

$$\bar{Z}^i = \bar{F}' \bar{Z}^i \bar{F} + \sum_{j \in N} \bar{P}^{j'} R^{ij} \bar{P}^j + Q^i, \text{ and } \bar{F} \triangleq A - \sum_{i \in N} B^i \bar{P}^i. \quad (4.2b)$$

Proposition 4.1 Suppose there exist two N -tuples of matrices $\{\bar{Z}^i, \bar{P}^i\}$ satisfying (4.2). Let $\bar{F}_i \triangleq A - \sum_{j \neq i} B^j \bar{P}^j$ and $\bar{Q}_i \triangleq Q^i + \sum_{j \neq i} \bar{P}^{j'} R^{ij} \bar{P}^j$. If the pair $(\bar{F}_i, B^i)$ is stabilizable¹³ and the pair $(\bar{F}_i, \bar{Q}_i)$ is detectable,¹⁴ then:
 (i) there is a Markov-stationary equilibrium where player i 's strategy is $\bar{\gamma}^{i*}(z) = -\bar{P}^i z$ and his (finite) payoff is $\frac{1}{2} z_1' \bar{Z}^i z_1$, and
 (ii) the resulting equilibrium system dynamics $z_{t+1} = \bar{F} z_t$ is stable (i.e., $\lim_{t \rightarrow \infty} D^t = 0$).

While (4.2) can be viewed as the limit of (4.1) as $T \rightarrow \infty$, (4.2) can have other solutions that are not related to the finite-horizon solution. Under the above assumptions of stabilizability and detectability, the latter would also constitute equilibria of the infinite-horizon game.

There is an extensive literature in various areas of economics analyzing models that constitute either a special case or a variant of the above framework (some in continuous-time). Furthermore, all infinite-horizon models have discounted rewards.

A very partial list of references follows. [29], [65], [66], [8] and [23] deal with dynamic oligopolistic competition. [56] discusses natural resources. [61], [51], [5], [17] and [44] deal with macroeconomic policy games. Some more examples are given in Section 6.

¹²By contrast, nice sufficient conditions are available in the one-player case (e.g., [11], pp. 73-80) and in the zero-sum case [6].

¹³This is defined as follows. The matrix $[B^i, \bar{F}_i B^i, \bar{F}_i^2 B^i, \dots, \bar{F}_i^{n-1} B^i]$ has full rank. Intuitively, this ensures the existence of a pair of strategies that will drive the state to 0 in finite time.

¹⁴This is defined by $(\bar{F}_i', \bar{Q}_i')$ being stabilizable.

5. Stochastic Games with Myopic Equilibrium

A stochastic game is said to have a myopic equilibrium if a static game can be constructed from the primitives of the stochastic game such that the infinite repetition of an equilibrium of the static game constitutes an equilibrium for the stochastic game. We provide sufficient conditions on the reward and transition functions ensuring the existence of a myopic equilibrium for a discounted stochastic game, and then list some applications of this approach. Our presentation follows [42].

Proposition 5.1 *Assume that a stochastic game is such that:*

- (i) *the reward function is additively separable: $r_i(z, a) = K_i(a) + L_i(z)$, $\forall z, a, i$.*
- (ii) *the transition law is state-independent: $\Pr(z_{t+1} = z' / z_t = z, a_t = a) = p(z' / a)$ or $z_{t+1} \sim \xi(a_t)$.*
- (iii) *the one-shot game where player i 's action set is A_i and his payoff is $\gamma_i(a) = K_i(a) + \beta_i E\{L_i[\xi(a)]\}$ has a pure-strategy equilibrium a^* .*
- (iv) *$P\{\xi(a^*) \in \{z : a^* \in A_z\}\} = 1$ (i.e., a^* is feasible in the next period for all current states).*

Then the strategy where, at every stage t , player i plays a_i^ if $z \in S(a^*)$ and any feasible action otherwise is a Markov-stationary equilibrium of the infinite-horizon game.*

There are several applications in economics and management science for which this class of games provides a natural framework of analysis. For an early attempt at bringing quantity and price competition together in an oligopoly model with inventory and uncertain demand, see [50]. Different one-player inventory control models have myopic optimal policies: see [42], Chapter 3 for references. In the context of fisheries, see [74]. Noncooperative advertising models with this special structure have also been analyzed by Monahan and Sobel [58]. A simple model of dynamic R&D competition with myopic equilibrium investment strategies is developed by Blonski [12].

