



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.

© 2018 Elsevier Inc. All rights reserved.

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).

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.

then, in Section 5, we show how these and other Riemannian dynamics can be provided with microfoundations using suitably constructed revision protocols.

A basic aim of our analysis is to demonstrate that many basic properties of the replicator and Euclidean projection dynamics extend to our substantially more general setting. In Section 6, we show that Riemannian dynamics satisfy the basic desiderata for evolutionary game dynamics: they heed a payoff monotonicity condition known as *positive correlation*, they ensure local stability of regular interior ESSs, and they converge globally in the class of potential games. In the latter context, Riemannian game dynamics also provide a broad generalization of *Kimura's maximum principle* (Kimura, 1958; Shahshahani, 1979). This principle states that when agents are matched to play a normal form common interest game, the replicator dynamics move in the direction of maximal increase in average payoffs, provided that lengths of displacement vectors are evaluated using the Shahshahani metric. Extending this principle, we observe that Riemannian dynamics track the direction of steepest ascent of potential in any potential game, provided that displacements are evaluated using the Riemannian metric at hand.

Obtaining further results on stability, convergence, and global behavior requires additional structure on our dynamics – and hence on the underlying Riemannian metric. This structure is provided by an *integrability condition*. In prior work on game dynamics, such conditions have been imposed on the *vector fields* used to convert the strategies' payoffs into vectors of choice probabilities.² By contrast, the integrability condition employed here is imposed on the *matrix field* that defines a Riemannian metric, requiring that it be expressible as the Hessian of a convex function. We call this function the *potential* of the metric, and we refer to the resulting dynamics as *Hessian game dynamics*.³ Both the replicator dynamics and the Euclidean projection dynamics are members of this class. As we explain in Section 7, Hessian dynamics are continuous when their potential function becomes infinitely steep at the boundary of the simplex, leading to the distinction between *continuous* and *discontinuous* Hessian dynamics.

The key tool that we employ for the analysis of Hessian dynamics is the *Bregman divergence* (Bregman, 1967), an asymmetric measure of the “remoteness” of a given population state from any fixed target state.⁴ By using the Bregman divergence as a Lyapunov function, we prove global convergence to Nash equilibrium in strictly contractive games and local stability of evolutionarily stable states under Hessian game dynamics. We also show that certain distinctive properties of the replicator dynamics in normal form games extend to *all* continuous Hessian dynamics – in particular, the convergence of time averages of interior solutions to the set of Nash equilibria, and the existence of simple sufficient conditions for permanence. Finally, we show that strictly dominated strategies are eliminated under continuous Hessian dynamics, a conclusion which does not extend to the discontinuous regime.⁵

Related work. There are very close connections between the dynamics considered here and dynamics studied by Hofbauer and Sigmund (1990), Hopkins (1999), and Harper (2011). In order

² See Hart and Mas-Colell (2001), Hofbauer and Sandholm (2007), and Sandholm (2010a).

³ In the context of convex programming, gradient flows generated by Hessian Riemannian (HR) metrics of this sort have been explored at depth by Bolte and Teboulle (2003), Alvarez et al. (2004), Mertikopoulos and Staudigl (2018), and many others. Laraki and Mertikopoulos (2015) also examine the long-term rationality properties of a class of second-order, *inertial* game dynamics derived from HR metrics.

⁴ In the Shahshahani case, this boils down to the Kullback–Leibler divergence, which has seen wide use in the analysis of the replicator dynamics (Hofbauer and Sigmund, 1998; Weibull, 1995).

⁵ See Sandholm et al. (2008) and Section 7.4.

to have the machinery in place to make these connections clear, we postpone this discussion until Section 2.3.

There is a more surprising connection between Hessian dynamics and models of reinforcement learning in normal form games. Rustichini (1999), Hofbauer et al. (2009) and Mertikopoulos and Moustakas (2010) show that if players track the cumulative payoffs (or *scores*) of their strategies and choose mixed strategies at each instant by applying the logit choice rule to these scores, the evolution of mixed strategies is described by the replicator dynamics.⁶ Combining our analysis here with that of Mertikopoulos and Sandholm (2016), we show that Hessian dynamics derived from a steep potential function also describe the evolution of mixed strategies under reinforcement learning. In addition to substantially generalizing existing results, our analysis provides an intuitive explanation for the tight links between the two processes. Section 8 describes these and other connections between Hessian dynamics and reinforcement learning in detail.

2. Population games and evolutionary dynamics

Notation. Let $\mathcal{A} = \{\alpha_1, \dots, \alpha_n\}$ be a finite set. The real space spanned by \mathcal{A} will be denoted by $\mathbb{R}^{\mathcal{A}}$ and we will write $\delta_{\alpha\beta}$ for the Kronecker deltas on \mathcal{A} . We will also write $\mathcal{K} \equiv \mathbb{R}_+^{\mathcal{A}}$ for the nonnegative orthant of $\mathbb{R}^{\mathcal{A}}$, $\mathcal{K}^\circ \equiv \mathbb{R}_{++}^{\mathcal{A}}$ for its interior (the positive orthant), and $\mathbb{R}_0^{\mathcal{A}} = \{z \in \mathbb{R}^{\mathcal{A}} : \sum_{\alpha} z_{\alpha} = 0\}$ for the subspace of vectors whose components sum to zero. Finally, in a slight abuse of notation, we will write $\mathbb{R}^{\text{supp}(x)} = \{z \in \mathbb{R}^{\mathcal{A}} : z_{\alpha} = 0 \text{ whenever } x_{\alpha} = 0\}$ for the set of vectors in $\mathbb{R}^{\mathcal{A}}$ whose support is contained in the support of $x \in \mathbb{R}^{\mathcal{A}}$.

2.1. Population games

Throughout this paper we focus on games played by a population of nonatomic agents. Our analysis extends to the multi-population setting without significant effort, but we focus on single-population games for simplicity and notational clarity.

During play, each agent chooses an *action* (or *pure strategy*) from a finite set \mathcal{A} , and their payoff is determined by their choice of action and by the proportions $x_{\alpha} \in [0, 1]$ of the population playing each action $\alpha \in \mathcal{A}$. Collectively, these proportions define a *population state* $x = (x_{\alpha})_{\alpha \in \mathcal{A}} \in \mathbb{R}^{\mathcal{A}}$, and we write $\mathcal{X} = \Delta(\mathcal{A}) = \{x \in \mathbb{R}_+^{\mathcal{A}} : \sum_{\alpha} x_{\alpha} = 1\}$ for the set of population states (or *state space*) of the game. The payoff to an agent playing $\alpha \in \mathcal{A}$ when the population state is $x \in \mathcal{X}$ is given by an associated *payoff function* $v_{\alpha} : \mathcal{X} \rightarrow \mathbb{R}$, which we assume to be Lipschitz continuous. Putting all this together, a *population game* may be identified with a set of actions and their associated payoff functions, and will be denoted by $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$.

A population state $x^* \in \mathcal{X}$ is a *Nash equilibrium* (NE) of a population game \mathcal{G} if

$$v_{\alpha}(x^*) \geq v_{\beta}(x^*) \quad \text{for all } \alpha \in \text{supp}(x^*) \text{ and for all } \beta \in \mathcal{A}. \tag{NE}$$

If x^* satisfies (NE) and is *pure* (i.e. $x^* = e_{\alpha}$ for some $\alpha \in \mathcal{A}$), it is called a *pure Nash equilibrium* of \mathcal{G} ; if, in addition, (NE) holds as a strict inequality for all $\beta \notin \text{supp}(x^*)$, x^* is said to be a *strict equilibrium* of \mathcal{G} .

A *restriction* of a game \mathcal{G} is a population game $\mathcal{G}' \equiv \mathcal{G}'(\mathcal{A}', v')$ that is defined by a subset $\mathcal{A}' \subseteq \mathcal{A}$ of the original game's action set and by payoff functions v_{α} obtained by restricting

⁶ For related results, see also Börgers and Sarin (1997), Posch (1997), and Hopkins (2002).

the original payoff functions to the reduced state space $\mathcal{X}' = \Delta(\mathcal{A}')$ of \mathcal{G}' . If $x \in \mathcal{X}$ is a Nash equilibrium of some restriction of \mathcal{G} , it will be called a *restricted equilibrium*; as such, $x \in \mathcal{X}$ is a restricted equilibrium of \mathcal{G} if all strategies in its support earn equal payoffs.

Example 2.1 (*Matching in normal form games*). The simplest example of a population game is obtained by uniformly matching a population of agents to play a two-player symmetric normal form game with payoff matrix $A = (A_{\alpha\beta})_{\alpha,\beta=1}^n$. Aggregating over all matches, the payoff to an α -strategist when the population is at state $x \in \mathcal{X}$ is $v_\alpha(x) = \sum_{\beta \in \mathcal{A}} A_{\alpha\beta} x_\beta$.

Example 2.2 (*Potential games*). A population game \mathcal{G} is called a *potential game* (Monderer and Shapley, 1996; Sandholm, 2001) if there exists a *potential function* f defined on a neighborhood of \mathcal{X} such that

$$\frac{\partial f}{\partial x_\alpha} = v_\alpha(x) \quad \text{for all } \alpha \in \mathcal{A} \text{ and all } x \in \mathcal{X}. \tag{2.1}$$

Example 2.3 (*Contractive games*). A population game \mathcal{G} is called (*weakly*) *contractive* (Hofbauer and Sandholm, 2009) if

$$\sum_{\alpha \in \mathcal{A}} (v_\alpha(x') - v_\alpha(x))(x'_\alpha - x_\alpha) \leq 0 \quad \text{for all } x, x' \in \mathcal{X}. \tag{2.2}$$

If (2.2) binds only when $x = x'$, \mathcal{G} is called *strictly contractive*, whereas if (2.2) binds for all $x, x' \in \mathcal{X}$, \mathcal{G} is called *conservative*.⁷

2.2. Evolutionary dynamics

The term *evolutionary dynamics* refers to rules that assign to each population game \mathcal{G} a dynamical system on its state space \mathcal{X} . This is usually done by mapping each game to a *law of motion*, i.e. a differential equation of the form

$$\dot{x} = V(x). \tag{D}$$

In most cases, the *motion field* $V(x)$ of (D) is defined by introducing a mapping $(x, \pi) \mapsto \tilde{V}(x, \pi)$ from state/payoff pairs to vectors, and then specifying that $V(x) \equiv \tilde{V}(x, v(x))$. In what follows, we will focus exclusively on such dynamics.

To ensure that solutions to (D) remain in \mathcal{X} for all $t \geq 0$, $V(x)$ should not point outward from \mathcal{X} ; formally, $V(x)$ should lie in the *tangent cone* of \mathcal{X} at x , defined here as

$$\text{TC}_{\mathcal{X}}(x) = \{z \in \mathbb{R}_0^{\mathcal{A}} : z_\alpha \geq 0 \text{ whenever } x_\alpha = 0\}. \tag{2.3}$$

Under many evolutionary dynamics (including the replicator dynamics and other imitative dynamics), the support of $x(t)$ remains invariant under (D), implying in turn that the interior of each face of \mathcal{X} remains invariant under (D). When this is the case, $V(x)$ actually lies in the *tangent space* to \mathcal{X} at x , defined as

$$\text{T}_{\mathcal{X}}(x) = \{z \in \mathbb{R}_0^{\mathcal{A}} : z_\alpha = 0 \text{ whenever } x_\alpha = 0\} \subseteq \text{TC}_{\mathcal{X}}(x). \tag{2.4}$$

Clearly, for every interior state $x \in \mathcal{X}^\circ$, we have $\text{T}_{\mathcal{X}}(x) = \text{TC}_{\mathcal{X}}(x) = \mathbb{R}_0^{\mathcal{A}}$.

⁷ Hofbauer and Sandholm (2009) use the name *stable games* instead of contractive, but Sandholm (2015) proselytizes for the terms employed here. In convex analysis, condition (2.2) is called *monotonicity*.

A basic monotonicity criterion linking (D) with the underlying game requires *positive correlation* between the strategies’ payoffs and growth rates. Concretely, this means that

$$\sum_{\alpha \in \mathcal{A}} v_{\alpha}(x) V_{\alpha}(x) \geq 0 \quad \text{for all } x \in \mathcal{X}, \tag{PC}$$

with equality only if $V(x) = 0$.⁸ If (D) satisfies (PC), every Nash equilibrium of \mathcal{G} is a rest point of (D). For a detailed discussion, see Sandholm (2010b).

We provide two prototypical examples of evolutionary dynamics below:

Example 2.4 (*The replicator dynamics*). The quintessential evolutionary game dynamics are the *replicator dynamics* of Taylor and Jonker (1978):

$$\dot{x}_{\alpha} = x_{\alpha} \left[v_{\alpha}(x) - \sum_{\beta \in \mathcal{A}} x_{\beta} v_{\beta}(x) \right]. \tag{RD}$$

Example 2.5 (*The Euclidean projection dynamics*). The other fundamental example we consider is the *Euclidean projection dynamics* of Nagurney and Zhang (1997) (see also Friedman, 1991, and Lahkar and Sandholm, 2008). These are defined by

$$\dot{x} = \arg \min_{z \in \text{TC}_{\mathcal{X}}(x)} \|v(x) - z\|_2^2, \tag{PD}$$

where $\|z\|_2 = (\sum_{\alpha} z_{\alpha}^2)^{1/2}$ denotes the ordinary Euclidean norm on $\mathbb{R}^{\mathcal{A}}$. Geometrically, the dynamics (PD) are defined by taking the Euclidean projection of the payoff field $v(x)$ onto the tangent cone $\text{TC}_{\mathcal{X}}(x)$. Since $\text{TC}_{\mathcal{X}}(x) = \mathbb{R}_0^{\mathcal{A}}$ on the interior \mathcal{X}° of the simplex, we obtain the simple formula

$$\dot{x}_{\alpha} = v_{\alpha}(x) - \frac{1}{|\mathcal{A}|} \sum_{\beta \in \mathcal{A}} v_{\beta}(x), \tag{2.5}$$

valid for all interior $x \in \mathcal{X}^{\circ}$. For an explicit formula on the boundary of \mathcal{X} , see Example 4.2.

2.3. Antecedents

The class of dynamics studied here is a substantial generalization of both the replicator dynamics and the projection dynamics. We now describe works from an assortment of fields that are antecedents of our approach.

The replicator equation (RD) for common interest games is a basic model from population genetics (Schuster and Sigmund, 1983). The *fundamental theorem of natural selection*, attributed to Fisher (1930), states that natural selection among genes increases overall population fitness. Kimura (1958) introduced a corresponding maximum principle, showing that population fitness increases at a maximum rate under (RD) provided that one imposes a certain nonlinear constraint on the set of feasible changes in population frequencies (see Remark 3.2 in Section 3.3). Later, Shahshahani (1979) and Akin (1979) put Kimura’s maximum principle on a firm mathematical footing using tools from differential geometry – specifically, by introducing a suitable

⁸ This and closely related conditions are considered by Friedman (1991), Swinkels (1993), Sandholm (2001), and Demichelis and Ritzberger (2003).

Riemannian metric (see Section 3.2). The derivation of the replicator dynamics in the latter papers provides a basic instance of the geometric construction of Riemannian dynamics developed in Section 3.6, while our construction based on balancing gains and costs can be viewed as an extension of Kimura’s analysis (cf. Remark 3.2).

Hofbauer and Sigmund (1990) model natural selection in populations of animals whose traits are represented by elements of a continuous set. They assume that all members of the population share the same trait x , except for an infinitesimal group of mutants whose traits differ infinitesimally from x . The evolution of the preponderant trait x follows a gradient-like process, moving in the direction that agrees with the play of the most successful local mutants. To obtain variations on this process, Hofbauer and Sigmund (1990) use a Riemannian metric to define the size and shape of the neighborhood of local mutants. When the trait space is \mathcal{X} and the fitness of mutant y takes the linear form $\sum_{\alpha} y_{\alpha} v_{\alpha}(x)$, they showed that the evolution of x on the interior of \mathcal{X} is given by

$$\dot{x}_{\alpha} = \sum_{\beta \in \mathcal{A}} \left[g_{\alpha\beta}^{-1}(x) - \frac{\sum_{\gamma} g_{\alpha\gamma}^{-1}(x) \sum_{\gamma} g_{\gamma\beta}^{-1}(x)}{\sum_{\gamma, \kappa} g_{\gamma\kappa}^{-1}(x)} \right] v_{\beta}(x), \tag{2.6}$$

where $g(x)$ is a field of symmetric positive definite matrices that defines the Riemannian metric in question (see Section 3.2). Hofbauer and Sigmund (1990) then observed that under the Shahshahani metric, the system (2.6) boils down to the replicator dynamics (RD). As we shall see, (2.6) describes the dynamics studied in this paper at all states $x \in \mathcal{X}$ in what we call the *minimal-rank* case (cf. Section 3.4).

In the course of analyzing perturbed best response dynamics (Fudenberg and Levine, 1998) and variants of fictitious play (Brown, 1951), Hopkins (1999) introduced a class of game dynamics that are defined on the interior of \mathcal{X} as

$$\dot{x}_{\alpha} = \sum_{\beta \in \mathcal{A}} M_{\alpha\beta}(x) v_{\beta}(x). \tag{2.7}$$

Here $M(x)$ is a smoothly-varying field of symmetric matrices that are positive definite on $\mathbb{R}_0^{\mathcal{A}}$ and map constant vectors to 0. Hopkins (1999) showed that the linearization of these dynamics agrees with that of perturbed best response dynamics up to a positive affine transformation. As a result, the local stability of rest points of (2.7) agrees with that of the corresponding rest points of perturbed best response dynamics with sufficiently small noise levels. As we show in Appendix A.1, the dynamics (2.6) satisfy Hopkins’ conditions; conversely, all dynamics satisfying Hopkins’ conditions can be expressed in the form (2.6). Thus, on the interior of \mathcal{X} , the dynamics of Hopkins (1999) are equivalent to the dynamics studied here (Proposition A.3).

More recently, Harper (2011) used ideas from information geometry to define generalizations of the replicator dynamics, and employed concepts from Riemannian geometry to state and prove certain properties of the induced dynamics. Ignoring boundary issues, these dynamics are an important special case of ours – specifically, the class of separable dynamics that we introduce in Example 4.4.

Finally, we note here that there is a surprising and deep connection between the Hessian subclass of Riemannian game dynamics and a model of reinforcement learning recently examined by Mertikopoulos and Sandholm (2016). We explore this relation in detail in Section 8.

3. Riemannian game dynamics

3.1. Gains, costs, and dynamics

We now define the dynamics we study as balancing a *gain from motion*, determined from the game’s payoffs, against a *cost of motion*, a new primitive that captures the difficulty of motion along a given direction from a given state. To streamline our presentation, we focus below on interior states $x \in \mathcal{X}^\circ \equiv \text{int}(\mathcal{X})$, postponing the treatment of boundary states until the machinery needed to handle them is in place.

Given a population game $\mathcal{G}(\mathcal{A}, v)$, the *gain from motion* from state $x \in \mathcal{X}$ along $z \in \mathbb{R}^A$ is defined as

$$G^v(z; x) = \sum_{\alpha \in \mathcal{A}} v_\alpha(x) z_\alpha. \tag{3.1}$$

In words, the gain of motion measures the agreement between payoffs and vectors of motion as in the standard monotonicity criterion (PC). For an alternative interpretation, recall that the defining property (2.1) of a potential game with potential function f can be expressed as

$$\sum_{\alpha \in \mathcal{A}} \frac{\partial f}{\partial x_\alpha} z_\alpha = \sum_{\alpha \in \mathcal{A}} v_\alpha(x) z_\alpha \quad \text{for all } z \in \mathbb{R}^A \text{ and all } x \in \mathcal{X}. \tag{3.2}$$

The left-hand side of (3.2) is the rate of change in the value of potential as the state moves away from x along z . Viewed in this light, the gain $G^v(z; x)$ extends the notion of “the rate of increase in potential” to games that do not admit a potential function.⁹ In particular, the gain captures the alignment between the direction of motion z and the payoffs at state x ; it is also linearly homogeneous in z , so it grows linearly as one increases the speed of motion in a fixed direction.

