Journal of Mathematical Economics 58 (2015) 46–60

Contents lists available at ScienceDirect

Journal of Mathematical Economics

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

Stochastic stability on general state spaces☆

Jonathan Newton

School of Economics, University of Sydney, Australia

ARTICLE INFO

Article history:
Received 3 October 2013
Received in revised form
30 November 2014
Accepted 25 March 2015
Available online 7 April 2015

Keywords:
Learning
Stochastic stability
General state space

ABSTRACT

This paper studies stochastic stability methods applied to processes on general state spaces. This includes settings in which agents repeatedly interact and choose from an uncountable set of strategies. Dynamics exist for which the stochastically stable states differ from those of any reasonable finite discretization. When there are a finite number of rest points of the unperturbed dynamic, sufficient conditions for analogues of results from the finite state space literature are derived and studied. Illustrative examples are given.

1. Introduction

The occurrence of social learning and the convergence of agents' behavior via processes of adaptive behavior is well-documented within economics (e.g. Chong et al., 2006; Selten and Apesteguia, 2005). The possibility of multiple resting points for such processes naturally leads one to question which of these stable states is more plausible from an economic perspective. Strongly influenced by evolutionary game theory (Smith and Price, 1973), a literature has grown that analyses the robustness of stable states of social learning dynamics to random errors made by players in their choice of action (Kandori et al., 1993; Young, 1993a). These ideas have been applied to a variety of economic situations, including bargaining (Binmore et al., 2003; Naidu et al., 2010), Nash demand games (Young, 1993b; Agastya, 1999), exchange economies (Serrano and Volij, 2008), local interaction on networks and the persistence of altruistic behavior (Eshel et al., 1998).

A common approach when assessing the robustness of stable states of social learning dynamics has been that pioneered by Kandori et al. (1993) and Young (1993a), building on the work of Freidlin and Wentzell (1984). Agents are assumed to make errors independently and when they do make an error are assumed to play a strategy chosen at random from a distribution with full support on a finite set of strategies. This imposes a mathematical structure on the process that leads to clear and appealing characterization results.

Unfortunately, such results cannot be straightforwardly applied when agents have non-finite sets of strategies.[1] Even assuming the convergence of the underlying social learning dynamic, the addition of random errors can lead to behavior which hinders efforts to obtain a clear cut characterization of the long run pattern of play. This paper takes up the task of analyzing the problems and intricacies which arise and, when there are a finite number of rest points of the underlying dynamic, determines a set of sufficient conditions which enable existing results to be applied to models with continuous state spaces. These conditions include a continuity requirement on error distributions and players' responses as a function of the current state, an asymptotic stability condition and a condition which ensures a specific type of discontinuity does not occur at rest points of the underlying dynamic. Examples are given showing how no subset of the conditions is sufficient on its own.

Fortunately, all of these conditions are satisfied for many common models found in economics. Typical error distributions of the kind described above coupled with the continuous best responses found in many models of industrial organization will often satisfy all of the conditions. This study applies the theory to linear quadratic games and to population models in the style of Kandori et al. (1993).

A related paper is that of Feinberg (2006), which also looks at discrete time, continuous state space processes. However, the paper in question imposes the strong condition that the perturbed

☆ I wish to thank Heather Battey, Yossi Feinberg, Johannes Hörner, George Mailath and the audience at the Econometric Society Australasian Meetings for helpful comments. The author is the recipient of a Discovery Early Career Researcher Award funded by the Australian Research Council (Grant Number: DE130101768).

E-mail address: jonathan.newton@sydney.edu.au.

[1] An early paper in the literature (Foster and Young, 1990) has an infinite state space and a continuous time dynamic in which perturbations are modeled as a Wiener process. However, it differs markedly from the majority of the literature, in which the error distributions are irrelevant to the stability results as long as they have full support.

http://dx.doi.org/10.1016/j.jmateco.2015.03.007
0304-4068/

process be governed by transition probabilities that are continuous functions of the current state of the process. The bulk of the analysis in the current paper concerns situations where this is not the case. Moreover, Feinberg considers a particular unperturbed dynamic and state space, whereas the current paper is more general in its scope. Schenk-Hoppé (2000) adapts the results of Freidlin and Wentzell (1984) and Ellison (2000) for use in finding stochastically stable states in a continuous strategy oligopoly model equipped with an imitation dynamic.

By considering finite state spaces, Young (1993a) dispenses with the need for regularity assumptions found in treatments of perturbed dynamics by Freidlin and Wentzell (1984), Kifer (1988, 1990). Specifically, all finite state spaces are compact, continuity requirements become unnecessary, and the probability of non-convergence to some stable state in given finite time need no longer be bounded by a function of error probabilities. The treatment of the current paper incorporates some finiteness in that the set of orders of magnitude of one step transition probabilities is taken to be finite. This allows us to use weaker continuity requirements on transition probabilities. We also dispense with compactness assumptions on the state space. From an economics perspective this enables, for example, the use of the Cartesian plane as the state space and the use of error probabilities which are independent across players.

