Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 177 (2018) 315–364

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Riemannian game dynamics ☆

Panayotis Mertikopoulos [a], William H. Sandholm [b,*]

[a] Univ. Grenoble Alpes, CNRS, Inria, LIG, F-38000 Grenoble, France

[b] Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA

Received 2 June 2017; final version received 10 April 2018; accepted 9 June 2018
Available online 28 June 2018

Abstract

We study a class of evolutionary game dynamics defined by balancing a *gain* determined by the game's payoffs against a *cost of motion* that captures the difficulty with which the population moves between states. Costs of motion are represented by a Riemannian metric, i.e., a state-dependent inner product on the set of population states. The replicator dynamics and the (Euclidean) projection dynamics are the archetypal examples of the class we study. Like these representative dynamics, all Riemannian game dynamics satisfy certain basic desiderata, including positive correlation, local stability of interior ESSs, and global convergence in potential games. When the underlying Riemannian metric satisfies a Hessian integrability condition, the resulting dynamics preserve many further properties of the replicator and projection dynamics. We examine the close connections between Hessian game dynamics and reinforcement learning in normal form games, extending and elucidating a well-known link between the replicator dynamics and exponential reinforcement learning.

JEL classification: C72; C73

☆ We thank Josef Hofbauer and Dai Zusai for helpful discussions and comments, and we thank Marciano Siniscalchi and an anonymous referee for very thoughtful reports. Part of this work was carried out during the authors' visit to the Hausdorff Research Institute for Mathematics at the University of Bonn in the framework of the Trimester Program "Stochastic Dynamics in Economics and Finance". PM is grateful for financial support from the French National Research Agency (ANR) under grant no. ANR-GAGA-13-JS01-0004-01 and the CNRS under grant no. PEPS-REAL.net-2016. WHS is grateful for financial support under NSF Grants SES-1155135, SES-1458992, and SES-1728853 and ARO Grant MSN201957.

* Corresponding author.
E-mail addresses: panayotis.mertikopoulos@imag.fr (P. Mertikopoulos), whs@ssc.wisc.edu (W.H. Sandholm).

https://doi.org/10.1016/j.jet.2018.06.002

Keywords: Evolutionary game theory; Learning in games; Projection dynamics; Riemannian metrics; Replicator dynamics; Reinforcement learning

1. Introduction

Viewed abstractly, evolutionary game dynamics assign to every population game a dynamical system on the game's set of population states. Under most such dynamics, the vector of motion at a given population state depends only on payoffs and behavior at that state, implying that changes in aggregate behavior are determined by current strategic conditions. Such dynamics may thus be viewed as state-dependent rules for transforming current payoffs into feasible directions of motion.

In this paper, we introduce a family of evolutionary game dynamics under which the vector of motion z from any state x is obtained by balancing two forces. The first, the *gain from motion*, is obtained by adding the products of the strategies' payoffs at x with their rates of change under z. This quantity is the measure of agreement between payoffs and motion used in the standard monotonicity condition for game dynamics.[1] The second, the *cost of motion*, captures the difficulty with which the population moves from state x along vector z. Different specifications of these quadratic costs define different members of our family of dynamics. These costs are usefully represented by means of a *Riemannian metric*, a state-dependent inner product used to evaluate lengths of and angles between vectors of motion. Accordingly, the dynamics studied here, defined by maximizing differences between gains and costs, are called *Riemannian game dynamics*.

The two archetypal examples of Riemannian game dynamics are the replicator dynamics (Taylor and Jonker, 1978) and the (Euclidean) projection dynamics (Nagurney and Zhang, 1997), both derived from fairly simple structures. First, the replicator dynamics are derived from the *Shahshahani metric* (Shahshahani, 1979), under which the cost of increasing a strategy's relative frequency in the population is inversely proportional to said frequency. Second, the projection dynamics are obtained by measuring the cost of motion in the standard Euclidean fashion, independently of the population's current state. Other Riemannian metrics can be used in applications where different strategies have clear affinities, allowing the presence and performance of one strategy to influence the use of similar alternatives.

The metric's boundary behavior is the source of a fundamental dichotomy that is best explained by looking at our two prototypical examples above. Under the replicator dynamics: (*i*) the law of motion for every game is continuous; (*ii*) the set of utilized strategies remains constant along every solution trajectory; and (*iii*) the dynamics' rest points are the restricted equilibria of the game – the states at which all strategies in use earn the same payoff. In contrast, under the Euclidean projection dynamics: (*i*) the law of motion is typically discontinuous at the boundary of the simplex; (*ii*) the set of utilized strategies may change infinitely often along the same solution trajectory; and (*iii*) the dynamics' rest points are the Nash equilibria of the underlying game. Based on this behavior, we obtain a natural distinction between *continuous* and *discontinuous Riemannian dynamics*, each category sharing the boundary behavior of its prototype. In Section 4, we introduce a variety of examples of Riemannian dynamics from both classes;

[1] See Friedman (1991), Swinkels (1993), Sandholm (2001), Demichelis and Ritzberger (2003), and condition (PC) below.