By contrast, the *cost of motion* $C(z; x)$ is a primitive that represents the intrinsic difficulty of moving from state x along a given displacement vector z . We assume that the costs of motion are positive, smoothly varying with the population state x , and quadratic in z . It is convenient to define costs $C(z; x)$ for states x in the positive orthant $\mathcal{K}^\circ \equiv \mathbb{R}_{++}^A$ and for displacement vectors z in \mathbb{R}^A .¹⁰ Then since costs are positive and quadratic in z , the cost function can be expressed as

$$C(x; z) = \frac{1}{2} z^\top g(x) z, \quad \text{for all } z \in \mathbb{R}^A \text{ and all } x \in \mathcal{K}^\circ, \tag{3.3}$$

where g is a smooth assignment of symmetric positive definite matrices $g(x)$ to states $x \in \mathcal{K}^\circ$.

To use the above to define the dynamics at interior population states, we posit that the vector of motion from state $x \in \mathcal{X}^\circ$ maximizes the difference between the gain of motion $G^v(x; z)$ and the cost of motion $C(x; z)$, subject to feasibility:

$$\dot{x} = \arg \max_{z \in \mathbb{R}_0^A} [G^v(z; x) - C(z; x)]. \tag{3.4}$$

⁹ The logic here is similar to the original motivation for the definition of contractive games, which extends the idea of a game with a concave potential function to games that do not admit a potential (Hofbauer and Sandholm, 2009). The gain (3.1) is referred to as the “aggregate gross gain” by Zusai (2018) in his general analysis of Lyapunov functions for contractive games and evolutionarily stable strategies.

¹⁰ We can interpret \mathcal{K}° as the set of population states that could arise if the population size were allowed to vary. We could instead define costs only for states in \mathcal{X}° and displacement vectors in \mathbb{R}_0^A , at the price of additional abstraction: see Remark 3.1 below.

We refer to the dynamics (3.4) as *Riemannian game dynamics*, for reasons that we will soon make clear. Before doing so, we show how the leading examples of these dynamics are derived from the ansatz (3.4) through suitable choices of the cost function $C(x; z)$:

Example 3.1. A straightforward, state-independent choice for the cost of motion is

$$C(x; z) = \frac{1}{2} \sum_{\alpha \in \mathcal{A}} z_\alpha^2. \tag{3.5}$$

To solve the resulting maximization problem in (3.4), consider the Lagrangian

$$\Lambda(z, \mu; x) = \sum_{\alpha \in \mathcal{A}} \left[v_\alpha(x) z_\alpha - \frac{1}{2} z_\alpha^2 - \mu z_\alpha \right], \tag{3.6}$$

where the last term is associated with the motion feasibility constraint $\sum_{\alpha \in \mathcal{A}} z_\alpha = 0$. A direct differentiation gives the optimality condition $z_\alpha = v_\alpha(x) - \mu$, and the feasibility constraint yields $\mu = |\mathcal{A}|^{-1} \sum_{\alpha \in \mathcal{A}} v_\alpha(x)$. Substituting back into in (3.4) yields

$$\dot{x}_\alpha = v_\alpha(x) - \frac{1}{|\mathcal{A}|} \sum_{\alpha \in \mathcal{A}} v_\alpha(x). \tag{3.7}$$

As we discussed in Section 2 (cf. Example 2.5), the system (3.7) describes the (Euclidean) projection dynamics of Nagurney and Zhang (1997) on \mathcal{X}° .

Example 3.2. For a basic state-dependent choice for the cost of motion, let

$$C(x; z) = \frac{1}{2} \sum_{\alpha \in \mathcal{A}} \frac{z_\alpha^2}{x_\alpha} \tag{3.8}$$

The Lagrangian for the maximization problem in (3.4) is now

$$\Lambda(z, \mu; x) = \sum_{\alpha \in \mathcal{A}} \left[v_\alpha(x) z_\alpha - \frac{z_\alpha^2}{2x_\alpha} - \mu z_\alpha \right]. \tag{3.9}$$

Differentiating now yields the optimality condition $z_\alpha = x_\alpha v_\alpha(x) - \mu x_\alpha$, and feasibility implies that $\mu = \sum_{\beta \in \mathcal{A}} x_\beta v_\beta(x)$. Substituting in (3.4), we obtain

$$\dot{x}_\alpha = x_\alpha \left[v_\alpha(x) - \sum_{\beta \in \mathcal{A}} x_\beta v_\beta(x) \right]. \tag{3.10}$$

The system (3.10) defines the replicator dynamics of Taylor and Jonker (1978) (cf. Example 2.4). Although the derivation above assumed that x is interior, the expression (3.10) actually describes the replicator dynamics on all of \mathcal{X} ; we explain why this is so in Section 3.4.

3.2. Costs of motion and Riemannian metrics

We now proceed with a reinterpretation of the costs of motion using notions from geometry. The fundamental notion here is that of a *Riemannian metric*, a position-dependent variant of the ordinary (Euclidean) scalar product between vectors.¹¹

¹¹ To be clear, a Riemannian metric is not a metric in the sense of measuring distances between points in a metric space, but it induces such a distance function in a canonical way. For a comprehensive introduction to this topic, see the masterful account of Lee (1997, 2003).

To start, we recall that a *scalar product* on a subspace W of \mathbb{R}^A is a bilinear pairing $\langle \cdot, \cdot \rangle: W \times W \rightarrow \mathbb{R}$ which satisfies the following for all $w, w' \in W$:

- (1) *Symmetry*: $\langle w, w' \rangle = \langle w', w \rangle$.
- (2) *Positive definiteness*: $\langle w, w \rangle \geq 0$, with equality if and only if $w = 0$.

The *norm* of a vector $w \in W$ is then defined as

$$\|w\| = \langle w, w \rangle^{1/2}. \tag{3.11}$$

When $W = \mathbb{R}^A$, the definition above becomes most transparent by writing $w = \sum_{\alpha} w_{\alpha} e_{\alpha}$ and $w' = \sum_{\beta} w'_{\beta} e_{\beta}$ in the standard basis $\{e_{\alpha}\}_{\alpha \in A}$ of \mathbb{R}^A . Since $\langle \cdot, \cdot \rangle$ is positive definite and bilinear, there exists a positive-definite matrix $g = (g_{\alpha\beta})_{\alpha, \beta \in A}$ such that

$$\langle w, w' \rangle = \sum_{\alpha, \beta \in A} w_{\alpha} g_{\alpha\beta} w'_{\beta} = w^{\top} g w' \tag{3.12a}$$

and

$$\|w\|^2 = \sum_{\alpha, \beta \in A} w_{\alpha} g_{\alpha\beta} w_{\beta} = w^{\top} g w. \tag{3.12b}$$

The matrix g is known as the *metric tensor* of $\langle \cdot, \cdot \rangle$ and its components are $g_{\alpha\beta} = \langle e_{\alpha}, e_{\beta} \rangle$. Clearly, a scalar product is represented uniquely by its metric tensor and vice versa, so we will move freely between the two representations in what follows.

With all this in mind, a *Riemannian metric* on an open set U of \mathbb{R}^A is a C^1 -smooth assignment of scalar products $\langle \cdot, \cdot \rangle_x$ to each $x \in U$ – or, equivalently, as a smooth field $g(x)$ of symmetric positive-definite matrices on U . In other words, a Riemannian metric prescribes a way of measuring lengths of and angles between displacement vectors at each $x \in U$.

The similarity in notation between the above and the definition of costs of motion is not a coincidence. Looking back at (3.4), we see that costs of motion and Riemannian metrics are both defined by means of a C^1 -smooth field of symmetric positive-definite matrices, with costs and norms being related via

$$C(x; z) = \frac{1}{2} z^{\top} g(x) z = \frac{1}{2} \|z\|_x^2. \tag{3.13}$$

We summarize this connection as follows:

Observation 3.1. Specifying a cost function on \mathcal{K}° is equivalent to endowing \mathcal{K}° with a Riemannian metric.

Remark 3.1. Defining costs of motion and Riemannian metrics on the positive orthant \mathcal{K}° allows us to work in standard coordinates, and simplifies passing from one to the other. That being said, we could equally well have taken a more parsimonious approach by defining costs of motion $C(x; z)$ only for states $x \in \mathcal{X}^{\circ}$ and feasible displacement vectors $z \in \mathbb{R}_0^A$, and similarly working with Riemannian metrics that specify inner products $\langle \cdot, \cdot \rangle_x$ on \mathbb{R}_0^A for each $x \in \mathcal{X}^{\circ}$. In this approach, the equivalence between cost functions and Riemannian metrics can be derived from a standard bijection between quadratic forms and bilinear forms (see e.g., Friedberg et al., 2002, p. 433), but at the cost of an extra degree of abstraction.

Before proceeding, it is instructive to recast our previous examples in terms of Riemannian metrics:

Example 3.3. The *Euclidean metric* is defined by choosing $g(x)$ to be the identity matrix:

$$g(x) = I = \text{diag}(1, \dots, 1) \quad \text{for all } x \in \mathcal{K}^\circ. \tag{3.14}$$

This metric corresponds to the cost function $C(z; x) = \frac{1}{2} \sum_{\alpha} z_{\alpha}^2$ of Example 3.1, and yields the standard expressions $\langle w, w' \rangle_x = w^{\top} w'$ and $\|w\|_x = \sqrt{w^{\top} w}$, all independent of x .

Example 3.4. The *Shahshahani metric* is defined as

$$g(x) = \text{diag}(1/x_1, \dots, 1/x_n) \quad \text{for all } x \in \mathcal{K}^\circ. \tag{3.15}$$

This metric corresponds to the cost function $C(z; x) = \frac{1}{2} \sum_{\alpha} z_{\alpha}^2/x_{\alpha}$ of Example 3.2, and yields the Shahshahani inner product $\langle w, w' \rangle_x = \sum_{\alpha} w_{\alpha} w'_{\alpha}/x_{\alpha}$. In contrast to its Euclidean counterpart, the Shahshahani metric is state-dependent: For instance, since $\|e_{\alpha}\|_x = x_{\alpha}^{-1/2}$, the set of vectors at x with Shahshahani norm 1 is squeezed toward the x_{α} axis as x_{α} becomes small (cf. Fig. 1(b)).

We now present two further classes of metrics to which we return in Section 4:

Example 3.5. For $p \geq 0$, the *p-Shahshahani metric* is defined as

$$g(x) = \text{diag}(1/x_1^p, \dots, 1/x_n^p) \quad \text{for all } x \in \mathcal{K}^\circ. \tag{3.16}$$

This definition includes the Euclidean metric ($p = 0$) and the standard Shahshahani metric ($p = 1$) as special cases, and corresponds to the cost function

$$C(x; z) = \frac{1}{2} \sum_{\alpha \in \mathcal{A}} \frac{z_{\alpha}^2}{x_{\alpha}^p}. \tag{3.17}$$

Since $\frac{1}{x_{\alpha}^p} / \frac{1}{x_{\beta}^p} = (x_{\beta}/x_{\alpha})^p$, specifying larger values of p means raising the relative cost of changes in the use of rare strategies. For instance, if strategy α is half as prevalent in the population as strategy β , then changes in the use of α cost 2^p times as much as changes in the use of β .

Figs. 1(a)–1(c) illustrate the effects of increasing the value of p on costs of motion: when x_{α} is small, increasing p increases the cost of moving toward and away from the $x_{\alpha} = 0$ boundary relative to the cost of moving along this boundary.

Example 3.6. Let $\mathcal{A}_1, \dots, \mathcal{A}_m$ be a partition of \mathcal{A} into m groups of intrinsically similar strategies, let $[\alpha]$ denote the group containing strategy α , let $x_{[\alpha]} = \sum_{\beta \in [\alpha]} x_{\beta}$ denote the population share of all strategies that are “similar” to α in the above partition, and let $s > 0$ be a parameter representing the “strength” of the similarity relation. The *nested Shahshahani metric* is then defined as

$$g_{\alpha\beta}(x) = \begin{cases} \frac{\delta_{\alpha\beta}}{x_{\alpha}} + s \frac{1}{x_{[\alpha]}} & \text{if } \beta \in [\alpha], \\ 0 & \text{otherwise.} \end{cases} \tag{3.18}$$

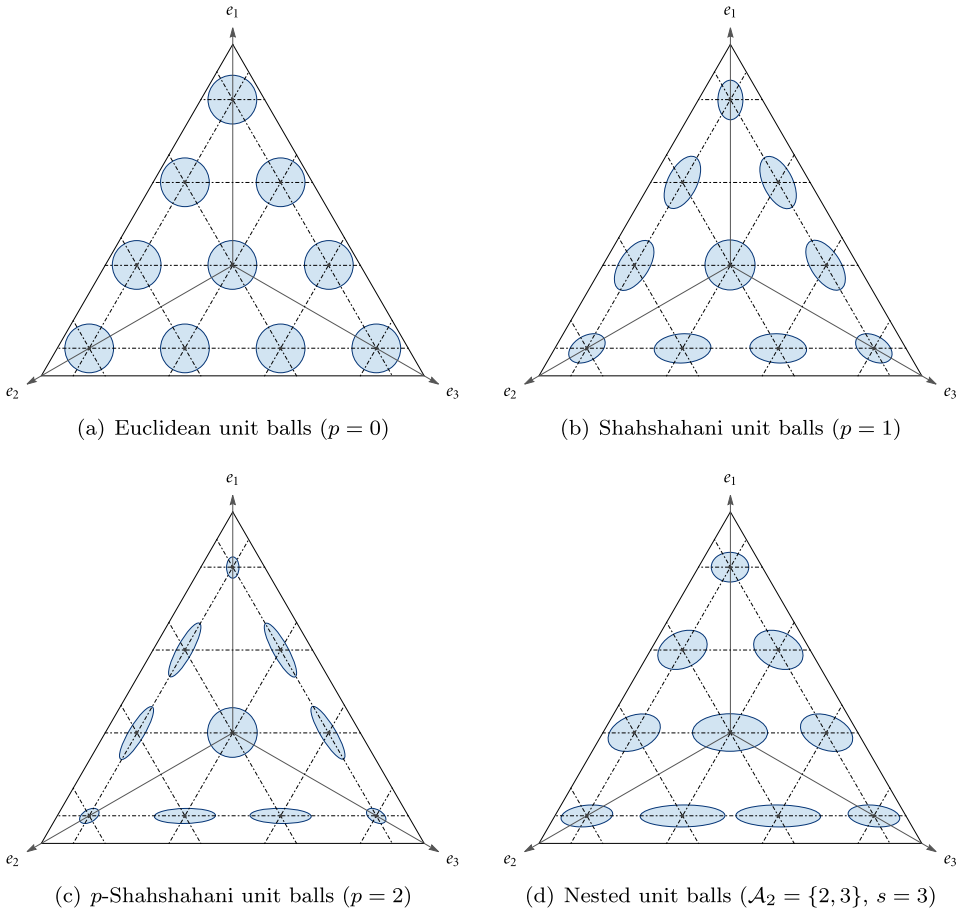


Fig. 1. Unit balls on the 3-simplex under the metrics of Examples 3.3–3.6. For each base point x shown, the shaded regions comprise all tangent vectors z based at x that satisfy $\|z\|_x^2 \leq 1$.

While the full expression for the cost function corresponding to the metric (3.18) is cumbersome, the cost of motion along the basic directions $e_\beta - e_\alpha$ takes a fairly simple form, namely

$$C(e_\beta - e_\alpha; x) = \begin{cases} \frac{1}{2} \left(\frac{1}{x_\alpha} + \frac{1}{x_\beta} \right) & \text{if } \beta \in [\alpha], \\ \frac{1}{2} \left(\frac{1}{x_\alpha} + \frac{1}{x_\beta} \right) + \frac{1}{2}s \left(\frac{1}{x_{[\alpha]}} + \frac{1}{x_{[\beta]}} \right) & \text{otherwise.} \end{cases} \quad (3.19)$$

Under (3.19), switches between strategies in the same group take the same form as under the Shahshahani cost function from Examples 3.2 and 3.4. On the other hand, switches between strategies in different groups are more costly, with the additional costs being inversely proportional to population shares of the groups and proportional to the strength of the similarity relation.

Fig. 1(d) illustrates the costs of motion (3.19) from various states in \mathcal{X} for partition $\mathcal{A}_1 = \{1\}$, $\mathcal{A}_2 = \{2, 3\}$ and similarity strength $s = 3$. The unit balls near the $x_1 = 0$ boundary are elongated along that boundary to a greater extent than the balls near the other boundaries. This reflects the fact that, mutatis mutandis, a unit of cost buys more motion between strategies 2 and 3 than between the other pairs of strategies.

3.3. Derivation of the dynamics: the interior case

With the above machinery at hand, we can provide an explicit description of the game dynamics under study on the interior \mathcal{X}° of \mathcal{X} . To do so, fix a population game $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$ and a Riemannian metric g on \mathcal{K}° . Then, by (3.13), the associated Riemannian game dynamics are

$$\dot{x} = \arg \max_{z \in \mathbb{R}_0^A} [G^v(z; x) - C(z; x)] = \arg \max_{z \in \mathbb{R}_0^A} \left[\sum_{\alpha \in \mathcal{A}} v_\alpha(x) z_\alpha - \frac{1}{2} \|z\|_x^2 \right]. \tag{3.20}$$

As in Examples 3.1 and 3.2, to obtain an explicit expression for the vector of motion that solves the maximization problem (3.20), consider the Lagrangian

$$\Lambda(z, \mu; x) = \sum_{\alpha \in \mathcal{A}} v_\alpha(x) z_\alpha - \frac{1}{2} \|z\|_x^2 - \mu \sum_{\alpha \in \mathcal{A}} z_\alpha. \tag{3.21}$$

Then, interpreting $(v_\alpha(x))_{\alpha \in \mathcal{A}}$ as a row vector (see Section 3.6) and writing $\mathbf{1} = (1, \dots, 1)^\top$ for the column vector of ones in the standard basis of \mathbb{R}^A , a simple differentiation yields the first-order optimality condition

$$v(x) = z^\top g(x) + \mu \mathbf{1}^\top. \tag{3.22}$$

Thus, after rearranging, we get

$$z = g^{-1}(x)[v(x)^\top - \mu \mathbf{1}], \tag{3.23}$$

where $g^{-1}(x)$ denotes the inverse of the matrix $g(x)$. Using the constraint $\sum_{\alpha} z_\alpha = 0$ to solve for μ and substituting in (3.20), some easy algebra leads to the explicit expression

$$\dot{x} = g^{-1}(x) \left[v(x)^\top - \frac{v(x)g^{-1}(x)\mathbf{1}}{\mathbf{1}^\top g^{-1}(x)\mathbf{1}} \mathbf{1} \right]. \tag{3.24}$$

Thus, if we set

$$v^\sharp(x) = g^{-1}(x)v(x)^\top \quad \text{and} \quad n(x) = g^{-1}(x)\mathbf{1}, \tag{3.25}$$

we obtain

$$\dot{x} = v^\sharp(x) - \frac{\langle v^\sharp(x), n(x) \rangle_x}{\|n(x)\|_x^2} n(x) = v^\sharp(x) - \frac{\sum_{\alpha \in \mathcal{A}} v_\alpha^\sharp(x)}{\sum_{\alpha \in \mathcal{A}} n_\alpha(x)} n(x) \tag{3.26}$$

We now revisit our two archetypal examples in the light of the explicit expression (3.26):

Example 3.7. If $g(x) = I$ is the Euclidean metric, we get $v^\sharp(x) = v(x)^\top$ and $n(x) = \mathbf{1}$, so (3.26) immediately boils down to (3.7). Therefore, when the cost of motion is defined using Euclidean lengths, the components of the displacement vector \dot{x} equal those of $v(x)$ up to a constant that ensures that $\dot{x} \in \mathbb{R}_0^A$.

Example 3.8. If $g(x) = \text{diag}(1/x_1, \dots, 1/x_n)$ is the Shahshahani metric of Example 3.4, we readily get $v_\alpha^\sharp(x) = x_\alpha v_\alpha(x)$ and $n_\alpha(x) = x_\alpha$, so (3.26) boils down to the replicator dynamics (3.10). Thus, when costs are defined using the Shahshahani norm of the population displacement vector, changes in the use of rare strategies are more costly than changes in the use of common ones. As a consequence, the initial term of \dot{x}_α is proportional to both the payoff $v_\alpha(x)$ of strategy α and to the mass x_α of agents playing strategy α . The second term ensures that $\dot{x} \in \mathbb{R}_0^A$, but here the normalization for strategy α is itself proportional to x_α .