6. Other Applications to Dynamic Oligopoly

In addition to the previously mentioned studies of dynamic inter-firm competition using stochastic games, this section describes other papers in industrial organization, clustered according to sub-area without regard to methodological considerations.¹⁵

¹⁵There is a very large body of literature in industrial economics dealing with two-stage games where firms typically make simultaneous long-term decisions in the first stage (such as R&D level, capacity, entry, or advertising, etc.), and, upon observing the outcome of the first stage, the firms make short-term decisions in the product market (price or output levels) in the second stage. While such games can generally be translated into the framework of (finite-horizon) stochastic games, we do not cover here the numerous examples available (see, e.g., [3] for one such example and some related discussion).

Among the models with truly dynamic strategic competition, there is one class characterized by price competition and some form of inertia on the part of consumers. Rosenthal [67] pioneered this literature with Bertrand duopoly competition under complete consumer loyalty, with one firm's market share as the natural state variable. He characterized a Markov-stationary equilibrium where prices remained above marginal costs indefinitely. By contrast, under less-than-complete consumer loyalty, Rosenthal [68] produces an ϵ -equilibrium in Markov-stationary strategies where prices converge with probability one to marginal costs. A distinctive feature of these papers, as well as of the follow-up piece by Chen and Rosenthal [16], is that a closed-form *mixed-strategy* equilibrium was considered and actually exhibited. A closely related strand of literature deals with long-run price competition when consumers face costs for switching between different buyers: see [28], [8], [59]. The latter paper also deals with mixed strategies.

Inter-firm racing models, which may be viewed to some extent as discrete-time extensions or analogs of the patent race models discussed in Section 2.2, have been investigated: [38], [39], [4], among others.

Learning-by-doing in Arrow's sense, whereby firms' production costs fall with production experience, also naturally gives rise to interesting phenomena of a dynamic character. Cabral and Riordan [14] characterize the long-term consequences of this feature in a Markov-stationary framework with firms' cumulative sales as the natural state variables.

Another natural source of strategic dynamic competition is also due to technological progress, but allows for market entry and exit. Here, technological progress is modelled as process R&D, with firms expending resources to lower their unit costs. For a rich model of industry dynamics featuring process R&D-type of strategic competition over time and allowing for endogenous entry and exit of firms, see [27]. Another strategic model of industry dynamics without process R&D is studied in [2]. Perfectly competitive (nonstrategic) models of industry dynamics are listed in Section 8.

7. Dynamic Games of Perfect Information

In some subfields of economics, another class of dynamic games that has been used with some frequency is characterized by perfect information: players move sequentially, with each player knowing the history of play, including the previous move. Perfect information has many simplifying features, leading to much more general existence results. A general framework, with uncountable action sets, has been developed and existence of pure-strategy subgame-perfect equilibrium proved in [37] and [40], generalizing the classical result erroneously attributed to Zermelo (see [70]).

An early application to duopoly is by Cyert and DeGroot [19], who

model long-term competition with firms moving alternately, each being committed to its choice in the off-period. This work has given rise to more general analysis, covered in the previous chapter, along with the theory of strategic bequests (under limited altruism.)

8. Stochastic Games with a Continuum of Players

This class of games has not been covered in the summer institute, aside from a brief mention. Our presentation here follows Bergin and Bernhardt [9], [10]. With a continuum of players, each player is identified by a characteristic, $\alpha \in \Lambda$, with α evolving stochastically over time. In addition, there is aggregate uncertainty, modeled as a Markov sequence of shocks over time $\{\theta_t\}_{t=1}^\infty$, $\theta_t \in \Theta$: at each period in time, an aggregate shock θ_t is realized in the per-period state space Θ . The full process is modeled as a joint distribution, ν , on the sequences of aggregate shocks, Θ^∞ . At time t , given a history of states $\theta^t = (\theta_1, \dots, \theta_t) \in \Theta^t$, the conditional distribution on Θ^∞ given θ^t is $\nu(\cdot/\theta^t)$.