The paper is organized as follows. Section 2 introduces the ideas of the paper via two motivating examples. Section 3 describes the processes of interest, gives convergence results, looks at transition probabilities between stable states, and defines a useful regularity property, showing how this property allows the problems associated with infinite state spaces to be circumvented. Section 4 gives sufficient conditions for this property to hold and discusses each of the conditions, giving examples of the problems which arise if any condition fails to hold. Section 5 gives examples. Section 6 solves an example from Section 2 for which our regularity condition fails to hold. Section 7 concludes. Formal proofs are relegated to the Appendix.

2. Motivating examples

This paper focuses on situations where agents follow some rule when deciding how to behave. The rule can be deterministic or random, cautious or hasty, imitative or best responding: any kind of behavioral bias or irregularity can be represented. Usually the rule is adaptive in the sense that an agent's behavior is intended to improve his lot. What really matters is that the rule has the Markov property: the past per se does not affect the future, although features of the present shaped by the past, including memories, are allowed to do so. We analyze situations where behavior over time will converge towards one of a number of stable states. As long as there is some probability of convergence to more than one stable state, this is predictively awkward. The possibility of random errors or idiosyncratic play justifies the introduction of perturbed versions of the process which help in obtaining long run predictions. There is a well-developed literature which deals with these problems for finite state spaces,[2] so the first question that must be addressed is whether there is benefit to be had from dealing directly with processes on general state spaces, rather than with finite discrete approximations.

2.1. Discretization can fail to represent the original process accurately

There is not always a suitable finite discretization of a process available. To illustrate, we present the following example. Consider a Markov process with state space $X = [0, 1] \subset \mathbb{R}$ endowed with the Euclidean distance metric. Let the process be governed by the Markov kernel $P(., .)$. The Markov kernel is a generalized analogue of transition probabilities on Markov chains. $P(x, A)$ gives the probability with which the process moves from state x to any state within a set of states A. For notational ease, for $y \in X$, we identify $P(., y) := P(., \{y\})$. Let $P(x, x^2) = 1$. This process has a set of stable states $\Lambda = \{0, 1\}$: from $x^* \in \Lambda$, $P(x^*, x^*) = 1$. We examine a perturbed variant of the process in which each period, with probability $1 - \varepsilon$ the unperturbed process is followed, and with probability ε the new state is drawn from the uniform distribution $U[\![0, 1]\!]$. This perturbed process has an invariant measure π_ε which converges to a measure with all weight on $\{0\}$ as $\varepsilon \to 0$: the set of *stochastically stable*[3] states is $\{0\}$.

Any discretized state space and process should satisfy some properties in order for it to be a reasonable representation of the original process. We suggest the following as reasonable restrictions on the discretized state space $X_\Delta \subseteq X$ and the discretized unperturbed process $P_\Delta(., .)$: (a) From a state $x \in X_\Delta$, if a set $A \subseteq X$ is reached with positive probability under the original process, then the closest states to A in X_Δ (under the original metric) are reached with positive probability under the discretized process $P_\Delta(., .)$; (b) If, from a state $x \in X_\Delta$, under the original process the set of states in X which are closer to $z \in X_\Delta$ than to any other point in X_Δ is never reached with positive probability, then z is never reached with positive probability under the discretized process; (c) Stable states of the original process are states of the discretized process and therefore stable states of the discretized process by (b).

We take as a discretization of the perturbation (the uniform distribution on X) any distribution on X_Δ that places positive probability on all states in X_Δ. Now, for any finite discretization satisfying our conditions, as $\varepsilon \to 0$, the limit of π_ε places positive probability on all states in $\{0, 1\}$: discretizing the process has given us one additional stochastically stable state.[4]

Finding the stochastically stable states of the original process in this section turns out to be simple. The reason for this is that far enough along any convergent path to a stable state, the probability under the perturbed process of moving to the basin of attraction of another given stable state is of constant order of ε. For example, from any convergent path to 0 under the unperturbed process $P(., .)$, at any given future period t the probability under the perturbed process of being in the basin of attraction of state is 1 is of order $\varepsilon^\infty = 0$. There do not exist convergent paths to 0 with escape probabilities of different orders of ε. We shall define Property C as the absence of multiple paths which converge to the same stable state and have different orders of escape probability. When Property C holds, we show that variants of results used heavily in the finite state space stochastic stability literature can be used. An important part of the current paper gives sufficient conditions under which Property C holds.

2.2. Multiple convergent paths

The next example can be considered as a model in which there are two possible focal points for a social norm. There are $n \geq 2$

2 See also Bergin and Lipman (1996), van Damme and Weibull (2002), Beggs (2005).

3 The use of the term 'stochastic stability' in the economics literature refers almost exclusively to states with positive weight under some limiting measure. Other uses of the term appear in the literature on dynamic processes. This paper follows the economic usage.

4 It may be remarked that, for this example, there exist sequences of finite discretizations such that the limit (of the sequence of discretizations) of the limits (as $\varepsilon \to 0$) of π_ε converges to the stochastically stable states of the original process. Such a sequence does not always exist, as is apparent from the example in Section 2.2.