Remark 3.2. The derivation above is closely related to Kimura’s (1958) derivation of the replicator dynamics in *common interest games*, i.e., games in which $v(x) = (Ax)^\top$ for some symmetric matrix A . Such games admit the potential function $f(x) = \frac{1}{2}x^\top Ax$ (cf. Eq. (2.1)), which reports one-half of the population’s average payoff. Using somewhat different language, Kimura (1958) proposed the population dynamics

$$\dot{x} = \arg \max \left\{ \sum_{\alpha} \frac{\partial f}{\partial x_{\alpha}} z_{\alpha} : z \in \text{TC}_{\mathcal{X}}(x) \text{ and } \|z\|_x^2 = \sigma_v^2(x) \right\}. \tag{3.27}$$

Here $\sum_{\alpha} \frac{\partial f}{\partial x_{\alpha}} z_{\alpha}$ is the rate of change of potential along z , $\|\cdot\|_x^2$ is the Shahshahani norm and $\sigma_v^2(x) = \sum_{\alpha} x_{\alpha} [v_{\alpha}(x) - \sum_{\beta} x_{\beta} v_{\beta}(x)]^2$ denotes the variance in the population’s payoffs at state x . It is easy to verify that (3.27) boils down to the replicator dynamics for the potential game $v(x) = (Ax)^\top$.

3.4. The boundary case

We now turn to an important dichotomy that arises when extending the definition of the dynamics (3.20) to the boundary of \mathcal{X} . To begin, recall from (2.4) that the *tangent space* $\text{T}_{\mathcal{K}}(x)$ to the nonnegative orthant \mathcal{K} at x is the linear subspace

$$\text{T}_{\mathcal{K}}(x) = \{z \in \mathbb{R}^A : z_{\alpha} = 0 \text{ whenever } x_{\alpha} = 0\} = \mathbb{R}^{\text{supp}(x)}. \tag{3.28}$$

We then say that a Riemannian metric g on \mathcal{K}° is *extendable* to \mathcal{K} if the map $x \mapsto g^{-1}(x)$ on \mathcal{K}° admits a (necessarily unique) C^1 -smooth extension to \mathcal{K} which we denote by g^{\sharp} (so $g^{\sharp}(x) \equiv g^{-1}(x)$ for all $x \in \mathcal{K}^{\circ}$), and which satisfies

$$\text{T}_{\mathcal{K}}(x) \subseteq \text{im } g^{\sharp}(x) \text{ for all } x \in \mathcal{K}. \tag{3.29}$$

In the above, $\text{im } g^{\sharp}(x)$ is the image (column space) of $g^{\sharp}(x)$; we henceforth call this set the *domain* of g at x and denote it by $\text{dom } g(x)$. Proposition B.1 in Appendix B shows that if g is extendable in the sense of (3.29), then the field of scalar products associated with g also admits a unique continuous extension from \mathcal{K}° to \mathcal{K} , with $\langle \cdot, \cdot \rangle_x$ defined on $\text{dom } g(x)$.

In what follows, we focus on two basic forms of extendability. First, if $\text{dom } g(x) = \mathbb{R}^A$ for all $x \in \mathcal{K}$, we say that g is *full-rank extendable*; instead, if $\text{dom } g(x) = \text{T}_{\mathcal{K}}(x)$ for all $x \in \mathcal{K}$, we say that g is *minimal-rank extendable*. We henceforth use the term “extendable” to refer to these two cases exclusively.

Example 3.9. The Euclidean metric has $g^{\sharp}(x) = g^{-1}(x) = I$ for all $x \in \mathcal{K}$, so it is full-rank extendable by default.

Example 3.10. The Shahshahani metric has $g^{\sharp}(x) = g^{-1}(x) = \text{diag}(x_1, \dots, x_n)$, so $\text{dom } g(x) = \text{T}_{\mathcal{K}}(x) = \mathbb{R}^{\text{supp}(x)}$ for all $x \in \mathcal{K}$. Thus, the Shahshahani metric is minimal-rank extendable, and the induced scalar product on $\text{dom } g(x) = \mathbb{R}^{\text{supp}(x)}$ is

$$\langle w, w' \rangle_x = \sum_{\alpha \in \text{supp}(x)} w_{\alpha} w'_{\alpha} / x_{\alpha} \text{ for all } w, w' \in \text{T}_{\mathcal{K}}(x). \tag{3.30}$$

Remark 3.3. Intuitively, minimal-rank extendable metrics partition \mathcal{K} into the relative interiors of each of its faces (including \mathcal{K}° itself). We will see that under the dynamics generated by such metrics, the relative interior of each face of \mathcal{X} is an invariant set.




Fig. 2. Admissible cones under the Euclidean and Shahshahani metrics. For $x \in \mathcal{X}^\circ$, we have $\text{Adm}_g(x) = \text{TC}\mathcal{X}(x)$. For $x \in \text{bd}(\mathcal{X})$, we still have $\text{Adm}_g(x) = \text{TC}\mathcal{X}(x)$ in the Euclidean case, but the Shahshahani metric can only be extended to the tangent space $\text{Adm}_g(x) = \text{T}\mathcal{X}(x)$.

To extend the definition of the dynamics to the boundary $\text{bd}(\mathcal{X})$ of \mathcal{X} , we introduce the cone of g -admissible vectors

$$\text{Adm}_g(x) = \text{TC}\mathcal{X}(x) \cap \text{dom } g(x). \tag{3.31}$$

This cone, which comprises all tangent vectors $z \in \text{TC}\mathcal{X}(x)$ that also lie in $\text{dom } g(x)$, specifies the possible directions of motion at a given state $x \in \mathcal{X}$. In particular, when $x \in \mathcal{X}^\circ$ is interior, we have $\text{TC}\mathcal{X}(x) = \text{T}\mathcal{X}(x) = \mathbb{R}_0^A$ and $\text{dom } g(x) = \mathbb{R}^A$. Thus, the g -admissible set is the hyperplane

$$\text{Adm}_g(x) = \mathbb{R}_0^A \cap \mathbb{R}^A = \mathbb{R}_0^A, \tag{3.32}$$

as anticipated in Eq. (3.20). Further instances of g -admissible cones are depicted in Fig. 2.

The restriction to $\text{dom } g(x)$ is needed because the norm $\|z\|_x^2$ is only defined for $z \in \text{dom } g(x)$. When g is extendable, the only case in which $\text{dom } g(x)$ is not all of \mathbb{R}^A occurs when $x \in \text{bd}(\mathcal{X})$ and g is minimal-rank extendable, in which case $\text{dom } g(x) = \text{T}\mathcal{X}(x) = \mathbb{R}^{\text{supp}(x)}$ (cf. Example 3.10).

3.5. Riemannian game dynamics

With all this at hand, we are finally in a position to extend the definition of the dynamics to all of \mathcal{X} . Concretely, building on (3.20), the Riemannian game dynamics induced by an extendable g are

$$\dot{x} = \arg \max_{z \in \text{Adm}_g(x)} [G^v(z; x) - C(z; x)] = \arg \max_{z \in \text{Adm}_g(x)} \left[\sum_{\alpha \in \mathcal{A}} v_\alpha(x) z_\alpha - \frac{1}{2} \|z\|_x^2 \right] \tag{RmD}$$

Equation (3.26) showed that (RmD) can be expressed at interior states as

$$\dot{x} = v^\sharp(x) - \frac{\langle v^\sharp(x), n(x) \rangle_x}{\|n(x)\|_x^2} n(x) = v^\sharp(x) - \frac{\sum_{\alpha \in \mathcal{A}} v_\alpha^\sharp(x)}{\sum_{\alpha \in \mathcal{A}} n_\alpha(x)} n(x) \tag{3.33a}$$

where $v^\sharp(x) = g^\sharp(x)v(x)^\top$ and $n(x) = g^\sharp(x)\mathbf{1}$. After a slight rearrangement, we can also express the dynamics as a linear transformation of payoffs $v(x)$:

$$\dot{x}_\alpha = \sum_{\beta \in \mathcal{A}} \left[g_{\alpha\beta}^\sharp(x) - \frac{n_\alpha(x)n_\beta(x)}{\sum_\gamma n_\gamma(x)} \right] v_\beta(x). \tag{3.33b}$$

Equation (A.2b) in Appendix A provides a concise third expression for the dynamics on \mathcal{X}° in terms of a pseudoinverse matrix.

If g is minimal-rank extendable, Proposition B.2 in Appendix B shows that (3.33) holds for all $x \in \mathcal{X}$, provided that one uses $g^\sharp(x)$ in the definition (3.25) of $v^\sharp(x)$ and $n(x)$. Proposition B.2 also shows that, in this case, one need only take the sums in the formulas (3.33) over the strategies in the support of x .

If instead g is full-rank extendable, extending (3.33) to boundary states requires solving a convex program whose inequality constraints may be active. For this reason, coordinate formulas for (RmD) may depend on the support of x – and, indeed, (RmD) may fail to be continuous at the boundary of \mathcal{X} (see Example 4.2 below). With this in mind, it will be convenient to call the dynamics generated by minimal-rank extendable metrics *continuous Riemannian dynamics*, and those generated by full-rank extendable metrics *discontinuous Riemannian dynamics*.

3.6. Geometric derivation of the dynamics

In (RmD), the dynamics’ vector of motion from x is defined to maximize the difference between the gain $G^v(z; x) = \sum_{\alpha \in \mathcal{A}} v_\alpha(x)z_\alpha$ and the cost of motion $C(z; x) = \frac{1}{2}\|z\|_x^2$ over the set of admissible vectors $z \in \text{Adm}_g(x)$. We now show how these dynamics can be derived using a purely geometric approach, generalizing Shahshahani’s (1979) derivation of the replicator dynamics in common interest games, and Nagurney and Zhang’s (1997) definition of the Euclidean projection dynamics. In what follows, we rely on some basic ideas from Riemannian geometry; for a comprehensive treatment, we refer again to Lee (1997).

To start, we introduce ideas about duality that explain our convention of writing payoffs using row vectors and the notations $v^\sharp(x)$ and $n(x)$ from Section 3.3. As in Section 3.2, let W be a subspace of $\mathbb{R}^{\mathcal{A}}$. A linear functional $\omega: W \rightarrow \mathbb{R}$ acting on vectors $w \in W$ is called a *covector*, and the space W^* of such functionals is called the *dual space* of W . We write $\langle \omega | w \rangle$ for the action of a covector $\omega \in W^*$ on a vector $w \in W$; to emphasize this pairing, the elements of W and W^* are also referred to as *primal* and *dual vectors* respectively. When $W = \mathbb{R}^{\mathcal{A}}$, we use the standard basis of $\mathbb{R}^{\mathcal{A}}$ to write everything in matrix notation, and distinguish vectors and covectors by writing primal vectors $w \in \mathbb{R}^{\mathcal{A}}$ as column vectors and dual vectors $\omega \in (\mathbb{R}^{\mathcal{A}})^*$ as row vectors. The action $\langle \omega | w \rangle$ of ω on w is then given by the matrix product $\omega w = \sum_\alpha \omega_\alpha w_\alpha$.

After mild manipulations, the definitions of Nash equilibrium (NE), positive correlation (PC), potential games (2.1) and contractive games (2.2) can be expressed in terms of expressions of the form $\sum_\alpha v_\alpha(x)z_\alpha$, where z is a tangent vector. Put differently, the payoff “vector” $(v_\alpha(x))_{\alpha \in \mathcal{A}}$ acts as a linear functional on displacement vectors, and so should be regarded as a *covector*. This is why we represent payoffs $v(x)$ in matrix notation as row vectors.

Example 3.11. The defining property (2.1) of potential games can be expressed as

$$\langle Df(x) | z \rangle = \langle v(x) | z \rangle \quad \text{for all } z \in \mathbb{R}^{\mathcal{A}} \text{ and all } x \in \mathcal{X}. \tag{3.34}$$

On the left-hand side, $Df(x)$ denotes the derivative of f at x , a linear functional that acts on tangent vectors $z \in \mathbb{R}^{\mathcal{A}}$ to yield the *directional derivative* $f'(x; z)$. Thus, (3.34) can be expressed as an equality between covectors, viz. $v(x) = Df(x)$.

Returning to our derivation of game dynamics, our aim in what follows is to find a vector field $x \mapsto V(x) \in \text{TC}(x)$ that agrees to the greatest possible extent with the payoff covector field $x \mapsto v(x)$, where this “agreement” is defined in terms of the Riemannian metric g . The derivation requires two steps: *i*) using a canonical transformation to convert the covector field into a vector field; and *ii*) projecting this field onto the cone of admissible vectors of motion.

For the first step, fix a Riemannian metric g on \mathcal{K}° that is extendable to \mathcal{K} as defined in Section 3.4. The *primal equivalent* of a covector $\omega \in (\mathbb{R}^A)^*$ at $x \in \mathcal{K}$ is the (necessarily unique) vector $\omega^\sharp \in \text{dom } g(x)$ such that

$$\langle \omega | w \rangle = \langle \omega^\sharp, w \rangle_x \quad \text{for all } w \in \text{dom } g(x). \tag{3.35}$$

In matrix notation, it is easy to verify that

$$\omega^\sharp = g^\sharp(x)\omega^\top, \tag{3.36}$$

in agreement with the definition $v^\sharp(x) = g^\sharp(x)v(x)^\top$ from (3.25).

For the second step, we transform each vector $v^\sharp(x)$ into a g -admissible vector by projecting it onto $\text{Adm}_g(x)$. Specifically, for all $x \in \mathcal{X}$ and $w \in \text{dom } g(x)$, the *projection* of w at x is defined as

$$\Pi_x(w) = \arg \min_{z \in \text{Adm}_g(x)} \|w - z\|_x. \tag{3.37}$$

The induced Riemannian dynamics are then defined as

$$\dot{x} = \Pi_x(v^\sharp(x)). \tag{3.38}$$

When $x \in \mathcal{X}^\circ$ is interior, we have $\text{Adm}_g(x) = \mathbb{R}_0^A$ by default, so $\Pi_x(w)$ is simply the orthogonal projection of $w \in \text{dom } g(x) = \mathbb{R}^A$ onto \mathbb{R}_0^A with respect to g . Accordingly, $\Pi_x(w)$ can be computed by finding a normal vector to \mathbb{R}_0^A and subtracting this vector’s contribution to w (as in the first step of the Gram–Schmidt orthonormalization process). To carry this out, observe that $\sum_\alpha z_\alpha = 0$ for all $z \in \mathbb{R}_0^A$, so the vector $n(x) = g^\sharp(x)\mathbf{1}$ defined in (3.25) satisfies

$$\langle n(x), z \rangle_x = n(x)^\top g(x)z = \mathbf{1}^\top g^\sharp(x)g(x)z = 0 \quad \text{for all } z \in \mathbb{R}_0^A. \tag{3.39}$$

This shows that $n(x)$ is a *normal vector* to $\text{T}_{\mathcal{X}}(x)$ with respect to $g(x)$. Thus, for all $x \in \mathcal{X}^\circ$, we can express the right-hand side of (3.38) as

$$\Pi_x(v^\sharp(x)) = v^\sharp(x) - \text{proj}_{n(x)} v^\sharp(x) = v^\sharp(x) - \frac{\langle n(x), v^\sharp(x) \rangle_x}{\|n(x)\|_x^2} n(x), \tag{3.40}$$

in agreement with (3.26).

More generally, for any state $x \in \mathcal{X}$ we have

$$\begin{aligned} \Pi_x(v^\sharp(x)) &= \arg \min_{z \in \text{Adm}_g(x)} \|v^\sharp(x) - z\|_x \\ &= \arg \min_{z \in \text{Adm}_g(x)} \left[\|v^\sharp(x)\|_x^2 + \|z\|_x^2 - 2\langle v^\sharp(x), z \rangle_x \right] \\ &= \arg \max_{z \in \text{Adm}_g(x)} \left[\langle v^\sharp(x), z \rangle_x - \frac{1}{2}\|z\|_x^2 - \frac{1}{2}\|v^\sharp(x)\|_x^2 \right] \\ &= \arg \max_{z \in \text{Adm}_g(x)} \left[\langle v(x) | z \rangle - \frac{1}{2}\|z\|_x^2 \right]. \end{aligned} \tag{3.41}$$

The dynamics (3.38) and (RmD) are therefore identical. We will take advantage of this geometric representation of (RmD) freely in what follows.

Remark 3.4. In addition to building on Kimura’s and Shahshahani’s derivations of the replicator dynamics, the dual representations of Riemannian game dynamics have a close analogue in a class of game dynamics called *target projection dynamics* (Friesz et al., 1994; Sandholm, 2005). These dynamics are defined on \mathcal{X} as

$$\dot{x} = \arg \min_{x' \in \mathcal{X}} \|v(x) - x'\|_2^2 - x. \tag{3.42}$$

Using a version of (3.41), Tsakas and Voorneveld (2009) showed that (3.42) can also be expressed as

$$\dot{x} = \arg \max_{x' \in \mathcal{X}} \left[\langle v(x) | x' \rangle - \frac{1}{2} \|x' - x\|_2^2 \right] - x. \tag{3.43}$$

4. Examples

We now present a variety of examples of Riemannian game dynamics. We start by extending the interior expressions (3.7) and (3.10) for our two prototypical dynamics to allow for boundary states:

Example 4.1 (*Replicator dynamics revisited*). Let g be the Shahshahani metric, so $g_{\alpha\beta}^\sharp(x) = \delta_{\alpha\beta} x_\beta$, $n_\alpha(x) = x_\alpha$, and $v_\alpha^\sharp(x) = x_\alpha v_\alpha(x)$. Since g is minimal-rank extendable, (3.33) yields the (continuous) Riemannian dynamics

$$\dot{x}_\alpha = x_\alpha \left[v_\alpha(x) - \sum_\beta x_\beta v_\beta(x) \right], \tag{RD}$$

which are the replicator dynamics of Taylor and Jonker (1978). The dynamics’ continuity is reflected in the fact that the formula (RD) is valid throughout \mathcal{X} .

Example 4.2 (*Projection dynamics revisited*). Let g be the Euclidean metric. Starting from formulation (3.38), Lahkar and Sandholm (2008) derived the following representation of the associated (discontinuous) Riemannian dynamics:

$$\dot{x}_\alpha = \begin{cases} v_\alpha(x) - |\mathcal{A}(x)|^{-1} \sum_{\beta \in \mathcal{A}(x)} v_\beta(x) & \text{if } \alpha \in \mathcal{A}(x), \\ 0 & \text{otherwise,} \end{cases} \tag{4.1}$$

where $\mathcal{A}(x)$ is a subset of \mathcal{A} that maximizes the average $|\mathcal{A}'|^{-1} \sum_{\beta \in \mathcal{A}'} v_\beta(x)$ over all subsets $\mathcal{A}' \subset \mathcal{A}$ that contain $\text{supp}(x)$. These are the projection dynamics (PD) of Nagurney and Zhang (1997). The discontinuity of (PD) is reflected in the appearance of $\text{supp}(x)$ in (4.1) via the definition of $\mathcal{A}(x)$.

Remark 4.1. The dynamics (RD) and (PD) highlight an important qualitative difference between Shahshahani and Euclidean projections, which is representative of continuous and discontinuous Riemannian dynamics respectively. The replicator dynamics (RD) comprise a Lipschitz continuous dynamical system on \mathcal{X} which preserves the face structure of \mathcal{X} , in that the relative interior of each face of \mathcal{X} remains invariant. By contrast, the projection dynamics (PD) may fail to be continuous at the boundary of \mathcal{X} . Thus, the relevant notion of a solution to (PD) is that of a

Carathéodory solution, which allows for kinks at a measure zero set of times. As a result, solutions of (PD) may leave and re-enter the relative interior of any face of \mathcal{X} in perpetuity.

The next example generalizes the previous two:

Example 4.3 (The p -replicator dynamics). For $p \geq 0$, let $g_{\alpha\beta}(x) = \delta_{\alpha\beta}x_{\beta}^{-p}$ denote the p -Shahshahani metric introduced in Example 3.5. We then have $g_{\alpha\beta}^{\sharp}(x) = \delta_{\alpha\beta}x_{\beta}^p$, $n_{\alpha}(x) = x_{\alpha}^p$, and $v_{\alpha}^{\sharp}(x) = x_{\alpha}^p v_{\alpha}(x)$. Thus, (3.33) yields the p -replicator dynamics

$$\dot{x}_{\alpha} = x_{\alpha}^p \left(v_{\alpha}(x) - \frac{\sum_{\beta \in \mathcal{A}} x_{\beta}^p v_{\beta}(x)}{\sum_{\beta \in \mathcal{A}} x_{\beta}^p} \right), \tag{4.2}$$

valid for all interior $x \in \mathcal{X}^{\circ}$. These dynamics were first defined by Harper (2011).