Each player $\alpha \in \Lambda$ chooses an action, a , from a common action space A . In the stage game, a distributional strategy for the population is a joint distribution over players and actions — $\tau \in M(\Lambda \times A)$, the set of probability measures on $\Lambda \times A$. Preferences of a player at time t are given by a (uniformly bounded) function $r(\alpha, a, \tau_t, \theta_t)$. The characteristic of a player α evolves stochastically over time according to a transition kernel $P(d\alpha/\alpha, a, \tau_t, \theta_t)$. For the distribution on characteristics, a current distribution τ_t implies that the next distribution on characteristics is given by $\mu_{t+1}(X) = \int P(X/\alpha, a, \tau_t, \theta_t) d\tau_t$, so that τ_{t+1} must have marginal distribution μ_{t+1} . Players seek to maximize the present discounted value of payoffs.

Under continuity assumptions on the payoff functions and transition kernels (the latter in the weak* topology), a Markov equilibrium is shown to exist with the state variable being the triplet (μ, θ, v) where $v : \Lambda \rightarrow R$ is a continuous function.¹⁶ In equilibrium, at any state (μ, θ, v) , the value of the distributional strategy depends only on these variables, where $v(\alpha)$ gives the expected payoff of α in the remainder of the game, and in equilibrium this is the actual payoff. (The role of v is similar to the role of sunspots as an alternative coordinating device in the state space, used to achieve existence of Markov equilibrium.)

The above discussion is in the context of an environment where the aggregate distribution evolves deterministically, conditional on the value of θ_t . The second result uses a reformulation of aggregate uncertainty and

¹⁶When the data of the game is history-dependent, existence of a (non-Markovian) equilibrium is also shown.

focuses on the case where the aggregate distribution evolves stochastically, but where aggregate uncertainty is not explicitly separated.

This class of games constitutes a natural game-theoretic framework for analyzing dynamic perfect competition. There are two main strands of economic literature that consider models related to this form. The first deals with dynamic market games: see [71], [46], [47]. The second strand deals with purely competitive industry dynamics (relying on a price-taking assumption in partial equilibrium), with entry and exit over time: see, e.g., [45], [43], [52].

9. Empirical and Experimental Work

Last but not least, some empirical and experimental studies based on stochastic games have been conducted in recent times. For some examples of empirical work, see [60], [72], [73], among others. Walker and Wooders [79], [80] test the minmax hypothesis in a zero-sum Markov game model of tennis serves using field data from Wimbledon, and find better support than earlier work based on laboratory data. There are also experimental studies testing the ability of laboratory subjects to play in dynamic games. In a simple oligopoly market game, Keser [48] finds little support for theoretical predictions based on a unique finite-horizon Markov equilibrium. On the other hand, Herr, Gardner and Walker [41] find the theoretical solutions quite well confirmed by laboratory behavior in a common-property resource game. Keser and Gardner [49] report a good fit for subgame perfect equilibrium predictions at the aggregate, but not at the individual, level.

Acknowledgement. I would like to thank Jim Bergin, Jean-François Mertens, Abraham Neyman, Robert Rosenthal, Matt Sobel, and Sylvain Sorin for helpful exchanges concerning the material covered here. This paper was prepared for presentation at the NATO Advanced Study Institute on Stochastic Games and Applications held at SUNY-Stony Brook in July 1999.

