Games and Economic Behavior 104 (2017) 92–130

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Continuous-time stochastic games ☆

Abraham Neyman

Institute of Mathematics, and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel

ARTICLE INFO

Article history:
Received 2 April 2015
Available online 22 March 2017

JEL classification:
C73
C70

Keywords:
Stochastic games
Continuous-time games
Equilibrium
Uniform equilibrium

ABSTRACT

We study continuous-time stochastic games, with a focus on the existence of their equilibria that are insensitive to a small imprecision in the specification of players' evaluations of streams of payoffs.

We show that the stationary, namely, time-independent, discounting game has a stationary equilibrium and that the discounting game and the more general game with time-separable payoffs have an epsilon equilibrium that is an epsilon equilibrium in all games with a sufficiently small perturbation of the players' valuations.

A limit point of discounting valuations need not be a discounting valuation as some of the "mass" may be pushed to infinity; it is represented by an average of a discounting valuation and a mass at infinity.

We show that for every such limit point there is a strategy profile that is an epsilon equilibrium of all the discounting games with discounting valuations that are sufficiently close to the limit point.

1. Introduction

1.1. Motivating continuous-time stochastic games

A fundamental shortfall of economic forecasting is its inability to make deterministic forecasts. A notable example is the repeated occurrence of financial crises. Even though each financial crisis was forecasted by many economic experts, many other experts claimed just before each occurrence that "this time will be different," as documented in the book *This Time Is Different: Eight Centuries of Financial Folly* by Reinhart and Rogoff. On the other hand, many economic experts have forecasted financial crises that never materialized. Finally, even those who correctly predicted a financial crisis could not predict the time of its occurrence and many of them warned that it was imminent when in fact it occurred many years later.

An analogous observation comes from sports, e.g., soccer. In observing the play of a soccer game, or knowing the qualities of both competing teams, we can recognize a clear advantage of one team over the other. Therefore, a correct forecast is that the stronger team is more likely than the other team to score the next goal. However, it is impossible to be sure which team will be the first to score the next goal, and it is impossible to forecast the time of the scoring.

In both cases, the state of the economy and the score of the game can dramatically change in a split second, and players' actions impact the likelihood of each state change.

☆ This research was supported in part by Israel Science Foundation grant 1596/10.
E-mail address: aneyman@math.huji.ac.il.

http://dx.doi.org/10.1016/j.geb.2017.02.004

A game-theoretic model that accounts for the change of a state between different stages of the interaction, and where such change is impacted by the players' actions, is a stochastic game. However, no single deterministic-time dynamic game can capture the important common feature of these two examples: namely, the probability of a state change in any short time interval can be positive yet arbitrarily small. This feature can be analyzed by studying the asymptotic behavior in a sequence of discrete-time stochastic games, where the individual stage represents short time intervals that converge to zero and the transition probabilities to a new state also converge to zero. The limit of such a sequence of discrete-time stochastic games is a continuous-time stochastic game; see Neyman (2012, 2013).

The analysis of continuous-time stochastic games enables us to model and analyze the important properties that (1) the state of the interaction can change, (2) the stochastic law of states is impacted by players' actions, and (3) in any short time interval the probability of a discontinuous state change can be positive (depending on the players' actions), but arbitrarily small for sufficiently short time intervals.

Accounting for stochastic state changes – both gradual (continuous) and instantaneous (discontinuous) but with infinitesimal probability in any short time interval – is common in the theory of continuous-time finance; see, e.g., Merton (1992). However, in continuous-time finance theory the stochastic law of states is not impacted by agents' actions. Accounting for (mainly deterministic and more recently also stochastic) continuous state changes that are impacted by agents' actions is common in the theory of differential games; see, e.g., Isaacs (1999).

In the present paper we study the three above-mentioned properties, with a special focus on discontinuous state changes. To do this we study the continuous-time stochastic game with finitely many states, since if there are finitely many states, then any state change is discontinuous.

1.2. Difficulties with continuous-time strategies

A classical game-theoretic analysis of continuous-time games entails a few unavoidable pathologies, as for some naturally defined strategies there is no distribution on the space of plays that is compatible with these strategies, and for other naturally defined strategies there are multiple distributions on the space of plays that are compatible with these strategies. In addition, what defines a strategy is questionable. In spite of these pathologies, we will describe unambiguously equilibrium payoffs and equilibrium strategies in continuous-time stochastic games.

Earlier game-theoretic studies overcome the above-mentioned pathologies by restricting the study either to strategies with inertia, or to Markov (memoryless) strategies that select an action as a function of (only) the current time and state. See, e.g., Bergin and MacLeod (1993) and Simon and Stinchcombe (1989) for the study of continuous-time supergames, Perry and Reny (1993) for the study of continuous-time bargaining, Isaacs (1999) for the study of differential games, and Zachrisson (1964), Lai and Tanaka (1982), and Guo and Hernandez-Lerma (2003, 2005) for the study of continuous-time stochastic games, termed in the above-mentioned literature Markov games, or Markov chain games. A more detailed discussion of the relation between these earlier contributions and the present paper appears in Section 7.

The common characteristic of these earlier studies is that they each consider only a subset of strategies, so that a profile[1] of strategies selected from the restricted class defines a distribution over plays of the games, and thus optimality and equilibrium are well defined, but only within the restricted class of strategies. Therefore, in an equilibrium, there is no beneficial unilateral deviation within only the restricted class of strategies. Therefore, there is neither optimality nor nonexistence of beneficial unilateral deviation claims for general strategies.

The restriction to Markov (memoryless) strategies (see, e.g., Zachrisson, 1964; Rykov, 1966; Miller, 1968a, 1968b; Kakumanu, 1969, 1971, 1975; Guo and Hernandez-Lerma, 2009) is (essentially) innocuous in the study of continuous-time, discounted or finite-horizon Markov decision processes and two-person zero-sum stochastic games.

The two properties of the discounted or finite-horizon payoffs that make this restriction innocuous are: (1) impatience, namely, the contribution to the payoff of the play in the distant future is negligible, and (2) time-separability of the payoff, namely, the payoff of a play on $[0, \infty)$ is, for every $s > 0$, a sum of a function of the play on the time interval $[0, s)$ and a function of the play on the time interval $[s, \infty)$.

The Markov (memoryless) assumption is restrictive in the Markov decision processes with a payoff that is not time-separable and in the zero-sum stochastic games with non-impatient payoffs.

The Markov (memoryless) assumption is restrictive in the non-zero-sum game-theoretic framework. It turns out that in discounted stochastic games (with finitely many states and actions) a Markov strategy profile that is an equilibrium in the universe of Markov strategies is also an equilibrium in the universe of history-dependent strategies. However, the set of equilibrium strategies and equilibrium payoffs in the universe of Markov strategies is a proper subset of those in the universe of history-dependent strategies.

A fundamental shortcoming of the restriction to memoryless strategies (and to oblivious strategies, which depend only on the state process) arises in the study of approximate equilibria that are robust to small changes of the discounting valuation. For example, there is a continuous-time (and discrete-time, even two-person zero-sum, e.g., the Big Match, Blackwell and Ferguson, 1968) stochastic game for which there is no Markov strategy profile that is an approximate equilibrium in all the discounted games with a sufficiently small discounting rate.

[1] I.e., a list, one for each player.