Since g is minimal-rank extendable if and only if $p > 0$, the dynamics are defined throughout \mathcal{X} via Eq. (4.2) for precisely these values of p . However, the dynamics are only Lipschitz continuous for $p \geq 1$; see Example 4.6 and Section 6.1 below. Three values of p are worth highlighting:

- (1) For $p = 0$, we obtain the projection dynamics (PD).
- (2) For $p = 1$, we obtain the replicator dynamics (RD).
- (3) For $p = 2$, we obtain the *log-barrier dynamics*, a system first examined by Bayer and Lagarias (1989) in the context of convex programming.¹²

In the above dynamics, the value of p parametrizes the costs of changes in the use of less common strategies. This is illustrated in Fig. 3, which presents a collection of p -replicator phase portraits in standard Rock-Paper-Scissors:

$$A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}. \tag{4.3}$$

When $p = 0$, displacement costs are independent of the current state; thus the circular form of the payoffs (4.3) generates circular closed orbits, subject to feasibility constraints (Fig. 3(a)). As p increases, the costs of motion for uncommon strategies become more important relative to the game’s payoffs (Figs. 3(b) and 3(c)). As a direct consequence, the closed orbits of the dynamics are “flattened” near each face of the simplex, and are ultimately reshaped into a nearly triangular form (Fig. 3(d)).¹³

Fig. 3 also illustrates a basic dichotomy between continuous and discontinuous Riemannian dynamics. In the discontinuous regime ($p = 0$), there is a unique forward solution from every initial condition in \mathcal{X} . However, solutions may enter and leave the boundary of \mathcal{X} , and solutions from different initial conditions can merge in finite time. In the smooth regime ($p \geq 1$), solutions exist and are unique in forward and backward time, and the support of the state remains fixed along each solution trajectory. Existence and uniqueness of solutions is treated formally in Section 6.1.

¹² The reason for this name is explained in Example 4.6; see especially Eq. (4.14).

¹³ That all of these dynamics feature closed orbits is not coincidental – see Proposition 7.5.




Fig. 3. Phase portraits of the p -replicator dynamics in standard Rock-Paper-Scissors. As $p \in [0, \infty)$ increases, the shape of the closed orbits changes from circular to triangular. When $p = 0$, solutions enter and leave the boundary of the simplex, but forward solutions exist and are unique. For $p \geq 1$, forward and backward solutions exist, are unique, and their support is constant.

Example 4.4 (*Separable metrics and their dynamics*). A Riemannian metric g on \mathcal{K}° is called *separable* if its metric tensor is of the form

$$g(x) = \text{diag}(1/\phi(x_1), \dots, 1/\phi(x_n)), \tag{4.4}$$

where $\phi: [0, \infty) \rightarrow [0, \infty)$ is a continuous *weighting function* that is strictly positive on $(0, \infty)$. For such metrics, we readily get

$$g^\sharp(x) = \text{diag}(\phi(x_1), \dots, \phi(x_n)), \tag{4.5}$$

so g is minimal-rank extendable if $\lim_{z \rightarrow 0^+} \phi(z) = 0$ and full-rank extendable otherwise.

When (3.33) applies, the dynamics induced by g take the form

$$\dot{x}_\alpha = \phi(x_\alpha) \left[v_\alpha(x) - \frac{\sum_\beta \phi(x_\beta) v_\beta(x)}{\sum_\beta \phi(x_\beta)} \right]. \tag{4.6}$$

Ignoring the dynamics’ behavior at the boundary, (4.6) was studied by Harper (2011) under the name *escort replicator dynamics*, and was further examined by Mertikopoulos and Sandholm (2016) and Bravo and Mertikopoulos (2017) in the context of game-theoretic learning (see Section 8). It is clear that the construction above generalizes immediately to allow different weighting functions for different strategies.

Moving beyond the separable case, Riemannian dynamics can also capture the effects of intrinsic relationships among the game’s strategies.

Example 4.5 (*Nested replicator dynamics*). In Example 3.6, we defined the nested Shahshani metric as

$$g_{\alpha\beta}(x) = \begin{cases} \frac{\delta_{\alpha\beta}}{x_\alpha} + s \frac{1}{x_{[\alpha]}} & \text{if } \beta \in [\alpha], \\ 0 & \text{otherwise,} \end{cases} \tag{4.7}$$

where $\mathcal{A}_1, \dots, \mathcal{A}_m$ is a partition of \mathcal{A} into groups of intrinsically similar strategies, $[\alpha]$ denotes the group containing strategy α , $x_{[\alpha]} = \sum_{\beta \in [\alpha]} x_\beta$, and s is a positive constant. A straightforward calculation shows that

$$g_{\alpha\beta}^\#(x) = \begin{cases} x_\alpha \delta_{\alpha\beta} - \frac{s}{1+s} \frac{x_\alpha x_\beta}{x_{[\alpha]}} & \text{if } \beta \in [\alpha], \\ 0 & \text{otherwise.} \end{cases} \tag{4.8}$$

It is evident from (4.8) that the metric g is minimal-rank extendable. Applying (3.33), we find that g generates the *nested replicator dynamics*:

$$\dot{x}_\alpha = x_\alpha \left[\frac{s}{1+s} \left(v_\alpha(x) - \frac{1}{x_{[\alpha]}} \sum_{\beta \in [\alpha]} x_\beta v_\beta(x) \right) + \frac{1}{1+s} \left(v_\alpha(x) - \sum_{\beta \in \mathcal{A}} x_\beta v_\beta(x) \right) \right] \tag{NRD}$$

if $x_\alpha > 0$ and $\dot{x}_\alpha = 0$ otherwise.

The imitative dynamics (NRD) were introduced by Mertikopoulos and Sandholm (2018) to model settings in which agents assess strategies using two distinct procedures: at rate $\frac{s}{1+s}$, an agent only compares the payoff of his current strategy α to those of strategies in group $[\alpha]$; at rate $\frac{1}{1+s}$, they compare the payoff of their current strategy to that of all other strategies.

Fig. 4 presents phase diagrams of the dynamics (NRD) with $s = 3$ in the standard Rock-Paper-Scissors game. The two panels illustrate the consequences of two similarity groupings. In each case, the longest “side” of each closed orbit corresponds to the pair of similar strategies, which are switched between more easily than the remaining pairs of dissimilar strategies.

The following class of dynamics incorporates all of our previous examples. It is examined at depth in Sections 7 and 8:

Example 4.6 (*Hessian Riemannian metrics and their dynamics*). A generalization of the above class of examples can be obtained by considering Riemannian metrics that are defined as Hessians of convex functions.¹⁴ To that end, let $h : \mathcal{K} \rightarrow \mathbb{R}$ be a continuous function on \mathcal{K} such that

¹⁴ For the origins of the idea in geometry, see Duistermaat (2001) and references therein; for applications to convex programming, see Bolte and Teboulle (2003) and Alvarez et al. (2004).




Fig. 4. Phase portraits of the nested replicator dynamics in standard Rock-Paper-Scissors with $s = 3$ for two similarity groupings.

- (i) h is C^3 -smooth on every positive suborthant of \mathcal{K} .
- (ii) $\text{Hess } h(x)$ is positive definite for all $x \in \mathcal{K}^\circ$.

Then, h induces a natural Riemannian metric on \mathcal{K}° defined as

$$g = \text{Hess } h, \tag{4.9}$$

or, in components:

$$g_{\alpha\beta}(x) = \frac{\partial^2 h(x)}{\partial x_\alpha \partial x_\beta}. \tag{4.10}$$

When this is the case, we say that g is a *Hessian Riemannian* (HR) metric and we refer to h as the *metric potential* of g .

As an example, the metric (3.18) that generates the nested dynamics (NRD) is an HR metric with potential

$$h(x) = \sum_{\alpha \in \mathcal{A}} x_\alpha (\log x_\alpha + s \log x_{[\alpha]}). \tag{4.11}$$

Moreover, every separable metric of the form (4.4) is an HR metric with potential

$$h(x) = \sum_{\alpha \in \mathcal{A}} \theta(x_\alpha) \tag{4.12}$$

for some smooth function $\theta: [0, +\infty) \rightarrow \mathbb{R}$ with $1/\theta''(z) = \phi(z)$. In particular, for $p \notin \{1, 2\}$, the p -replicator dynamics are generated by the potential

$$h_p(x) = \sum_{\alpha \in \mathcal{A}} \theta_p(x_\alpha) \quad \text{with} \quad \theta_p(z) = \frac{1}{(p-1)(p-2)} z^{2-p}, \tag{4.13}$$

and for $p = 1$ and $p = 2$, the corresponding potential functions are

$$h_1(x) = \sum_{\alpha \in \mathcal{A}} x_\alpha \log x_\alpha \quad \text{and} \quad h_2(x) = - \sum_{\alpha \in \mathcal{A}} \log x_\alpha, \tag{4.14}$$

Table 1
Regularity of the p -replicator dynamics and behavior of the metric potential function h_p .

p	Name	Regularity	Potential
0	projection	discontinuous	quadratic
(0, 1)	–	not Lipschitz	power law
1	replicator	smooth	Gibbs entropy
(1, 2)	–	smooth	Tsallis entropy
2	log-barrier	smooth	logarithmic
(2, ∞)	–	smooth	inverse power law

respectively.¹⁵ The values $p = 0$, $p = 1$, and $p = 2$ partition the class of p -replicator dynamics into six cases whose properties we summarize in Table 1.¹⁶

Definition (4.9) is an integrability condition on the matrix field g . As with vector fields on simply connected domains, integrability of the matrix field can be characterized by a symmetry condition on the derivatives of g ,¹⁷ namely

$$\frac{\partial g_{\alpha\gamma}}{\partial x_\beta} = \frac{\partial g_{\beta\gamma}}{\partial x_\alpha} \quad \text{for all } \alpha, \beta, \gamma \in \mathcal{A}. \tag{4.15}$$

Conditions (4.9) and (4.15) differ fundamentally from integrability conditions appearing in previous work on game dynamics, which are imposed on the *vector fields* that define the dynamics.¹⁸ In Section 7, we show that this integrability property provides important theoretical tools for the analysis of the induced Riemannian dynamics, which we call *Hessian game dynamics*.

5. Microfoundations via revision protocols

To provide microfoundations for deterministic game dynamics (D), one typically specifies a stochastic revision process that induces (D) in the so-called “mean field” limit. To do so, suppose that agents in the population are recurrently chosen at random and given the opportunity to switch strategies. What agents do when facing such opportunities is described by a *revision protocol* ρ whose components $\rho_{\alpha\beta}(x, \pi)$ describe the rates at which α -strategists who have received revision opportunities switch to strategy β , as a function of the current population state x and payoff vector π .¹⁹

Together, a population game $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$ and a revision protocol ρ induce the *mean dynamics*:

$$\dot{x}_\alpha = \sum_{\beta \in \mathcal{A}} [x_\beta \rho_{\beta\alpha}(x, v(x)) - x_\alpha \rho_{\alpha\beta}(x, v(x))], \tag{MD}$$

¹⁵ It is possible to define the potential h_p for all values of p using a single formula. Let $\theta_p(z) = (z^{2-p} + p(p-2)z - (p-1)^2)/((p-1)(p-2))$ when $p \neq \{1, 2\}$, and define θ_1 and θ_2 by analytic continuation. Linear and constant terms do not affect the resulting metric, and the explicit formulas for θ_1 and θ_2 follow from the fact that $\lim_{a \rightarrow 0} (z^a - 1)/a = \log z$.

¹⁶ When $p \geq 2$, the potential h_p becomes infinite on the boundary of \mathcal{K} , violating a standing assumption for h ; we address this technicality in Remark 7.1. Also, the (negative) *Tsallis entropy* (Tsallis, 1988) mentioned in Table 1 is defined as $S_q(x) = (q-1)^{-1} \sum_\alpha (x_\alpha^q - x_\alpha)$ for $q \in (0, 1)$.

¹⁷ This characterization follows from the integrability condition for ordinary vector fields (i.e., symmetry of the Jacobian matrix) and the symmetry of $g(x)$.

¹⁸ See Hart and Mas-Colell (2001), Hofbauer and Sandholm (2009), and Sandholm (2010a, 2014).

¹⁹ Weibull (1995) and Björnerstedt and Weibull (1996) introduce revision protocols for imitative dynamics. Sandholm (2010b, 2015) and Izquierdo et al. (2018) extend this approach to more general classes of dynamics.

which describe the rate of change in the use of each strategy α as the difference between inflows into α from other strategies and outflows from α to other strategies. For a fixed protocol ρ , (MD) can be viewed as a map from population games v to laws of motion on \mathcal{X} , as described in Section 2.2.²⁰

The prototype for this construction is, again, the replicator dynamics (RD). Three well-known protocols that generate (RD) are:

$$\rho_{\alpha\beta}(x, \pi) = x_\beta \pi_\beta, \tag{5.1a}$$

$$\rho_{\alpha\beta}(x, \pi) = -x_\beta \pi_\alpha, \tag{5.1b}$$

$$\rho_{\alpha\beta}(x, \pi) = x_\beta [\pi_\beta - \pi_\alpha]_+, \tag{5.1c}$$

with π assumed nonnegative in (5.1a) and nonpositive in (5.1b).²¹ The x_β appearing in the right-hand sides allows us to interpret (5.1a)–(5.1c) as imitative protocols, with a revising agent picking a candidate strategy by observing the choice and the payoff of a randomly chosen opponent. The protocols differ in how payoffs determine the rates at which switches are consummated. Protocols (5.1a) and (5.1b), due to Weibull (1995) and Björnerstedt and Weibull (1996), are respectively called *imitation of success* and *imitation driven by dissatisfaction*. In the former, imitation rates increase linearly in the opponent’s payoff; in the latter, imitation rates decrease linearly in the revising agent’s own payoff. Protocol (5.1c) is due to Helbing (1992) and Schlag (1998), and is called *pairwise proportional imitation*. Under (5.1c), a revising agent only considers switching if the opponent’s payoff is higher than their own, and then does so at a rate proportional to the payoff difference. Substituting any of these protocols into (MD) and rearranging yields the replicator dynamics (RD).

We now show that the revision protocols from this example can be generalized to cover wider ranges of Riemannian game dynamics, focusing again on interior population states:²²

Proposition 5.1. *Let g be an extendable Riemannian metric such that $g^\sharp(x)$ is nonnegative for all $x \in \mathcal{X}^\circ$. Then up to a change of speed, the following protocols generate (RmD) as their mean dynamics on \mathcal{X}° :*

$$\rho_{\alpha\beta}(x, \pi) = \frac{(g^\sharp(x)\mathbf{1})_\alpha}{x_\alpha} (\pi g^\sharp(x))_\beta, \tag{5.2a}$$

$$\rho_{\alpha\beta}(x, \pi) = -\frac{(\pi g^\sharp(x))_\alpha}{x_\alpha} (g^\sharp(x)\mathbf{1})_\beta, \tag{5.2b}$$

where π is assumed nonnegative in (5.2a) and nonpositive in (5.2b). If $g(x)$ is separable, the dynamics (RmD) are also generated (up to a change of speed) by the protocol

$$\rho_{\alpha\beta}(x, \pi) = \frac{g_{\alpha\alpha}^\sharp(x)}{x_\alpha} g_{\beta\beta}^\sharp(x) [\pi_\beta - \pi_\alpha]_+. \tag{5.2c}$$

²⁰ Solutions to (MD) may further be viewed as approximations to the sample paths of stochastic evolutionary models generated by the game \mathcal{G} and protocol ρ : for a comprehensive treatment, see Benaïm and Weibull (2003) and Roth and Sandholm (2013).

²¹ Since the replicator dynamics (and all Riemannian game dynamics) are invariant to equal shifts in all strategies’ payoffs, these assumptions about payoffs are innocuous.

²² Under minimal-rank extendible metrics, the result to follow also applies on the boundary. Handling boundary states under full-rank extendible metrics requires modifications of the sort described in Lahkar and Sandholm (2008), a direction we do not pursue here.

Proof. Substitute (5.2a)–(5.2c) with $\pi = v(x)$, $v^\sharp(x) = (v(x)g^\sharp(x))^\top$, and $n(x) = (\mathbb{1}g^\sharp(x))^\top$ into (MD) to obtain

$$\dot{x} = v^\sharp(x) \sum_{\beta \in \mathcal{A}} n_\beta(x) - n(x) \sum_{\beta \in \mathcal{A}} v_\beta^\sharp(x). \tag{5.3}$$

Changing the speed at state x by dividing the right-hand side of (5.3) by $s(x) = \sum_\beta n_\beta(x)$ yields form (3.33a) of (RmD). \square

After a change of speed, the Riemannian dynamics (RmD) take the symmetric form (5.3). This symmetry is a source of the appealing properties of the dynamics established below. By contrast, the random assignment of revision opportunities implies that, under the mean dynamics (MD), the outflow rate from each strategy α to other strategies is proportional to the popularity x_α of the original strategy, resulting in an expression that is not symmetric. The factor x_α appearing in the denominators in (5.2) also lets us recover the symmetric expression (5.3) from (MD).

The asymmetric treatment of current and candidate strategies under (5.2) is illustrated by our running examples:

Example 5.1 (*p-replicator dynamics*). Since p -replicator dynamics are generated by the Riemannian metric $g(x) = \text{diag}(1/x_1^p, \dots, 1/x_n^p)$, (5.2c) implies that these dynamics are induced by the revision protocols

$$\rho_{\alpha\beta}(x, \pi) = x_\alpha^{p-1} x_\beta^p [\pi_\beta - \pi_\alpha]_+. \tag{5.4}$$

When $p = 1$, we have $x_\alpha^{p-1} = 1$ and $x_\beta^p = x_\beta$, so (5.4) boils down to the pairwise proportional imitation protocol (5.2c) and induces the replicator dynamics (RD). When $p = 0$, we have $x_\alpha^{p-1} = x_\alpha^{-1}$ and $x_\beta^p = 1$, so (5.4) gives

$$\rho_{\alpha\beta}(x, \pi) = x_\alpha^{-1} [\pi_\beta - \pi_\alpha]_+, \tag{5.5}$$

and induces the projection dynamics (PD) on \mathcal{X}° . Protocol (5.5) was introduced by Lahkar and Sandholm (2008), who interpret it as a model of “revision driven by insecurity”: agents playing rare strategies are particularly likely to consider revising, while candidate strategies are chosen without regard for their current levels of use.

While the revision protocols (5.2) are capable of generating many Riemannian dynamics (RmD), one can sometimes construct simpler protocols that take advantage of the structure of smaller classes of Riemannian dynamics. For the microfoundations of the nested replicator dynamics (NRD) and extensions thereof, we refer the reader to Mertikopoulos and Sandholm (2018).

6. General properties

In this section, we derive some general results for (RmD). In Section 6.1 we state a basic but technically challenging result on the existence and uniqueness of solutions. In Section 6.2 we show that the dynamics exhibit positive correlation with the game’s payoffs, and we characterize the dynamics’ rest points as either restricted equilibria or Nash equilibria. In Section 6.3 we study the global behavior of the dynamics in potential games. Finally, in Section 6.4 we establish the asymptotic stability of regular interior ESSs under (RmD).

6.1. *Existence and uniqueness of solutions*

To illustrate the possibilities for existence and uniqueness of solutions, it is useful to start with a simple example. Specifically, consider the p -replicator dynamics of Example 4.3 for a 2-strategy game with action set $\mathcal{A} = \{1, 2\}$ and payoff functions $v_1(x) = 1, v_2(x) = 0$.

When $p = 1$, we obtain the toy replicator equation

$$\dot{x}_1 = x_1(1 - x_1). \tag{6.1}$$

Solutions to this equation exist and are unique for all $t \in (-\infty, \infty)$, and the support of $x(t)$ is invariant. The pure states 0 and 1 are both rest points, and it is easy to check that the unique solution with initial condition $x_1(0) = a \in (0, 1)$ is $x_1(t) = a/[a + (1 - a)e^{-t}]$.

When $p = 0$, we obtain the Euclidean projection dynamics

$$\dot{x}_1 = \begin{cases} 1/2 & \text{if } x_1 < 1 \\ 0 & \text{if } x_1 = 1. \end{cases} \tag{6.2}$$

For every initial condition $x_1(0) \in [0, 1]$, this equation admits the unique forward solution $x_1(t) = x_1(0) + t/2$ for $t \in [0, 2(1 - x_1(0))]$ and $x_1(t) = 1$ thereafter. Evidently, the support of $x(t)$ is not invariant; also, backward solutions are not defined for all time, and solutions are not smooth in t when $x_1 = 1$ is reached.