References

1. Amir, R. (2003) Stochastic games in economics: The lattice-theoretic approach, in A. Neyman and S. Sorin (eds.), *Stochastic Games and Applications*, NATO Science Series C, Mathematical and Physical Sciences, Vol. 570, Kluwer Academic Publishers, Dordrecht, Chapter 29, pp. 443–453.
2. Amir, R. and Lambson, V. (2002) Entry, exit and imperfect competition in the long run, *Journal of Economic Theory*, to appear.
3. Amir, R. and Wooders, J. (2001) One-way spillovers, endogenous innovator/imitator roles and research joint ventures, *Games and Economic Behavior* **31**, 1–25.
4. Athey, S. and Schmutzler, A. (2001) Investment and market dominance, *RAND Journal of Economics* **32**, 1–26.

5. Barro, R. and Gordon, D. (1983) A positive theory of monetary policy in a natural rate model, *Journal of Political Economy* **91**, 589–610.
6. Basar, T. and Bernhardt, P. (1995) *H[∞] Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, 2nd ed., Birkhauser, Boston, MA.
7. Basar, T. and Olsder, G. (1999) *Dynamic Noncooperative Game Theory*, SIAM Classics, Philadelphia.
8. Beggs, A. and Klemperer, P. (1992) Multi-period competition with switching costs, *Econometrica* **60**, 651–666.
9. Bergin, J. and Bernhardt, D. (1992), Anonymous sequential games with aggregate uncertainty, *Journal of Mathematical Economics* **21**, 543–562.
10. Bergin, J. and Bernhardt, D. (1995) Anonymous sequential games: Existence and characterization, *Economic Theory* **5**, 461–489.
11. Bertsekas, D. (1976) *Dynamic Programming and Stochastic Control*, Academic Press, New York.
12. Blonski, M. (1999) When is rational behavior consistent with rules of thumb? A link between evolutionary terminology and neoclassical methodology, *Journal of Mathematical Economics* **32**, 131–144.
13. Brander, J. and Spencer, B. (1983) Strategic commitment with R&D: The symmetric case, *Bell Journal of Economics* **14**, 225–235.
14. Cabral, L. and Riordan, M. (1994) The learning curve, market dominance, and predatory pricing, *Econometrica* **56**, 1115–1140.
15. Cave, J. (1987) Long-term competition in a dynamic game: The Cold Fish War, *Rand Journal of Economics* **18**, 596–610.
16. Chen, Y. and Rosenthal, R. (1996) Dynamic duopoly with slowly changing consumer loyalties, *International Journal of Industrial Organization* **14**, 269–296.
17. Cohen, D. and Michel, P. (1988) How should control theory be used to calculate a time-consistent government policy?, *Review of Economic Studies* **55**, 263–274.
18. Crombs, D., Sevy, D. and Ponsard, J. (1987) Selection in dynamic entry games, *Games and Economic Behavior* **21**, 62–84.
19. Cyert, R. and DeGroot, M. (1970) Multiperiod decision models with alternating choice as the solution to the duopoly problem, *Quarterly Journal of Economics* **84**, 419–429.
20. Dana, R.-A. and Montrucchio, L. (1993) Stationary Markovian strategies in dynamic games, in R. Becker et al. (eds.), *General Equilibrium, Growth, and Trade. Vol 2. The Legacy of Lionel McKenzie*, Academic Press, San Diego, CA, pp. 331–351.
21. Dasgupta, P. and Heal, G. (1979) *Economic Theory and Exhaustible Resources*, Cambridge University Press, Cambridge.
22. Dockner, E., Jorgensen, S., Van Long, N. and Sorger, G. (2000) *Differential Games in Economics and Management Science*, Cambridge University Press, Cambridge.
23. Doganoglu, T. (2001) Dynamic price competition with persistent consumer taste, University of Kiel, Germany.
24. Dutta, P. and Sundaram, R. (1991) How different can strategic models be?, *Journal of Economic Theory* **60**, 42–61.
25. Dutta, P. and Sundaram, R. (1992) Markovian equilibrium in a class of stochastic games: Existence theorems for discounted and undiscounted models, *Economic Theory* **2**, 197–214.
26. Dutta, P. and Sundaram, R. (1993) The tragedy of the commons?, *Economic Theory* **3**, 413–26.
27. Ericson, R. and Pakes, A. (1995) Markov-perfect industry dynamics: A framework for empirical work, *Review of Economic Studies* **62**, 53–82.
28. Farrell, J. and Shapiro, C. (1988) Dynamic competition with switching costs, *RAND Journal of Economics* **19**, 123–137.
29. Fershtman, C. and Kamien, M. (1987) Dynamic duopolistic competition with sticky prices, *Econometrica* **55**, 1151–1164.
30. Fisher, R. and Mirman, L. (1994) Strategic dynamic interaction, *Journal of Eco-*