Finally, when $p = 1/2$, we obtain the differential equation

$$\dot{x}_1 = \frac{\sqrt{x_1(1 - x_1)}}{\sqrt{x_1} + \sqrt{1 - x_1}}. \tag{6.3}$$

Although this equation admits forward (and backward) solutions from every initial condition, these are no longer unique. Starting at $x_1(0) = 0$, we have the stationary solution $x_1(t) = 0$ for $t \in [0, \infty)$; furthermore, one can verify by a direct – albeit tedious – calculation that there is another solution, namely $x_1(t) = \frac{1}{2} + \frac{t-2}{4}\sqrt{1 + t - t^2/4}$ for $t \in [0, 4)$ and $x_1(t) = 1$ thereafter. Additional solutions may linger at $x_1 = 0$ before emulating the previous solution trajectory.

The differences in behavior in the three cases above can be traced back to the properties of the underlying Riemannian metrics. First, the replicator dynamics are generated by the Shahshahani metric, which is minimal-rank extendable to all of \mathcal{X} . In this case the induced dynamics (RmD) are Lipschitz continuous, so existence and uniqueness of solutions is guaranteed by the Picard–Lindelöf theorem (along with an argument to account for \mathcal{X} being closed). Moreover, the support of $x(t)$ is constant, and solutions exist in both forward and backward time (Sandholm, 2010b, Theorems 4.A.5 and 5.4.7).

On the other hand, the Euclidean projection dynamics (PD) are generated by a full-rank extendable metric. In such cases, the induced dynamics (RmD) are typically discontinuous, so the relevant solution notion is that of a *Carathéodory solution*, an absolutely continuous trajectory that satisfies (RmD) for almost all $t \geq 0$. In the case of (PD), Lahkar and Sandholm (2008) showed that every initial condition admits a unique Carathéodory forward solution, but that different solution orbits can merge in finite time, as illustrated in the previous example and in Fig. 3(a).

The following proposition shows that this behavior of (RD) and (PD) is representative of the minimal-rank and full-rank extendable cases respectively:

Proposition 6.1. *Let g be an extendable Riemannian metric.*

- (i) *If g is minimal-rank extendable, (RmD) admits a unique global solution from every initial condition in \mathcal{X} ; moreover, each solution has constant support.*
- (ii) *If g is full-rank extendable, (RmD) admits a unique forward Carathéodory solution from every initial condition in \mathcal{X} .*

Proposition 6.1 justifies the terminology *continuous* and *discontinuous* that we introduced in Section 3.5 to refer to dynamics induced by minimal-rank and full-rank metrics. The nontrivial part of Proposition 6.1 is the proof of part (ii): despite an apparent similarity, this result is considerably harder than the corresponding result of Lahkar and Sandholm (2008) for (PD), so we relegate its proof to Appendix D. The main reason for this difficulty is that known uniqueness proofs for projected differential equations depend crucially on the Riemannian metric being constant throughout the dynamics’ state space, an assumption that obviously fails here.

Of course, as can be seen from the continuous – but not Lipschitz continuous – system (6.3), (RmD) may fail to admit unique solutions from initial conditions at the boundary of \mathcal{X} if the underlying metric does not admit a Lipschitz continuous extension to the boundary of \mathcal{X} . To avoid the resulting complications, we do not consider dynamics that are continuous but not Lipschitz continuous in the rest of the paper.

6.2. Basic properties

We now establish some basic relationships between (RmD) and the payoffs of the underlying game. We first show that (RmD) respects positive correlation:

Proposition 6.2. *The dynamics (RmD) satisfy (PC).*

Proof. Let $V(x) = \Pi_x(v^\sharp(x))$. We then claim that

$$\langle v(x)|V(x) \rangle = \langle v^\sharp(x), V(x) \rangle_x \geq \langle \Pi_x(v^\sharp(x)), V(x) \rangle_x = \|V(x)\|_x^2 \geq 0, \tag{6.4}$$

with equality if and only if $V(x) = 0$. The only step in (6.4) needing justification is the first inequality. For this step, we split the analysis into three cases. First, if $x \in \mathcal{X}^\circ$, the inequality binds because Π_x orthogonally projects \mathbb{R}^A onto $\mathbb{R}_0^A = \text{Adm}_g(x)$, which contains $V(x)$. Second, if $x \in \text{bd}(\mathcal{X})$ and g is minimal-rank extendable, then $v^\sharp(x) \in \mathbb{R}^{\text{supp}(x)}$, so the inequality binds because Π_x projects $\mathbb{R}^{\text{supp}(x)}$ orthogonally onto $\mathbb{R}_0^A \cap \mathbb{R}^{\text{supp}(x)} = \text{Adm}_g(x)$, which contains $V(x)$. Finally, if $x \in \text{bd}(\mathcal{X})$ and g is full-rank extendable, Π_x is the closest point projection of \mathbb{R}^A onto the tangent cone $\text{TC}_{\mathcal{X}}(x)$. Hence, by Moreau’s decomposition theorem (Hiriart-Urruty and Lemaréchal, 2001), we infer that $v^\sharp(x) - \Pi_x(v^\sharp(x))$ lies in the normal cone

$$\text{NC}_{\mathcal{X}}(x) = \{w \in \mathbb{R}^A : \langle w, z \rangle_x \leq 0 \text{ for all } z \in \text{TC}_{\mathcal{X}}(x)\}. \tag{6.5}$$

Since $V(x) \in \text{Adm}_g(x) = \text{TC}_{\mathcal{X}}(x)$, the first inequality in (6.4) is immediate. \square

Proposition 6.2 is not particularly surprising: after all, the basic postulate behind (RmD) is that the dynamics’ vector of motion is the closest feasible approximation to the game’s payoff field, with the notion of closeness determined by the underlying Riemannian metric (or, equivalently, cost function). As we show below, this alignment can be exploited further to characterize the dynamics’ rest points.

To that end, recall that one of the main attributes of the Euclidean projection dynamics (PD) is *Nash stationarity*:

$$x^* \in \mathcal{X} \text{ is a rest point if and only if it is a Nash equilibrium.} \tag{NS}$$

This property does not hold under the replicator dynamics: for instance, every pure state of \mathcal{X} is stationary under (RD). In this case, (NS) is replaced by the notion of *restricted stationarity*:²³

$$x^* \in \mathcal{X} \text{ is a rest point if and only if it is a restricted equilibrium.} \tag{RS}$$

Our next result shows that this difference between the projection and the replicator dynamics is representative of the discontinuous and continuous cases, and highlights one advantage of the former over the latter:

Proposition 6.3.

- (i) *Continuous Riemannian dynamics satisfy (RS).*
- (ii) *Discontinuous Riemannian dynamics satisfy (NS).*

Proof. For (i), recall that the coordinate expression (3.33) for (RmD) always holds when g is minimal-rank extendable, and $\dot{x}_\alpha = 0$ whenever $x_\alpha = 0$. Therefore, it suffices to check that $x^* \in \mathcal{X}^\circ$ is a rest point if and only if all the components of $v(x^*)$ are equal. To that end, note that x^* is a rest point of (3.33) if and only if

$$v^\sharp(x^*) = \frac{\sum_\gamma v_\gamma^\sharp(x^*)}{\sum_\gamma n_\gamma(x^*)} n(x^*) \propto n(x^*). \tag{6.6}$$

In turn, this means that x^* is a rest point of (RmD) if and only if $v^\sharp(x^*) \propto n(x^*)$; our claim then follows from the fact that $g^\sharp(x^*)$ is invertible.

For (ii), assume that g is full-rank extendable and fix some $x^* \in \mathcal{X}$. It is easy to show that x^* is a Nash equilibrium if and only if it satisfies the variational characterization

$$0 \leq \langle v(x^*) | x - x^* \rangle = \langle v^\sharp(x^*), x - x^* \rangle_{x^*} \quad \text{for all } x \in \mathcal{X}, \tag{6.7}$$

which says that $v^\sharp(x^*)$ lies in the normal cone $\text{NC}_{\mathcal{X}}(x^*)$ of \mathcal{X} at x^* (cf. Eq. (6.5) above). Moreau’s decomposition theorem then yields $v^\sharp(x^*) \in \text{NC}_{\mathcal{X}}(x^*)$ if and only if $\Pi_{x^*}(v^\sharp(x^*)) = 0$, so our assertion follows. \square

Remark 6.1. We note without proof that shifting all strategies’ payoffs by the same amount has no effect on (RmD), and rescaling all strategies’ payoffs by the same factor only changes the speed at which solution paths are traversed. In addition, on the face of \mathcal{X} spanned by a subset \mathcal{A}' of \mathcal{A} , continuous dynamics are invariant to changes in the payoffs of strategies outside of \mathcal{A}' .

6.3. *Global convergence in potential games*

Recall that $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$ is a potential game if $v_\alpha(x) = \partial_\alpha f(x)$ for some potential function $f: \mathcal{X} \rightarrow \mathbb{R}$ (cf. Example 2.2). It then follows from Proposition 6.2 that f is a *strict global Lya-*

²³ Recall here that x^* is a *restricted equilibrium* if all strategies in its support earn equal payoffs.

punov function for (RmD), meaning that its value increases along (RmD) whenever the dynamics are not at rest.²⁴

For continuous Riemannian dynamics, a standard Lyapunov argument implies that all ω -limit points of (RmD) are rest points – and hence, by Proposition 6.3, restricted equilibria of \mathcal{G} . However, this argument does not extend to discontinuous dynamics and Nash equilibria because it requires continuity of solutions with respect to initial conditions, a requirement which is difficult to prove in our case. To circumvent this obstacle, we establish a lower semi-continuous (l.s.c.) bound on the rate of change of the game’s potential function. This bound then allows us to apply Proposition C.1 in Appendix C, which shows that, for dynamics on a compact set, such a bound on the rate of change of a Lyapunov function guarantees global convergence.

Proposition 6.4. *Let \mathcal{G} be a potential game with potential function f . Then, f is a strict Lyapunov function for (RmD) and every ω -limit point of (RmD) is a rest point of (RmD). These are restricted equilibria if (RmD) is continuous, and Nash equilibria if (RmD) is discontinuous.*

Proof. Let $V(x) = \Pi_x(v^\sharp(x))$ and let $x(t)$ be a solution of (RmD). Then, Proposition 6.2 yields

$$\frac{d}{dt} f(x(t)) = \langle Df(x(t)) | \dot{x}(t) \rangle = \langle v(x(t)) | V(x(t)) \rangle \geq 0, \tag{6.8}$$

with equality if and only if $V(x(t)) = 0$. Hence, f is a strict global Lyapunov function for (RmD).

When (RmD) is (Lipschitz) continuous, a standard argument shows that every ω -limit point of (RmD) is a rest point thereof (see e.g., Sandholm, 2010b, Theorem 7.B.3). The discontinuous case however requires a different treatment. To start, note that

$$\frac{d}{dt} f(x(t)) = \langle v(x(t)) | V(x(t)) \rangle = \langle v^\sharp(x), \Pi_x(v^\sharp(x)) \rangle_x \geq \|V(x(t))\|_x^2 \geq 0, \tag{6.9}$$

where the first inequality follows from Moreau’s decomposition theorem. Both inequalities bind if and only if $V(x(t)) = 0$; since the speed function $x \mapsto \|V(x)\|_x$ is lower semi-continuous (cf. Lemma C.2), Proposition C.1 shows that every ω -limit point of (RmD) is a rest point. \square

The classic analyses of Kimura (1958) and Shahshahani (1979) showed that in common interest games, average payoffs are increased at a maximal rate under the replicator dynamics, provided that “maximal” is defined with respect to the Shahshahani metric. We conclude this section by deriving an analogous principle for all Riemannian game dynamics. To state it, define the *gradient* of a smooth function $f : \mathcal{K}^\circ \rightarrow \mathbb{R}$ with respect to g by

$$\text{grad } f(x) = (Df(x) g^{-1}(x))^\top, \tag{6.10}$$

that is, as the (necessarily unique) vector satisfying

$$\langle Df(x) | z \rangle = \langle \text{grad } f(x), z \rangle_x \quad \text{for all } z \in \mathbb{R}^{\mathcal{A}}, x \in \mathcal{K}^\circ. \tag{6.11}$$

Geometrically, the vector $\text{grad } f(x)$ represents the direction of maximal increase of the function f at x with respect to the metric g .²⁵ We then have:

²⁴ Definitions concerning stability and convergence are collected in Appendix C.

²⁵ Specifically, this means that $\text{grad } f(x) = \arg \max\{D_z f(x) : \|z\|_x = 1\}$; that this is so follows from the definition of $\text{grad } f(x)$ and the Cauchy–Schwarz inequality.

Proposition 6.5. *Let \mathcal{G} be a potential game with potential function f and let g be an extendable Riemannian metric. Then, for all $x \in \mathcal{X}^\circ$, the vector field that defines (RmD) is the projection of $\text{grad } f$ onto $T_{\mathcal{X}}(x)$ with respect to g .*

Proof. Since $v(x) = Df(x)$, we have $\Pi_x(v^\sharp(x)) = \Pi_x(\text{grad } f(x))$, as claimed. \square

Hence, at interior states, the dynamics (RmD) increase the value of potential at a maximal rate under the geometry defined by g , subject to feasibility. For discontinuous dynamics, this conclusion remains true even at boundary states. For continuous dynamics, the interior of each face of \mathcal{X} is invariant under (RmD), so this conclusion holds provided that feasibility is understood to incorporate this additional constraint.

6.4. Local stability of regular interior ESSs

To state the final general property of (RmD), we recall that an interior population state $x^* \in \mathcal{X}^\circ$ is a regular evolutionarily stable state (or regular ESS) of game $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$ (Maynard Smith and Price, 1973; Taylor and Jonker, 1978) if x^* is a Nash equilibrium of \mathcal{G} and if the derivative of v at x^* exists and is negative definite with respect to $\mathbb{R}_0^{\mathcal{A}}$.²⁶ The following proposition on the stability of such states is essentially due to Hopkins (1999), who builds on a result of Taylor and Jonker (1978) for the replicator dynamics.

Proposition 6.6. *Let x^* be a regular interior ESS of \mathcal{G} . Then x^* is asymptotically stable under (RmD).*

Proof. Using (3.33b) (or see Proposition A.3), we can write (RmD) in the form

$$V(x) = M(x)v(x)^\top \tag{6.12}$$

for $x \in \mathcal{X}^\circ$, where $M(x)$ is symmetric, positive definite with respect to $\mathbb{R}_0^{\mathcal{A}}$, and maps $\mathbf{1}$ to 0. In fact, (3.33b) implies that these properties of $M(x)$ are satisfied whenever $x \in \mathcal{K}^\circ$.

To streamline the computation to follow, extend the C^1 function v from a neighborhood of x^* in \mathcal{X}° to a neighborhood of x^* in \mathcal{K}° (for instance, by setting $v(x + c\mathbf{1}) = v(x)$). Then using the chain rule, we obtain

$$\frac{\partial V_\alpha}{\partial x_\beta}(x^*) = \sum_\gamma \frac{\partial M_{\alpha\gamma}}{\partial x_\beta}(x^*)v_\gamma(x^*) + \sum_\gamma M_{\alpha\gamma}(x^*) \frac{\partial v_\gamma}{\partial x_\beta}(x^*). \tag{6.13}$$

Since x^* is an interior Nash equilibrium, $v(x^*)$ is a constant covector, and the fact that $M(x)$ maps $\mathbf{1}$ to 0 implies that $\sum_\gamma M_{\alpha\gamma}(x) \equiv 0$, and hence that $\sum_\gamma \frac{\partial M_{\alpha\gamma}}{\partial x_\beta}(x^*) = 0$. Thus the initial sum in (6.13) is 0, and we may rewrite what remains of (6.13) in matrix form as

$$DV(x^*) = M(x^*)Dv(x^*)^\top, \tag{6.14}$$

where $Dv(x^*)^\top \equiv (\frac{\partial v_\alpha}{\partial x_\beta}(x^*))_{\alpha, \beta \in \mathcal{A}}$.

By the definition of regular interior ESS, the matrix $Dv(x^*)^\top$ is negative definite with respect to $\mathbb{R}_0^{\mathcal{A}}$. Combining this fact and the previously mentioned properties of $M(x^*)$ with Hines’s

²⁶ The last condition requires that $\sum_\alpha z_\alpha v'_\alpha(x^*; z) < 0$ for all $z \in \mathbb{R}_0^{\mathcal{A}} \setminus \{0\}$, where $v'_\alpha(x^*; z)$ denotes the directional derivative of v_α . An equivalent condition is used in the proof below.

lemma (Hines, 1980; see also Sandholm, 2010b, Lemma 8.5.1), we conclude from (6.14) that the eigenvalues of $DV(x^*)$ corresponding to eigenvectors in \mathbb{R}_0^A have negative real part. Thus x^* is asymptotically stable under (6.12). \square

Remark 6.2. When the Riemannian metric $g(x)$ is separable (Example 4.4), the asymptotic stability of regular boundary ESSs can be established using a more elaborate linearization argument; for the case of the replicator dynamics, see Taylor and Jonker (1978) or Sandholm (2010b, Theorem 8.5.5). But separable metrics are instances of Hessian–Riemannian metrics (Example 4.6), and we will prove a general asymptotic stability result for (not necessarily regular) ESSs under Hessian game dynamics in the next section.

7. Hessian game dynamics

By virtue of the integrability property that defines them, potential games have desirable convergence properties under a wide range of evolutionary dynamics. By contrast, convergence results for other classes of games – for instance, contractive games – require additional structure, often taking the form of integrability properties built into the dynamics themselves.²⁷

In this section, we show that the integrability of Hessian Riemannian metrics allows us to generalize several properties of the replicator dynamics and the Euclidean projection dynamics to a substantially broader class of dynamics. These *Hessian game dynamics*, introduced in Example 4.6, take the form

$$\dot{x} = \arg \max_{z \in \text{Adm}_g(x)} \left[\langle v(x) | z \rangle - \frac{1}{2} \|z\|_x^2 \right], \quad g = \text{Hess } h, \tag{HD}$$

where the continuous function $h: \mathcal{K} \rightarrow \mathbb{R}$ is C^3 -smooth on every positive suborthant of \mathcal{K} , and where $\text{Hess } h(x)$ is positive definite for all $x \in \mathcal{K}^\circ$. As we demonstrate below, the integrability built into the dynamics (HD) is the source of a variety of stability and convergence results.

A key element of our analysis is the so-called *Bregman divergence*, which we introduce in Section 7.1. In Section 7.2, we establish global convergence to equilibrium in contractive games and local stability of evolutionarily stable states (ESSs), while Section 7.3 demonstrates the convergence of time-averages of interior trajectories to Nash equilibrium and provides sufficient conditions for permanence. Finally, Section 7.4 establishes the elimination of strictly dominated strategies under continuous Hessian dynamics.

7.1. Bregman divergences

When used as a tool for establishing convergence, Lyapunov functions typically measure some sort of “distance” between the current state and a target state x^* . For Hessian dynamics, a natural point of departure is the potential function h of the metric $g = \text{Hess } h$ that defines them. However, since the (game-specific) target state x^* is independent of g , there is no reason that h itself should serve as a Lyapunov function. Instead, taking advantage of the convexity of h , we consider the difference between $h(x^*)$ and the best linear approximation of $h(x^*)$ from x .

Formally, the *Bregman divergence* of h (Bregman, 1967) is defined as

$$D_h(x^*, x) = h(x^*) - h(x) - h'(x; x^* - x), \quad x^*, x \in \mathcal{X}, \tag{7.1}$$

²⁷ See Hofbauer and Sandholm (2007), Sandholm (2010a), and Zusai (2018).

where $h'(x; x^* - x)$ is the one-sided derivative of h at x along $x^* - x$, i.e.,

$$h'(x; x^* - x) = \lim_{t \rightarrow 0^+} t^{-1} [h(x + t(x^* - x)) - h(x)]. \tag{7.2}$$

Since h is convex, we have

$$D_h(x^*, x) \geq 0, \text{ with equality if and only if } x^* = x. \tag{7.3}$$

On the other hand, D_h is not symmetric in x^* and x , so it is not a bona fide distance function on \mathcal{X} ; rather, $D_h(x^*, x)$ describes the remoteness of x from the base point x^* , hence the name “divergence”.

Revisiting our two archetypal examples, the Euclidean metric is generated by the quadratic potential $h(x) = \frac{1}{2} \sum_{\alpha} x_{\alpha}^2$. Definition (7.1) then yields the *Euclidean divergence*

$$D_{\text{Eucl}}(x^*, x) = \frac{1}{2} \sum_{\alpha} (x_{\alpha} - x_{\alpha}^*)^2, \tag{7.4a}$$

which is (uncharacteristically) symmetric in x^* and x . Analogously, the Shahshahani metric is generated by the (negative) entropy $h(x) = \sum_{\alpha} x_{\alpha} \log x_{\alpha}$. A short calculation shows that the corresponding divergence function is the *Kullback–Leibler (KL) divergence*

$$D_{\text{KL}}(x^*, x) = \sum_{\alpha: x_{\alpha}^* > 0} x_{\alpha}^* \log(x_{\alpha}^*/x_{\alpha}), \tag{7.4b}$$

which has been used extensively in the analysis of the replicator dynamics (Hofbauer and Sigmund, 1998; Weibull, 1995).

The key qualitative difference between the Euclidean divergence (7.4a) and the KL divergence (7.4b) is that the former is finite for all $x, x^* \in \mathcal{X}$, whereas the latter blows up to $+\infty$ when $\text{supp}(x^*) \not\subseteq \text{supp}(x)$. The reason for this blow-up is that the entropy function $h(x) = \sum_{\alpha} x_{\alpha} \log x_{\alpha}$ becomes infinitely steep as any boundary point x of \mathcal{X} is approached from the interior of \mathcal{X} , i.e.,

$$\sup_{\alpha \in A} |\partial_{\alpha} h(x_n)| \rightarrow \infty \text{ for every interior sequence } x_n \text{ converging to } x. \tag{7.5}$$

When this is the case for all $x \in \text{bd}(\mathcal{X})$, we say that h is *steep* (Alvarez et al., 2004; Hofbauer and Sandholm, 2002). At the opposite end of the spectrum, if $Dh(x)$ exists for all $x \in \mathcal{X}$, we say that h is *nonsteep*.

The link between the steepness of h and the finiteness of the associated Bregman divergence is provided by the following lemma:

Lemma 7.1. *Fix $x^* \in \mathcal{X}$ and let $\mathcal{D}(x^*)$ denote the union of the relative interiors of the faces of \mathcal{X} that contain x^* , i.e.,*

$$\mathcal{D}(x^*) \equiv \{x \in \mathcal{X} : \text{supp}(x^*) \subseteq \text{supp}(x)\}. \tag{7.6}$$

If h is steep, we have $D_h(x^, x) < \infty$ for all $x \in \mathcal{D}(x^*)$; by contrast, if h is nonsteep, we have $D_h(x^*, x) < \infty$ for all $x \in \mathcal{X}$.*

Proof. If h is steep and $x \in \mathcal{D}(x^*)$, the smoothness of h on the face of \mathcal{X} spanned by $\text{supp}(x) \supseteq \text{supp}(x^*)$ implies that the directional derivative $h'(x; x^* - x)$ exists and is finite, so $D_h(x^*, x)$

is itself finite. If instead h is nonsteep, $h'(x; x^* - x)$ exists and is finite for all $x \in \mathcal{X}$, so again $D_h(x^*, x) < \infty$. \square

Beyond the positive definiteness property (7.3), the attribute of the Bregman divergence that recommends it as a Lyapunov function for (HD) is that the level sets of $D_h(x^*, \cdot)$ are perpendicular to all rays emanating from x^* under $g = \text{Hess } h$. Formally, we have:

Lemma 7.2. *Let $g = \text{Hess } h$ be an extendable HR metric and let $x^* \in \mathcal{X}$. Then, for every smooth curve $x(t)$ with constant support containing that of x^* , we have:*

$$\frac{d}{dt} D_h(x^*, x(t)) = \langle \dot{x}(t), x(t) - x^* \rangle_{x(t)}. \tag{7.7}$$

In particular, if $D_h(x^*, x(t))$ is constant, $\dot{x}(t)$ is perpendicular to $x(t) - x^*$. Finally, if h is nonsteep, the above conclusions hold for every smooth curve $x(t)$ on \mathcal{X} .

Proof. The proof is a direct application of the chain rule:

$$\begin{aligned} \frac{d}{dt} D_h(x^*, x) &= - \sum_{\alpha} \left[\frac{\partial h}{\partial x_{\alpha}} \dot{x}_{\alpha} + \frac{\partial h}{\partial x_{\alpha}} \frac{d}{dt} (x_{\alpha}^* - x_{\alpha}) + \sum_{\beta} \frac{\partial^2 h}{\partial x_{\alpha} \partial x_{\beta}} (x_{\alpha}^* - x_{\alpha}) \dot{x}_{\beta} \right] \\ &= \sum_{\alpha, \beta} (x_{\alpha}^* - x_{\alpha}) g_{\alpha\beta}(x) \dot{x}_{\beta} = \langle \dot{x}, x - x^* \rangle_x, \end{aligned} \tag{7.8}$$

where all summations are taken over the (constant) support $\mathcal{A}' \equiv \text{supp}(x(t))$ of $x(t)$ and we used the fact that $\dot{x}_{\alpha} = 0$ for $\alpha \notin \mathcal{A}'$. Finally, in the nonsteep case, h is smooth throughout \mathcal{X} , so the above holds for every smooth curve $x(t)$. \square

Within the class of Hessian Riemannian metrics, steepness of h roughly corresponds to minimal-rank extendability of the metric $g = \text{Hess } h$, and nonsteepness to full-rank extendability. These analogies fail when the steepness of h does not adequately control the regularity of g near the boundary of \mathcal{X} , or when g is minimal-rank extendable but generates non-Lipschitz dynamics. Bearing this in mind, we use the term *continuous Hessian dynamics* for Riemannian dynamics generated by a minimal-rank extendable metric $g = \text{Hess } h$ with steep h , and the term *discontinuous Hessian dynamics* for Riemannian dynamics generated by a full-rank extendable metric $g = \text{Hess } h$ with nonsteep h . In what follows, we will tacitly assume that the dynamics (HD) are either continuous or discontinuous.

7.2. Contractive games and evolutionarily stable states

Recall that a population game $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$ is called *contractive* if $\langle v(x') - v(x) | x' - x \rangle \leq 0$ for all $x, x' \in \mathcal{X}$, *strictly contractive* if the inequality is strict whenever $x \neq x'$, and *conservative* if the inequality always binds (cf. Example 2.3). As is well known, the set of Nash equilibria of any contractive game is convex, and every strictly contractive game admits a unique Nash equilibrium (Hofbauer and Sandholm, 2009).

Combining the defining inequality of strictly contractive games with the variational characterization of Nash equilibria (6.7), it follows that the (necessarily unique) Nash equilibrium of a strictly contractive game satisfies the inequality

$$\langle v(x) | x - x^* \rangle \leq 0 \quad \text{with equality only if } x = x^*. \tag{7.9}$$

Hofbauer and Sandholm (2009) call a state satisfying (7.9) a *globally evolutionarily stable state* (GESS). This is the global version of the seminal local solution concept of Maynard Smith and Price (1973): if (7.9) holds for all $x \neq x^*$ in a neighborhood of x^* , then x^* is called an *evolutionarily stable state* (ESS).²⁸

It is well known that the GESS x^* of a strictly contractive game attracts all solutions of the replicator dynamics whose initial support contains that of x^* ; by comparison, x^* attracts all orbits of the Euclidean projection dynamics (PD). Theorem 7.3 extends these results to all Hessian dynamics (HD).

Theorem 7.3. *Let x^* be the (necessarily unique) Nash equilibrium of a strictly contractive game \mathcal{G} . Then:*

- (i) *For all continuous Hessian dynamics, $D_h(x^*, \cdot)$ is a strict decreasing Lyapunov function on $\mathcal{D}(x^*)$, and x^* is asymptotically stable with basin $\mathcal{D}(x^*)$.*
- (ii) *For all discontinuous Hessian dynamics, $D_h(x^*, \cdot)$ is a strict decreasing global Lyapunov function, and x^* is globally asymptotically stable.*

Proof. We begin with the continuous case. By Proposition 6.1(i), every solution $x(t)$ of (HD) has constant support. Hence, if $x(0) \in \mathcal{D}(x^*)$, Lemma 7.2 yields:

$$\begin{aligned} \frac{d}{dt} D_h(x^*, x) &= \langle \dot{x}, x - x^* \rangle_x = \langle \Pi_x(v^\sharp(x)), x - x^* \rangle_x \\ &= \langle v^\sharp(x), x - x^* \rangle_x = \langle v(x) | x - x^* \rangle, \end{aligned} \tag{7.10a}$$

$$\leq 0, \tag{7.10b}$$

where we used the definition of Π_x for minimal-rank extendable metrics (cf. Section 3.6) to obtain (7.10a) and the definition (7.9) of a GESS for (7.10b). Since equality in (7.10) holds if and only if $x = x^*$, we conclude that $D_h(x^*, x)$ is a strict Lyapunov function on $\mathcal{D}(x^*)$. If we can show in addition that $x(t)$ has no ω -limit points in $\mathcal{X} \setminus \mathcal{D}(x^*)$, then asymptotic stability with basin $\mathcal{D}(x^*)$ follows from standard arguments (see e.g., Sandholm, 2010b, Theorem 7.B.3).

Assume therefore that $x(t)$ admits an ω -limit point $x^\omega \neq x^*$, so $x(t_n) \rightarrow x^\omega$ for some sequence of times $t_n \uparrow \infty$. Since $|\dot{x}_\alpha(t)|$ is bounded from above by $V_{\max} \equiv \sup_{x \in \mathcal{X}} \max_\beta |V_\beta(x)| < \infty$, there exists an open neighborhood U of x^ω and positive $a, \delta, n_0 > 0$ such that $x(t) \in U$ and $\langle v(x(t)) | x(t) - x^* \rangle \leq -a < 0$ for all $t \in [t_n, t_n + \delta]$ and all $n \geq n_0$. Hence, by (7.10), we get

$$D_h(x^*, x(t_n + \delta)) - D_h(x^*, x(0)) \leq \int_0^{t_n + \delta} \langle v(x(s)) | x(s) - x^* \rangle ds \leq -a(n - n_0)\delta. \tag{7.11}$$

We thus get $\liminf_{t \rightarrow \infty} D_h(x^*, x(t)) = -\infty$, a contradiction.

For the discontinuous case, note first that since $x^* - x \in \text{TC}_{\mathcal{X}}(x)$, Moreau’s decomposition theorem implies that $\langle \Pi_x(v^\sharp(x)), x - x^* \rangle_x \leq \langle v^\sharp(x), x - x^* \rangle_x$. Thus, replacing the first equality in (7.10a) by an inequality, (7.10) shows that $D_h(x^*, \cdot)$ is a strict global Lyapunov function for (HD). Global asymptotic stability then follows from Proposition C.1. \square

²⁸ This concise characterization of evolutionary stability is due to Hofbauer et al. (1979).

The only implication of \mathcal{G} being strictly contractive used in the previous proof is that its Nash equilibrium is a GESS. More generally, if a game admits an ESS x^* , applying the above arguments in a neighborhood of x^* defined by a level set of $D_h(x^*, \cdot)$ yields the following result:

Theorem 7.4. *Evolutionarily stable states are asymptotically stable under (HD).*

Hofbauer and Sigmund (1990), Hopkins (1999), and Harper (2011) all offer results on the local stability of interior evolutionarily stable states under Riemannian game dynamics. Hofbauer and Sigmund (1990) show that under all Riemannian (not necessarily Hessian) dynamics (RmD), the function $L(x) = \langle x - x^*, x - x^* \rangle_{x^*}$ is a strict local Lyapunov function for interior ESSs x^* , implying that x^* is asymptotically stable. As noted in Section 6.4, Hopkins (1999) uses linearization to establish local stability of regular interior ESSs under (RmD). Finally, Harper (2011) employs a version of the argument above to prove asymptotic stability of ESSs for separable Hessian dynamics of the form (4.4).

An important case of contractive games that do not admit an ESS is the class of conservative games, which include population games generated by matching in symmetric zero-sum games. Under the replicator dynamics, the KL divergence does not provide a strict Lyapunov function for conservative games, but rather a constant of motion. The following result extends this conclusion to all Hessian dynamics:

Proposition 7.5. *Let x^* be a Nash equilibrium of a conservative game \mathcal{G} . Then, $D_h(x^*, \cdot)$ is a constant of motion along any interior solution segment of (HD).*

Proof. Simply note that (7.10) binds if \mathcal{G} is conservative and $x \in \mathcal{X}^\circ$. \square

Remark 7.1. In the definition of (HD), we required that h be finite throughout \mathcal{X} . This requirement is unnecessary for the preceding results when x^* is interior; however, if x^* lies on the boundary of \mathcal{X} , the proofs of Theorems 7.3 and 7.4 do not go through because $D_h(x^*, \cdot)$ is no longer well-defined throughout $\mathcal{D}(x^*)$. Nevertheless, the results themselves remain true if $g = \text{Hess} h$ is separable, allowing us to handle the p -replicator dynamics for $p \geq 2$ (Example 4.3). To prove this, it suffices to replace the implicit summation over all strategies in $h'(x; x^* - x)$ with a sum extending over only the strategies that lie in the support of x^* .

7.3. Convergence of time-averaged trajectories and permanence

We now extend two classic results for the replicator dynamics in random matching games (Example 2.1) to Hessian dynamics. The results for these games take advantage of the linearity of payoffs $v_\alpha(x) = \sum_{\beta \in \mathcal{A}} A_{\alpha\beta} x_\beta$ in the population state.

The first such result states that if a solution $x(t)$ of the replicator dynamics stays a positive distance away from the boundary of the simplex, then the time-averaged orbit $\bar{x}(t) = t^{-1} \int_0^t x(s) ds$ converges to the set of Nash equilibria of the underlying game (Schuster et al., 1981). The class of games to which this result applies includes zero-sum games (cf. Proposition 7.5) and games satisfying sufficient conditions for permanence (cf. Proposition 7.7 below). The following proposition shows that this convergence property extends to all Hessian dynamics:

Proposition 7.6. *Let \mathcal{G} be a random matching game and let $x(t)$ be a solution orbit of (HD). If $x(t)$ is contained in a compact subset of \mathcal{X}° , the time-averaged orbit $\bar{x}(t) = t^{-1} \int_0^t x(s) ds$ converges to the set of Nash equilibria of \mathcal{G} .*

In the case of the replicator dynamics, this is proved by introducing the auxiliary variables $y_\alpha = \log x_\alpha$ and using the fact that $\dot{y} = \dot{x}/x$. To extend this proof to (HD), we instead define y via the Bregman divergence of h :

Proof of Proposition 7.6. Let $y_\alpha = D_h(e_\alpha, x)$, so $\dot{y}_\alpha = \langle v(x)|x - e_\alpha \rangle = \langle v(x)|x \rangle - v_\alpha(x)$ by Lemma 7.2. Then, for all $\alpha, \beta \in \mathcal{A}$, we get

$$y_\alpha(t) - y_\beta(t) = c_{\alpha\beta} + \int_0^t [v_\beta(x(s)) - v_\alpha(x(s))] ds, \tag{7.12}$$

where $c_{\alpha\beta} = y_\alpha(0) - y_\beta(0)$. Since $x(t)$ is contained in a compact subset of \mathcal{X}° , Lemma 7.1 implies that $\sup_t y_\alpha(t) < \infty$ for all $\alpha \in \mathcal{A}$. Thus, dividing both sides of (7.12) by t and taking the limit $t \rightarrow \infty$, we obtain

$$\lim_{t \rightarrow \infty} [v_\alpha(\bar{x}(t)) - v_\beta(\bar{x}(t))] = 0, \tag{7.13}$$

where we have used the linearity of $v_\alpha(x) = \sum_\beta A_{\alpha\beta} x_\beta$ in x to bring the integral into the arguments of v_α and v_β .

Equation (7.13) implies that if \bar{x}^* is an ω -limit point of $\bar{x}(t)$, then $v_\alpha(\bar{x}^*) = v_\beta(\bar{x}^*)$ for all $\alpha, \beta \in \mathcal{A}$, so \bar{x}^* is a Nash equilibrium of \mathcal{G} . Since \mathcal{X} is compact, every solution of (HD) converges to its ω -limit set, and our assertion follows. \square

Proposition 7.6 applies when the population share of each strategy remains bounded away from zero along all interior solution trajectories, a property known as *permanence*. Formally, a dynamical system on \mathcal{X} is called *permanent* if there exists a threshold $\delta > 0$ such that every interior solution satisfies $\liminf_{t \rightarrow \infty} x_\alpha(t) \geq \delta$ for all $\alpha \in \mathcal{A}$.

Hofbauer and Sigmund (1998) establish a sufficient condition for permanence under the replicator dynamics. Proposition 7.7 extends this result to all continuous Hessian dynamics, providing a sufficient condition for Proposition 7.6 to apply:

Proposition 7.7. *Let \mathcal{G} be a random matching game. Assume that the dynamics (HD) are continuous and there exists some $p \in \mathcal{X}^\circ$ such that*

$$\langle v(x^*)|p - x^* \rangle > 0 \text{ for all boundary rest points } x^* \text{ of (HD)}. \tag{7.14}$$

Then, the dynamics (HD) are permanent.

The proof of Proposition 7.7 follows the proof technique of Theorem 13.6.1 of Hofbauer and Sigmund (1998), and is presented in Appendix D.

7.4. Dominated strategies

We conclude by considering the elimination and survival of dominated strategies under (HD). To that end, recall that $\alpha \in \mathcal{A}$ is *strictly dominated* by $\beta \in \mathcal{A}$ if $v_\alpha(x) < v_\beta(x)$ for all $x \in \mathcal{X}$. More generally, $p \in \mathcal{X}$ is *strictly dominated* by $q \in \mathcal{X}$ if $\langle v(x)|p \rangle < \langle v(x)|q \rangle$ for all $x \in \mathcal{X}$, meaning that the average payoff of a small influx of mutants is always higher when the mutants are distributed according to q rather than p . We then say that $p \in \mathcal{X}$ *becomes extinct* along $x(t)$ if $\min\{x_\alpha(t) : \alpha \in \text{supp}(x^*)\} \rightarrow 0$ as $t \rightarrow \infty$ – or equivalently, if there are no ω -limit points of $x(t)$ in $\mathcal{D}(p)$.

Under the replicator dynamics, it is well known that dominated strategies become extinct along every interior solution trajectory (Akin, 1980). As we show below, this elimination result extends to all continuous Hessian dynamics (HD):

Proposition 7.8. *Under all continuous Hessian dynamics (HD), strictly dominated strategies become extinct along every interior solution orbit.*

Proof. The proof follows a standard argument for the replicator dynamics, replacing the KL divergence (7.4b) with the Bregman divergence (7.1). Specifically, Lemma 7.2 implies that along any interior solution $x(t)$,

$$\begin{aligned} \frac{d}{dt} (D_h(p, x) - D_h(q, x)) &= \langle \dot{x}, x - p \rangle_x - \langle \dot{x}, x - q \rangle_x \\ &= \langle \Pi_x(v^\#(x)), q - p \rangle_x = \langle v^\#(x), q - p \rangle_x = \langle v(x)|q - p \rangle, \end{aligned} \quad (7.15)$$

where the penultimate equality uses the fact that $x \in \mathcal{X}^\circ$. Since q strictly dominates p and v is continuous, we have $\langle v(x)|q - p \rangle \geq a$ for some positive constant $a > 0$, implying that $D_h(p, x(t)) \rightarrow \infty$. Hence, by Lemma 7.1, we conclude that $x(t)$ has no ω -limit points in $\mathcal{D}(p)$. \square

In general, the conclusion of Proposition 7.8 is false for discontinuous Hessian dynamics: Sandholm et al. (2008) construct a four-strategy game with a strictly dominated strategy that is played recurrently by a nonnegligible fraction of the population under the Euclidean projection dynamics (PD). The argument above shows that this strategy must become less common when the state is in the interior of \mathcal{X} ; however, solutions to (PD) are able to enter and leave the boundary of \mathcal{X} , and while there, dominated strategies may become more common. We conjecture that this construction can be suitably extended to all discontinuous Hessian dynamics, but we do not tackle this issue here.

8. Links with reinforcement learning

Under a variety of reinforcement learning processes for normal form games, mixed strategies evolve according to the replicator dynamics – see e.g., Börgers and Sarin (1997), Posch (1997), Rustichini (1999), Hopkins (2002) and Hofbauer et al. (2009). We conclude the paper by describing a broader connection between reinforcement learning and Hessian game dynamics.

8.1. Reinforcement learning

Our starting point is a class of reinforcement learning dynamics for N -player normal form games introduced by Coucheny et al. (2015) and Mertikopoulos and Sandholm (2016). Over the course of play, each player maintains a *score vector* representing the cumulative payoffs of each of his strategies; then, at each moment in time, the player selects a mixed strategy by applying a *choice map* to this score vector, similar in function to the perturbed best response maps used in stochastic fictitious play and perturbed best response dynamics (Fudenberg and Levine, 1998; Hofbauer and Sandholm, 2002, 2007).