- nomic Dynamics and Control* **16**, 267–287.
31. Fisher, R. and Mirman, L. (1996) The complete fish wars: Biological and dynamic interactions, *Journal of Environmental Economics and Management* **30**, 34–42.
32. Flaherty, M. T. (1980) Dynamic limit pricing, barriers to entry, and rational firms, *Journal of Economic Theory* **23**, 160–182.
33. Flaherty, M. T. (1980) Industry structure and cost-reducing investment, *Econometrica* **48**, 1187–1209.
34. Fudenberg, D. and Tirole, J. (1983) Learning by doing and market performance, *Bell Journal of Economics* **14**, 522–530.
35. Fudenberg, D. and Tirole, J. (1986) *Dynamic Models of Oligopoly*, Harwood Academic Publishers, Chur, Switzerland.
36. Fudenberg, D. and Tirole, J. (1991) *Game Theory*, MIT Press, Cambridge, MA.
37. Harris, C. (1986) Existence and characterization of perfect equilibrium in games of perfect information, *Econometrica* **53**, 613–628.
38. Harris, C. and Vickers, J. (1985) Perfect equilibrium in a model of a race, *Review of Economic Studies* **52**, 193–209.
39. Harris, C. and Vickers, J. (1987) Racing with uncertainty, *Review of Economic Studies* **54**, 1–21.
40. Hellwig, M. and Leininger, W. (1987) On the existence of subgame-perfect equilibrium in infinite-action games of perfect information, *Journal of Economic Theory* **43**, 55–75.
41. Herr, A., Gardner, R. and Walker, J. (1997) An experimental study of time-independent and time-dependent externalities in the commons, *Games and Economic Behavior* **19**, 77–96.
42. Heyman, D. and Sobel, M. (1984) *Stochastic Models in Operations Research, Volume II: Stochastic Optimization*, McGraw Hill, New York.
43. Hopenhayn, H. (1992) Entry, exit, and firm dynamics in long-run equilibrium, *Econometrica* **60**, 1127–50.
44. Jensen, H. and Lockwood, B. (1998) A note on discontinuous value functions and strategies in affine-quadratic differential games, *Economics Letters* **61**, 301–306.
45. Jovanovic, B. (1982) Selection and the evolution of industry, *Econometrica* **50**, 649–670.
46. Karatzas, I., Shubik, M. and Sudderth, W. (1994) Construction of stationary Markov equilibria in a strategic market game, *Mathematics of Operations Research* **4**, 975–1006.
47. Karatzas, I., Shubik, M. and Sudderth, W. (1997) A strategic market game with secured lending, *Journal of Mathematical Economics* **28**, 207–247.
48. Keser, C. (1993) Some results of experimental duopoly markets with demand inertia, *Journal of Industrial Economics* **41**, 133–151.
49. Keser, C. and Gardner, R. (1999) Strategic behavior of experienced subjects in a common pool resource game, *International Journal of Game Theory* **28**, 241–252.
50. Kirman, A. and Sobel, M. (1974) Dynamic oligopoly with inventories, *Econometrica* **42**, 279–287.
51. Kydland, F. and Prescott, E. (1977) Rules rather than discretion: The inconsistency of optimal plans, *Journal of Political Economy* **85**, 473–491.
52. Lambson, V. E. (1992) Competitive profits in the long run, *Review of Economic Studies* **59**, 125–142.
53. Lee, T. and Wilde, L. (1980) Market structure and innovation: A reformulation, *Quarterly Journal of Economics* **94**, 429–436.
54. Levhari, D. and Mirman, L. (1980) The Great Fish War: An example using a dynamic Cournot-Nash solution, *Bell Journal of Economics* **11**, 322–334.
55. Lewis, T. and Schmalensee, R. (1980) On oligopolistic markets for nonrenewable natural resources, *Quarterly Journal of Economics* **95**, 475–491.
56. Lindsey, R. (1989) Import disruptions, exhaustible resources and intertemporal security of supply, *Canadian Journal of Economics* **22**, 340–363.