Formally, let $v_{k\alpha}(x)$ denote the expected payoff of the α -th strategy of player k at mixed strategy profile $x = (x_1, \dots, x_N)$ in an N -player normal form game. The *choice map* Q_k of player k is then defined as

$$Q_k(y_k) = \arg \max_{x_k \in \mathcal{X}_k} \{ \langle y_k | x_k \rangle - h_k(x_k) \}, \tag{8.1}$$

where $\mathcal{X}_k \equiv \Delta(\mathcal{A}_k)$ denotes the mixed strategy space of player k (\mathcal{A}_k being the corresponding strategy set), and $h_k: \mathcal{X}_k \rightarrow \mathbb{R}$ is a smooth, strongly convex *penalty function*. The reinforcement learning process described above can then be written as

$$\begin{aligned} \dot{y}_k &= v_k(x) \\ x_k &= Q_k(y_k). \end{aligned} \tag{RL}$$

Now, let $g_k = \text{Hess } h_k$ and write $n_{k\alpha}(x) = \sum_{\beta \in \mathcal{A}_k} g_{k,\alpha\beta}^{-1}(x)$. Mertikopoulos and Sandholm (2016) showed that when the mixed strategy profile $x(t)$ is interior, its evolution under (RL) is given by

$$\dot{x}_{k\alpha} = \sum_{\beta \in \mathcal{A}_k} \left[g_{k,\alpha\beta}^{-1}(x) - \frac{n_{k\alpha}(x)n_{k\beta}(x)}{\sum_{\gamma} n_{k\gamma}(x)} \right] v_{k\beta}(x). \tag{RLD}$$

A comparison with (3.33b) shows that the dynamics of mixed strategies under (RL) agree with the Hessian dynamics (HD) at interior states.²⁹

8.2. A common derivation of (HD) and (RLD)

The derivation of (HD) here and of (RLD) in Mertikopoulos and Sandholm (2016) have very different starting points, leaving the reasons behind their equivalence somewhat mysterious. We now make these reasons clearer by deriving both dynamics using a common set of tools. Here we present the basic idea behind the argument, using the theory of convex duality to establish the equivalence of versions of (HD) of (RLD) defined on the entire positive orthant \mathcal{K}° .³⁰ Establishing the equivalence of the original processes on \mathcal{X}° using this approach requires further ideas, which we present in Appendix A.2.

The starting point for both (HD) and (RLD) is the potential function h , assumed here to be smooth, strongly convex, and *steep* at the boundary of \mathcal{K} in the sense of (7.5). The *convex conjugate* of h is then defined as³¹

$$h^*(y) = \sup_{x \in \mathcal{K}^\circ} \{ \langle y | x \rangle - h(x) \}, \quad y \in (\mathbb{R}^{\mathcal{A}})^*. \tag{8.2}$$

²⁹ While we have defined Riemannian game dynamics for single population games and reinforcement learning for N -person normal form games, this difference is of no consequence. One can similarly define Riemannian game dynamics for multipopulation games, or a symmetrized reinforcement learning process for symmetric two-player normal form games (cf. Appendix A.2).

³⁰ The usefulness of convex conjugates in analyzing maps of the form (8.1) is well known in learning and optimization – see e.g., Nemirovski and Yudin (1983), Hofbauer and Sandholm (2002), Shalev-Shwartz (2011), and Mertikopoulos and Zhou (2018).

³¹ For a comprehensive treatment, see Rockafellar (1970) or Hiriart-Urruty and Lemaréchal (2001).

Since h is steep and strongly convex, the supremum above is attained at a unique point $Q_{\mathcal{K}}(y) \in \mathcal{K}^\circ$ (Rockafellar, 1970, Theorem 26.5).³² By the first-order optimality conditions for (8.2), we then get

$$Q_{\mathcal{K}}(y) \equiv \arg \max_{x \in \mathcal{K}^\circ} \{ \langle y | x \rangle - h(x) \} = (Dh)^{-1}(y), \tag{8.3}$$

where $(Dh)^{-1}$ is the inverse function of Dh . Applying the envelope theorem to (8.2), we then obtain

$$Dh^*(y) = Q_{\mathcal{K}}(y). \tag{8.4}$$

It follows from (8.3) and (8.4) that $Dh(Dh^*(y)) \equiv y$, so differentiating and rearranging yields

$$\text{Hess } h^*(y) = (\text{Hess } h(x))^{-1} \tag{8.5}$$

with the inverse of the Hessian matrix of h being evaluated at $x = Q_{\mathcal{K}}(y)$.

We now introduce and compare full-dimensional analogues of continuous Hessian dynamics and (symmetric) reinforcement learning, taking the orthant \mathcal{K}° as the state space (and so defining payoffs v on \mathcal{K}°). In the case of (HD), the full-dimensional domain makes projections redundant, so the induced dynamics take the form

$$\dot{x} = v^\sharp(x) = g^{-1}(x) v(x)^\top = (\text{Hess } h(x))^{-1} v(x)^\top. \tag{HD}_{\mathcal{K}}$$

For its part, the full-dimensional analogue of (RL) is

$$\begin{aligned} \dot{y} &= v(x), \\ x &= Q_{\mathcal{K}}(y), \end{aligned} \tag{RL}_{\mathcal{K}}$$

with $Q_{\mathcal{K}}$ defined as in (8.3) above. Differentiating (RL)_ℳ and applying (8.4) and (8.5) then yields

$$\dot{x} = DQ_{\mathcal{K}}(y) v(x)^\top = \text{Hess } h^*(y) v(x)^\top = (\text{Hess } h(x))^{-1} v(x)^\top. \tag{RLD}_{\mathcal{K}}$$

We can express this argument in words. By definition, (full-dimensional) Riemannian dynamics are obtained by transforming payoffs at an interior population state x using the matrix $g^{-1}(x)$. In the Hessian case, the metric g admits a potential function h , so we have $g^{-1}(x) = (\text{Hess } h(x))^{-1}$ by definition. As for reinforcement learning, differentiation shows that the dynamics of the mixed strategy x are obtained by transforming payoffs at x by the derivative matrix $DQ_{\mathcal{K}}(y)$ of the choice map (8.3). Basic facts about convex conjugacy imply that this derivative is equal to $(\text{Hess } h(x))^{-1}$, where h is the penalty function that generates $Q_{\mathcal{K}}$. This argument establishes the equivalence of (HD)_ℳ and (RLD)_ℳ. For a corresponding argument for the original processes (HD) and (RLD), see Appendix A.2.

8.3. Boundary behavior in the nonsteep regime

Since continuous Hessian dynamics coincide with the reinforcement scheme (RL) when the penalty functions h_k are steep, certain results for interior trajectories – Propositions 7.6 and 7.8 in particular – can be obtained directly from the analysis of Mertikopoulos and Sandholm (2016).

³² Since h is strongly convex, it is bounded below by a strictly convex quadratic function (Hiriart-Urruty and Lemaréchal, 2001, Theorem B.4.1.1), which in turn implies that the domain of h^* is $(\mathbb{R}^n)^*$ (Rockafellar, 1970, Corollary 13.1).

On the other hand, in the nonsteep regime, (RLD) and (HD) agree at interior states, but their behaviors at the boundary differ in a fundamental way. Specifically, while the boundary behavior of discontinuous Hessian dynamics is defined using closest point projections, the reinforcement learning process (RL) for nonsteep h can no longer be reduced to mixed strategy dynamics at all. Instead one must work explicitly with the score variables y_k , which continue to aggregate payoff data of all strategies, even those that are not used. Among other things, this means that a strong cumulative performance of an unused strategy will return it to use.

This difference between how the processes are defined on the boundary has important consequences. For instance, while Sandholm et al. (2008) show that strictly dominated strategies may survive under the Euclidean projection dynamics, Mertikopoulos and Sandholm (2016) prove that the reinforcement learning process (RL) *always* eliminates dominated strategies, whether the penalty functions are steep or not. Under both processes, dominated strategies are initially eliminated along interior solution trajectories. But while they may resurface under (HD), the score variables of (RL) continue to register the poor performance of these strategies, ensuring that they remain extinct for all time.

Appendix A. Connections with other game dynamics

Throughout this appendix, we write $\mathbf{1} \in \mathbb{R}^n$ for the n -dimensional column vector of ones and $\Phi = I - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$ for the Euclidean orthogonal projection of \mathbb{R}^n onto $\mathbb{R}_0^n = \{z \in \mathbb{R}^n : \mathbf{1}^\top z = 0\} = \text{span}(\mathbf{1})^\perp$.³³ Recall also that the (Moore–Penrose) pseudoinverse of a matrix $M \in \mathbb{R}^{n \times n}$ is the unique matrix $M^+ \in \mathbb{R}^{n \times n}$ such that (i) $M^+y = 0$ whenever $y \in \text{range}(M)^\perp$; and (ii) $M^+y = x$ whenever $x \in \text{ker}(M)^\perp$, $y \in \text{range}(M)$, and $Mx = y$ (Friedberg et al., 2002, Sec. 6.7). Since a symmetric matrix $M \in \mathbb{R}^{n \times n}$ satisfies $\text{range}(M) = \text{ker}(M)^\perp$, we have the following well-known algebraic characterization of pseudoinverses:

Lemma A.1. *Let $M \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Then, M^+ is the unique matrix that (i) inverts M on $\text{ker}(M)^\perp = \text{range}(M)$; and (ii) satisfies $\text{ker}(M^+) = \text{ker}(M)$.*

A.1. Interior equivalence of Riemannian dynamics and Hopkins’ dynamics

We now derive the equivalence between (RmD) and Hopkins’ dynamics (2.7) on \mathcal{X}° , as noted in Section 2.3. To begin with, we say that $M \in \mathbb{R}^{n \times n}$ is a Hopkins matrix if it is symmetric, positive definite with respect to \mathbb{R}_0^n , and maps $\mathbf{1}$ to 0. The following lemma establishes a basic characterization of Hopkins matrices:

Lemma A.2.

(i) *If $S \in \mathbb{R}^{n \times n}$ is symmetric positive-definite, then $(\Phi S \Phi)^+$ is a Hopkins matrix and*

$$(\Phi S \Phi)^+ = S^{-1} - \frac{S^{-1}\mathbf{1}\mathbf{1}^\top S^{-1}}{\mathbf{1}^\top S^{-1}\mathbf{1}}. \tag{A.1}$$

³³ In the above and what follows, W^\perp denotes the orthogonal complement of a subspace W of \mathbb{R}^n , defined with respect to the ordinary Euclidean metric. Even though this might seem to suggest that the Euclidean metric plays a special role in what follows, it is just an artifact of writing everything in coordinates instead of abstractly; for a detailed discussion, see Lee (1997).

(ii) Conversely, if M is a Hopkins matrix, then $S = (M + c\mathbf{1}\mathbf{1}^\top)^{-1}$ with $c > 0$ is symmetric positive definite, and $(\Phi S\Phi)^+ = M$.

Proof. To prove part (i), let S be symmetric positive-definite. Then the symmetric matrix $\Phi S\Phi$ is positive definite with respect to \mathbb{R}_0^n and maps $\mathbf{1}$ to 0, so $\text{range}(\Phi S\Phi) = \ker(\Phi S\Phi)^\perp = \mathbb{R}_0^n$. If we denote the right-hand side of (A.1) by \bar{S} , a straightforward calculation shows that $\bar{S}\Phi S\Phi z = z$ for all $z \in \mathbb{R}_0^n$ and $\bar{S}\mathbf{1} = 0$. Thus Lemma A.1 implies that $\bar{S} = (\Phi S\Phi)^+$. That this is a Hopkins matrix is immediate from the fact that $\text{range}(\Phi S\Phi) = \ker(\Phi S\Phi)^\perp = \mathbb{R}_0^n$ and Lemma A.1.

To prove part (ii), let M be a Hopkins matrix. Clearly, the matrix $M + c\mathbf{1}\mathbf{1}^\top$ with $c > 0$ is symmetric positive-definite, so its inverse S is as well. Writing out $(\Phi S\Phi)^+$ using the right-hand side of (A.1) and simplifying the result yields $(\Phi S\Phi)^+ = M$. \square

Proposition A.3 below establishes the equivalence between (2.7) and (RmD) on \mathcal{X}° , and provides a concise third representation for both dynamics. Observe that (A.2c) is expression (3.33b) for (RmD) on \mathcal{X}° , but written in matrix form.

Proposition A.3. *Let $\dot{x} = V(x)$ be a dynamical system on \mathcal{X}° . Then, the following are equivalent:*

(i) *There is a smooth field of Hopkins matrices $M: \mathcal{X}^\circ \rightarrow \mathbb{R}^{A \times A}$ such that*

$$V(x) = M(x)v(x)^\top \quad \text{for all } x \in \mathcal{X}^\circ. \tag{A.2a}$$

(ii) *There is a smooth field of symmetric positive-definite matrices $H: \mathcal{X}^\circ \rightarrow \mathbb{R}^{A \times A}$ such that*

$$V(x) = (\Phi H(x)\Phi)^+v(x)^\top \quad \text{for all } x \in \mathcal{X}^\circ. \tag{A.2b}$$

(iii) *There is a smooth Riemannian metric g on \mathcal{K}° such that*

$$V(x) = \left(g^{-1}(x) - \frac{g^{-1}(x)\mathbf{1}\mathbf{1}^\top g^{-1}(x)}{\mathbf{1}^\top g^{-1}(x)\mathbf{1}} \right) v(x)^\top \quad \text{for all } x \in \mathcal{X}^\circ. \tag{A.2c}$$

Proof. The equivalence of (i) and (ii) follows from Lemma A.2, and the equivalence of (ii) and (iii) follows from Lemma A.2(i) with $S = H(x) = g(x)$.³⁴ \square

A.2. Continuous Hessian dynamics and reinforcement learning

We now complete the common derivation of (HD) and (RLD) on \mathcal{X}° initiated in Section 8.2. Proposition A.3 above shows that the continuous Hessian dynamics (HD) can be expressed as

$$\dot{x} = (\Phi H(x)\Phi)^+v(x)^\top. \tag{A.3}$$

We now extend the argument from Section 8.2 to show that the reinforcement learning dynamics (RLD) also take this form.

To that end, let $\mathcal{Z} = \{x - \frac{1}{n}\mathbf{1} : x \in \mathcal{X}^\circ\} \subset \mathbb{R}_0^n$, and define $h_{\mathcal{Z}}: \mathcal{Z} \rightarrow \mathbb{R}$ as

$$h_{\mathcal{Z}}(z) = h(z + \frac{1}{n}\mathbf{1}). \tag{A.4}$$

³⁴ To formally complete the argument that (ii) implies (iii), we observe without proof that the field S on \mathcal{X}° can be smoothly extended to \mathcal{K}° .

Since \mathcal{Z}° is open relative to \mathbb{R}_0^n , the derivative $Dh_{\mathcal{Z}}$ of $h_{\mathcal{Z}}$ at $z \in \mathcal{Z}^\circ$ can be represented by a covector $Dh_{\mathcal{Z}}(z)$ in $(\mathbb{R}_0^n)^*$. Imposing the Euclidean metric on \mathbb{R}_0^n for convenience, we can identify $(\mathbb{R}_0^n)^*$ with the set of covectors whose components sum to zero, and write the derivative and Hessian of $h_{\mathcal{Z}}$ at $z \in \mathcal{Z}^\circ$ as

$$Dh_{\mathcal{Z}}(z) = Dh(z + \frac{1}{n}\mathbf{1})\Phi, \tag{A.5}$$

$$\text{Hess } h_{\mathcal{Z}}(z) = \Phi H(z + \frac{1}{n}\mathbf{1})\Phi. \tag{A.6}$$

The convex conjugate $h_{\mathcal{Z}}^*$ of $h_{\mathcal{Z}}$ is then defined as

$$h_{\mathcal{Z}}^*(y_0) = \max_{z \in \mathcal{Z}} \{ \langle y | z \rangle - h_{\mathcal{Z}}(z) \}, \quad y_0 \in (\mathbb{R}_0^n)^*, \tag{A.7}$$

where the fact that the domain of $h_{\mathcal{Z}}^*$ is all of $(\mathbb{R}_0^n)^*$ follows from the compactness of \mathcal{Z} . By the basics of convex conjugation (Rockafellar, 1970, Chapter 26), the maps $Dh_{\mathcal{Z}}^*: (\mathbb{R}_0^n)^* \rightarrow \mathcal{Z}$ and $Dh_{\mathcal{Z}}: \mathcal{Z} \rightarrow (\mathbb{R}_0^n)^*$ are inverse to one another, so we have

$$\text{Hess } h_{\mathcal{Z}}^*(y_0) = (\text{Hess } h_{\mathcal{Z}}(dh_{\mathcal{Z}}^*(y_0)))^{-1}, \tag{A.8}$$

with both sides of the equality in (A.8) understood as linear maps from \mathbb{R}_0^n to itself.

Now recall that the (symmetric) reinforcement learning process is defined by

$$\begin{aligned} \dot{y} &= v(x) \\ x &= Q(y). \end{aligned} \tag{RL}$$

The choice map Q , defined in (8.1) in terms of h , satisfies $Q(y) = Q(y\Phi)$ for all $y \in \mathbb{R}^n$. Thus, applying definition (A.4), we can express Q in terms of $h_{\mathcal{Z}}$ as

$$Q(y) = \arg \max_{x \in \mathcal{X}} \{ y\Phi x - h(x) \} = \arg \max_{z \in \mathcal{Z}} \{ y\Phi z - h_{\mathcal{Z}}(z) \} + \frac{1}{n}\mathbf{1} = Dh_{\mathcal{Z}}^*(y\Phi) + \frac{1}{n}\mathbf{1}. \tag{A.9}$$

Hence, substituting (A.9) into (RL), differentiating, and using Eqs. (RL), (A.6), (A.8) and (A.9) yields

$$\begin{aligned} \dot{x} &= \text{Hess } h_{\mathcal{Z}}^*(v(x)\Phi) \Phi v(x)^\top = (\text{Hess } h_{\mathcal{Z}}(Dh_{\mathcal{Z}}^*(v(x)\Phi)))^{-1} \Phi v(x)^\top \\ &= (\text{Hess } h_{\mathcal{Z}}(x - \frac{1}{n}\mathbf{1}))^{-1} \Phi v(x)^\top = (\Phi H(x)\Phi)^+ v(x)^\top, \end{aligned} \tag{A.10}$$

as specified in (A.3).

Appendix B. Extensions of Riemannian metrics

Here we present some technical results concerning the extension of Riemannian metrics from \mathcal{K}° to \mathcal{K} . Proposition B.1 shows that an extendable metric g on \mathcal{K} induces a well-defined scalar product at all points of \mathcal{K} :

Proposition B.1. *Let g be an extendable Riemannian metric on \mathcal{K} . Then, for all $x \in \mathcal{K}$, there exists a unique scalar product $\langle \cdot, \cdot \rangle_x$ on $\text{dom } g(x)$ such that $\langle w, w' \rangle_{x_k} \rightarrow \langle w, w' \rangle_x$ for all $w, w' \in \text{dom } g(x)$ and for every interior sequence $x_k \rightarrow x$. Moreover, if g is minimal-rank extendable, we have $g_{\alpha\beta}^\sharp(x) = 0$ whenever $\alpha, \beta \notin \text{supp}(x)$.*

Proof. Fix some $x \in \mathcal{K}$ and write $g^\sharp(x) = Q^\top \Lambda Q$ where the diagonal matrix Λ consists of the eigenvalues of $g^\sharp(x)$ and Q is an orthogonal matrix ($Q^\top = Q^{-1}$) whose columns are the eigenvectors of $g^\sharp(x)$. Since $g^\sharp(x)$ is positive-semidefinite, its eigenvalues are nonnegative. Furthermore, since $\text{dom } g(x) = \text{im } g^\sharp(x)$, every eigenvector of a nonzero eigenvalue of $g^\sharp(x)$ must lie in $\text{dom } g(x)$: indeed, if $g^\sharp(x)z = \lambda z$ for some $\lambda > 0$, we will also have $z = g^\sharp(x)z\lambda^{-1}$, i.e., $z \in \text{im } g^\sharp(x) = \text{dom } g(x)$. As a result, we may write $g^\sharp(x) = \sum_{\lambda>0} \lambda u_\lambda u_\lambda^\top$ where the summation is taken over all positive eigenvalues $\lambda > 0$ of $g^\sharp(x)$ (assumed for convenience to be distinct) and u_λ is the corresponding column of Q . The metric tensor of the induced scalar product $\langle \cdot, \cdot \rangle_x$ at x is then defined as the pseudoinverse $(g^\sharp(x))^+$ of $g(x)$ (see Appendix A), given here by

$$(g^\sharp(x))^+ = \sum_{\lambda>0} \lambda^{-1} u_\lambda u_\lambda^\top. \tag{B.1}$$

Our continuity and uniqueness claims are then immediate.