57. Loury, G. (1979) Market structure and innovation, *Quarterly Journal of Economics* **93**, 395–410.
58. Monahan, G. and Sobel, M. (1994) Stochastic dynamic market share attraction games, *Games and Economic Behavior* **6**, 130–149.
59. Padilla, J. (1995), Revisiting dynamic duopoly with switching costs, *Journal of Economic Theory* **67**, 520–530.
60. Pakes, A. and Ericson, R. (1998) Empirical implications of alternative models of firm dynamics, *Journal of Economic Theory* **79**, 1–45.
61. Pindyck, R. (1977) Optimal economic stabilization policies under decentralized control and conflicting objectives, *IEEE Transactions on Automatic Control* **22**, 517–530.
62. Reinganum, J. (1981) Dynamic games of innovation, *Journal of Economic Theory* **25**, 21–41.
63. Reinganum, J. (1982) A dynamic game of R&D: Patent protection and competitive behavior, *Econometrica* **50**, 671–688.
64. Reinganum, J. and Stokey, N. (1985) Oligopoly extraction of a common-property natural resource: The importance of the period of commitment in dynamic games, *International Economic Review* **26**, 161–173.
65. Reynolds, S. (1987) Capacity investment, preemption and commitment in an infinite-horizon model, *International Economic Review* **28**, 69–88.
66. Reynolds, S. (1991), Oligopoly with capacity adjustment costs, *Journal of Economic Dynamics and Control* **15**, 491–514.
67. Rosenthal, R.W. (1982) A dynamic model of duopoly with consumer loyalties, *Journal of Economic Theory* **27**, 69–76.
68. Rosenthal, R.W. (1986) Dynamic duopoly with incomplete consumer loyalties, *International Economic Review* **27**, 399–406.
69. Salant, S. (1976) Exhaustible resources and industrial structure: A Cournot-Nash approach to the world oil market, *Journal of Political Economy* **84**, 1079–1093.
70. Schwalbe, U. and Walker, P. (2001) Zermelo and the early history of game theory, *Games and Economic Behavior* **34**, 123–137.
71. Shubik, M. and Whitt, W. (1973) Fiat money in an economy with one nondurable good and no credit, in A. Blaquiere (ed.), *Topics in Differential Games*, North-Holland, Amsterdam, pp. 401–448.
72. Slade, M. (1995) Empirical games: The oligopoly case, *Canadian Journal of Economics* **28**, 368–402.
73. Slade, M. (1998) Optimal pricing with costly adjustment: Evidence from retail grocery prices, *Review of Economic Studies* **65**, 87–107.
74. Sobel, M. (1982) Stochastic fishery game with myopic equilibria, in L. Mirman and D. Spulber (eds.), *Essays in the Economics of Renewable Resources*, North-Holland, Amsterdam, pp. 259–268.
75. Spence, M. (1979) Investment strategy and growth in a new market, *Bell Journal of Economics* **10**, 1–19.
76. Spence, M. (1981) The learning curve and competition, *Bell Journal of Economics* **12**, 49–70.
77. Spence, M. (1984) Cost reduction, competition, and industry performance, *Econometrica* **52**, 101–121.
78. Stokey, N. and Lucas, R. with E. Prescott (1989) *Recursive Methods in Economics Dynamics*, Harvard University Press, Cambridge, MA.
79. Walker, M. and Wooders, J. (2000) Equilibrium play in matches: Binary Markov games, mimeo, University of Arizona.
80. Walker, M. and Wooders, J. (2001) Minmax play at Wimbledon, *American Economic Review*, to appear.