Finally, to show that $g_{\alpha\beta}^\sharp(x) = 0$ if $\alpha \notin \text{supp}(x)$ and g is minimal-rank extendable, simply note that $g_{\alpha\beta}^\sharp(x) = e_\alpha^\top g^\sharp(x) e_\beta = \sum_\lambda \lambda e_\alpha^\top u_\lambda u_\lambda^\top e_\beta = 0$ because all eigenvectors of $g^\sharp(x)$ with positive eigenvalues lie in $\mathbb{R}^{\text{supp}(x)} = \text{dom } g(x)$. \square

Remark B.1. In the minimal-rank case, the scalar product $\langle \cdot, \cdot \rangle_x$ on $\text{dom } g(x)$ can be represented by the matrix $g(x) = (g^\sharp(x))^+$. This means that the submatrix $(g_{\alpha\beta}(x))_{\alpha, \beta \in \text{supp}(x)}$ is the inverse of the submatrix $(g_{\alpha\beta}^\sharp(x))_{\alpha, \beta \in \text{supp}(x)}$ while the remaining components (which have no geometric significance) are set to 0.

Next we establish the correctness of the formulas (3.33) for dynamics generated by a minimal-rank extendable metrics.

Proposition B.2. *The formulas (3.33) describe (RmD) generated by a minimal-rank extendable metric g at all states $x \in \mathcal{X}$; moreover, the sums in (3.33) need only be taken over $\beta \in \text{supp}(x)$.*

Proof. The case $x \in \mathcal{X}^\circ$ is covered in the text. Otherwise, if $x \in \text{bd}(\mathcal{X})$, the corresponding g -admissible set is

$$\text{Adm}_g(x) = \text{TC}_{\mathcal{X}}(x) \cap \text{T}_{\mathcal{K}}(x) = \mathbb{R}_0^A \cap \mathbb{R}^{\text{supp}(x)} = \text{T}_{\mathcal{X}}(x). \tag{B.2}$$

By definition, $v^\sharp(x)$ and $n(x)$ lie in $\text{im } g^\sharp(x) = \mathbb{R}^{\text{supp}(x)}$. Moreover, $n(x)$ is normal to $\text{T}_{\mathcal{X}}(x) = \mathbb{R}_0^A \cap \mathbb{R}^{\text{supp}(x)}$ because

$$\langle n(x), z \rangle_x = \mathbf{1}^\top g^\sharp(x) g(x) z = \mathbf{1}^\top z = 0 \quad \text{for all } z \in \mathbb{R}_0^A, \tag{B.3}$$

where the second equality follows from Remark B.1.

Eq. (3.41) shows that the right-hand side of (RmD) is equal to $\Pi_x(v^\sharp(x))$. In light of the facts above, (3.40) shows that $\Pi_x(v^\sharp(x))$ is equal to the right-hand side of (3.33a), the novelty being that $v_\alpha^\sharp(x)$ and $n_\alpha(x)$ vanish whenever $\alpha \notin \text{supp}(x)$. These claims prove the first statement in the proposition. To establish the second statement, simply observe that $g_{\alpha\beta}^\sharp(x) = 0$ whenever α or β is not in $\text{supp}(x)$ by Proposition B.1, so including $\beta \notin \text{supp}(x)$ in the sums in (3.33) is irrelevant. \square

Appendix C. Convergence and stability in dynamical systems

C.1. Definitions

Throughout this appendix, we focus on the dynamics

$$\dot{x} = V(x), \quad x \in \mathcal{X}, \tag{D}$$

and we assume that they admit unique solutions from every initial condition. With this in mind, we say that x^* is an ω -limit point of the solution orbit $x(t)$ if there is an increasing sequence of times $t_n \uparrow \infty$ such that $x(t_n) \rightarrow x^*$. We further say that x^* is *Lyapunov stable* if, for every neighborhood U of x^* , there exists a neighborhood U' of x^* such that every solution orbit $x(t)$ that starts in U' is contained in U for all $t \geq 0$. Finally, we say that x^* is *attracting* if there is a neighborhood U of x^* such that every solution that starts in U converges to x^* , and x^* is called *asymptotically stable* if it is Lyapunov stable and attracting. In this case, the maximal (relatively) open set of states from which solutions converge to x^* is called the *basin* of x^* ; if the basin of x^* is all of \mathcal{X} , we say that x^* is *globally asymptotically stable*.

C.2. A global convergence result

A standard result from dynamical systems states that if a smooth dynamical system on a compact set admits a strict global Lyapunov function, all ω -limit points are rest points (see e.g., Sandholm, 2010b, Theorem 7.B.3). The proof of this result relies on the continuity of solutions on initial conditions, a property which is not easily established for discontinuous dynamics. In Proposition C.1 below, we present a global convergence result that does not require continuity of solutions in initial conditions, but instead relies on a l.s.c. lower bound on the derivative of the Lyapunov function. To state it, let

$$\text{RP} = \{x \in \mathcal{X} : V(x) = 0\} \tag{C.1}$$

denote the set of rest points of the dynamics (D). We then have:

Proposition C.1. *Let $x(t)$ be an absolutely continuous solution orbit of (D) and let $\Gamma_+ = x(\mathbb{R}_+)$ denote the set of points visited by $x(t)$. Assume further that RP is closed and there exist functions $L : \mathcal{X} \rightarrow \mathbb{R}$ and $\phi : \mathcal{X} \rightarrow \mathbb{R}_+$ such that*

- (i) L is differentiable in a neighborhood of Γ_+ .
- (ii) ϕ is lower semi-continuous and $\phi(x) = 0$ if and only if $x \in \text{RP}$.
- (iii) $\langle DL(x) | V(x) \rangle \geq \phi(x)$ for all $x \in \Gamma_+$.

Then, $x(t)$ converges to RP.

Proof. By absolute continuity and Conditions (i) and (iii) above, we get

$$L(x(t)) - L(x(0)) = \int_0^t \langle DL(x(s)) | V(x(s)) \rangle ds \geq \int_0^t \phi(x(s)) ds \geq 0, \tag{C.2}$$

i.e., L is nondecreasing along $x(t)$. Furthermore, since \mathcal{X} is compact, $x(t)$ admits at least one ω -limit point $x^* \in \text{cl}(\Gamma_+) \subseteq \mathcal{X}$. Assume now that $x(t)$ admits an ω -limit point x^ω such that

$V(x^\omega) \neq 0$. Since RP is closed, Condition (ii) implies that there is a compact neighborhood K of x^ω and some $a > 0$ such that $\phi(x) \geq a > 0$ for all $x \in K$. With this in mind, we consider two complementary cases below:

Case 1. Suppose there exists some $T \geq 0$ such that $x(t) \in K$ for all $t \geq T$. Then $\phi(x(t)) \geq a$ for all $t \geq T$, so (C.2) yields $\lim_{t \rightarrow \infty} L(x(t)) = \infty$, a contradiction.

Case 2. Assume instead that, for all $T \geq 0$, we have $x(t) \notin K$ for some $t \geq T$. In this case, there exist open neighborhoods U and U' of x^ω with $\text{cl}(U) \subseteq U' \subset K$, and interlaced sequences $t_n, t'_n \uparrow \infty$ such that, for all n : (i) $t_n < t'_n < t_{n+1}$; (ii) $x(t_n) \in U, x(t'_n) \in K \setminus U'$; and (iii) $x(t) \in K$ whenever $t \in [t_n, t'_n]$. Then, since $|\dot{x}_\alpha(t)|$ is bounded from above by $V_{\max} \equiv \sup_{x \in \mathcal{X}} \max_\beta |V_\beta(x)| < \infty$, the time intervals $\delta_n \equiv t'_n - t_n$ will be bounded from below by $\delta_{\min} \equiv \text{dist}(\text{cl}(U), K \setminus U') / V_{\max} > 0$. We thus get

$$L(x(t'_n)) - L(x(0)) \geq \int_0^{t'_n} \phi(x(s)) ds \geq \sum_{j=1}^n \int_{t_j}^{t'_j} \phi(x(s)) ds \geq n\delta_{\min}, \tag{C.3}$$

i.e., $L(x(t'_n)) \rightarrow \infty$, a contradiction. \square

To apply Proposition C.1 to discontinuous Riemannian dynamics in potential games, we need the following result:

Lemma C.2. *The speed of motion $\|V(x)\|_x$ of the dynamics (RmD) is l.s.c. on \mathcal{X} .*

Proof. If the dynamics (RmD) are continuous, our claim follows immediately from the continuity of the underlying metric – in fact, $\|V(x)\|_x$ is continuous in this case. Otherwise, if (RmD) is discontinuous, recall that $V(x) \equiv \Pi_x(v^\sharp(x))$ is simply the projection of $v^\sharp(x)$ on the tangent cone $\text{TC}_{\mathcal{X}}(x)$ to \mathcal{X} at x (because $\text{Adm}_g(x) = \text{TC}_{\mathcal{X}}(x)$ in that case). Therefore, if we write $V^\perp(x) = v^\sharp(x) - V(x)$ for the projection of $v^\sharp(x)$ on the normal cone $\text{NC}_{\mathcal{X}}(x)$ to \mathcal{X} at x , Moreau’s decomposition theorem and the Cauchy–Schwarz inequality yield

$$\langle v^\sharp(x), z \rangle_x = \langle V(x) + V^\perp(x), z \rangle_x = \langle V(x), z \rangle_x \leq \|V(x)\|_x \|z\|_x, \tag{C.4}$$

for all $z \in \text{TC}_{\mathcal{X}}(x)$, with the inequality binding if and only if $z \propto V(x)$. We thus obtain the characterization

$$\|V(x)\|_x = \max_{z \in \text{TC}_{\mathcal{X}}(x) \cap B(x)} \langle v^\sharp(x), z \rangle_x, \tag{C.5}$$

where $B(x) = \{z \in \mathbb{R}^A : \|z\|_x \leq 1\}$.

Note now that the correspondence $x \mapsto \text{TC}_{\mathcal{X}}(x)$ is l.s.c. because it is constant on the interior of each face of \mathcal{X} and $\text{TC}_{\mathcal{X}}(x) \subseteq \text{TC}_{\mathcal{X}}(y)$ whenever $\text{supp}(x) \subseteq \text{supp}(y)$. This shows that the constraint correspondence $x \mapsto \text{TC}_{\mathcal{X}}(x) \cap B(x)$ of (C.5) is l.s.c.; since the objective function $\langle v^\sharp(x), z \rangle_x$ of (C.5) is jointly continuous in x and z , a precursor to the maximum theorem (Aliprantis and Border, 1999, Lemma 16.30) implies that $x \mapsto \|V(x)\|_x$ is itself l.s.c., as claimed. \square

Appendix D. Additional proofs

In this appendix, we collect some proofs that are too technical for the main text. We begin with the proof of Proposition 6.1 regarding the existence and uniqueness of solutions to (RmD).

As noted in Section 6, we only need to prove part (ii), which concerns the case of full-rank extendable metrics. Existence of forward solutions of (RmD) on \mathcal{X} follows from general results of Aubin and Cellina (1984) on solutions to discontinuous differential equations; see Lahkar and Sandholm (2008) for a summary of their argument. Thus, it remains to show that forward solutions to (RmD) from each initial condition in \mathcal{X} are unique. This conclusion follows from the following lemma:

Lemma D.1. *Let $x(t)$ and $x'(t)$ be solutions to (RmD), and let*

$$P(t) = \|x'(t) - x(t)\|_{x(t)}^2 e^{-\lambda t}. \tag{D.1}$$

If $\lambda > 0$ is large enough, then $P(t)$ is nonincreasing.

Given this lemma, we immediately obtain:

Proof of Proposition 6.1. Let $x(t)$ and $x'(t)$ be solutions of (RmD) with $x(0) = x'(0)$. We then get $P(t) = P(0) = 0$ for all $t \geq 0$, so $x(t) = x'(t)$ for all $t \geq 0$. \square

To prove Lemma D.1, we need one final auxiliary result. Let $V(x) = \Pi_x(v^\sharp(x))$ denote the right-hand side of (RmD). Then, as we show below, V satisfies a one-sided Lipschitz condition with respect to the underlying metric:

Lemma D.2. *There exists some $K_V > 0$ such that*

$$\langle V(x') - V(x), x' - x \rangle_x \leq K_V \|x' - x\|_x^2 \text{ for all } x, x' \in \mathcal{X}. \tag{D.2}$$

Proof of Lemma D.2. Write $w(x) = v^\sharp(x)$ and $w^\perp(x) = w(x) - \Pi_x(w(x))$. Since g is full-rank extendable, Π_x is the orthogonal projection onto $\text{TC}_{\mathcal{X}}(x)$ with respect to $\langle \cdot, \cdot \rangle_x$. We thus obtain

$$\begin{aligned} \langle V(x') - V(x), x' - x \rangle_x &= \langle w(x') - w(x), x' - x \rangle_x - \langle w^\perp(x') - w^\perp(x), x' - x \rangle_x \\ &= \langle w(x') - w(x), x' - x \rangle_x + \langle w^\perp(x), x' - x \rangle_x \\ &\quad + \langle w^\perp(x'), x - x' \rangle_{x'} + (w^\perp(x'))^\top (g(x') - g(x))(x' - x) \\ &\leq K_w \|x' - x\|_x^2 + (w^\perp(x'))^\top (g(x') - g(x))(x' - x), \end{aligned} \tag{D.3}$$

where the bound for the first term in the last line follows from the Cauchy–Schwarz inequality and the Lipschitz continuity of w , while the rest follows from Moreau’s decomposition theorem. To bound the last term, write $g_\alpha(x)$ for the α -th row of $g(x)$, let $W_{\max}^\perp = \max_{\alpha \in \mathcal{A}} \max_{x \in \mathcal{X}} w_\alpha^\perp(x)$, and let $\|\cdot\|_2$ denote the standard Euclidean norm. Then, if $C > 0$ is chosen sufficiently large, we get

$$\begin{aligned} &(w^\perp(x'))^\top (g(x') - g(x))(x' - x) \\ &\leq \sum_{\alpha \in \mathcal{A}} w_\alpha^\perp(x') \|g_\alpha(x') - g_\alpha(x)\|_2 \|x' - x\|_2 \\ &\leq W_{\max}^\perp \|x' - x\|_2 \sum_{\alpha \in \mathcal{A}} \|g_\alpha(x') - g_\alpha(x)\|_2 \leq W_{\max}^\perp C \|x' - x\|_x^2. \end{aligned} \tag{D.4}$$

In the above, the first inequality is an immediate corollary of the Cauchy–Schwarz inequality; the last one follows from the equivalence of norms on \mathbb{R}^A and the fact that g is C^1 on \mathcal{X} ; finally, C can be chosen independently of x and x' because \mathcal{X} is compact. Combining (D.3) and (D.4) completes our proof. \square

With Lemma D.2 at hand, we finally obtain:

Proof of Lemma D.1. Define $\dot{g}(x) = (\dot{g}_{\alpha\beta}(x))_{\alpha,\beta \in \mathcal{A}}$ by $\dot{g}_{\alpha\beta}(x) = \langle Dg_{\alpha\beta}(x) | V(x) \rangle = \sum_{\kappa \in \mathcal{A}} V_{\kappa}(x) \partial_{\kappa} g_{\alpha\beta}(x)$, and let $K_g = \max_{\alpha,\beta \in \mathcal{A}} \max_{x \in \mathcal{X}} \dot{g}_{\alpha\beta}(x) < \infty$ (recall that g is C^1). Then, for all $t \geq 0$ such that $x(t)$ and $x'(t)$ are differentiable, we have

$$\begin{aligned} \dot{P} &= 2\langle V(x') - V(x), x' - x \rangle_x e^{-\lambda t} + (x' - x)^{\top} \dot{g}(x)(x' - x) e^{-\lambda t} - \lambda \|x' - x\|_x^2 e^{-\lambda t} \\ &\leq -(\lambda - 2K_V - K_g) \|x' - x\|_x^2 e^{-\lambda t}, \end{aligned} \tag{D.5}$$

where we used Lemma D.2 to bound the second term in the first line. Taking $\lambda > 2K_V + K_g$ then yields $\dot{P} \leq 0$; since $x(t)$ and $x'(t)$ are absolutely continuous, we conclude that $P(t)$ is nondecreasing. \square

We close this appendix with the proof of our permanence criterion:

Proof of Proposition 7.7. Define $P: \mathcal{K} \rightarrow \mathbb{R}$ as $P(x) = -\exp(\sum_{\alpha} p_{\alpha} D_h(e_{\alpha}, x))$ for $x \in \mathcal{K}^{\circ}$ and $P(x) = 0$ for $x \in \text{bd}(\mathcal{K})$. The steepness of h implies that P is continuous, while Lemma 7.2 implies that $\frac{d}{dt} \log(P(x)) = \Psi(x) \equiv \langle v(x) | p - x \rangle$ for all $x \in \mathcal{X}^{\circ}$. Hence, by Theorem 12.2.1 of Hofbauer and Sigmund (1998), it suffices to show that the function Ψ is an average Lyapunov function for (HD), meaning that, for every initial condition $x(0) \in \text{bd}(\mathcal{X})$, there is a $t > 0$ such that

$$\frac{1}{t} \int_0^t \Psi(x(s)) ds > 0. \tag{D.6}$$

We proceed by induction on the cardinality of the support of the initial condition. The claim is trivial if this cardinality is 1. For the inductive step, suppose that (D.6) holds when the cardinality is $k \in \{1, \dots, |\mathcal{A}| - 2\}$, and consider an initial condition $x(0)$ whose support \mathcal{A}' has cardinality $k + 1$. If $x(t)$ converges to the boundary of the face \mathcal{X}' of \mathcal{X} spanned by \mathcal{A}' , then our claim follows from the inductive hypothesis and the same arguments as in the proof of Theorem 12.2.2 in Hofbauer and Sigmund (1998). If instead $x(t)$ does not converge to the boundary of \mathcal{X}' , then there exists a $\delta > 0$ and an increasing sequence of times $t_n \uparrow \infty$ with $x_{\alpha}(t_n) \geq \delta > 0$ for all $\alpha \in \mathcal{A}'$. Then, letting $\bar{x}_{\alpha}(t) = t^{-1} \int_0^t x_{\alpha}(s) ds$ and $\bar{u}(t) = t^{-1} \int_0^t \langle v(x(s)) | x(s) \rangle ds$, we may assume (by descending to a subsequence of t_n if necessary) that $\bar{x}(t_n)$ and $\bar{u}(t_n)$ converge to some \bar{x}^* and \bar{u}^* respectively as $n \rightarrow \infty$.

We now claim that $v_{\alpha}(\bar{x}^*) = \bar{u}^*$ for all $\alpha \in \mathcal{A}'$, implying that \bar{x}^* is a restricted equilibrium of \mathcal{G} . Indeed, let $y_{\alpha}(t) = D_h(e_{\alpha}, x(t))$ for all $\alpha \in \mathcal{A}'$. Then, $\dot{y}_{\alpha} = \langle v(x) | x \rangle - v_{\alpha}(x)$ by Lemma 7.2, so the linearity of $v(x)$ in x implies that

$$\frac{1}{t} \int_0^t \langle v(x(s)) | x(s) \rangle ds - v_{\alpha}(\bar{x}(t)) = \frac{y_{\alpha}(t) - y_{\alpha}(0)}{t}. \tag{D.7}$$

Given that $x(t_n)$ remains a minimal positive distance away from $\text{bd}(\mathcal{X}')$, it follows that $y_{\alpha}(t_n)$ is bounded from above for all $\alpha \in \mathcal{A}'$. Therefore, the right-hand side of (D.7) vanishes as $t_n \rightarrow \infty$, implying in turn that $v_{\alpha}(\bar{x}^*) = \bar{u}^*$ for all $\alpha \in \mathcal{A}'$, as claimed.

Now, since \bar{x}^* is a restricted equilibrium of \mathcal{G} , Proposition 6.3 implies that it is a boundary rest point of (HD), so $\bar{u}^* = \langle v(\bar{x}^*) | \bar{x}^* \rangle < \langle v(\bar{x}^*) | p \rangle$ by (7.14). Moreover, since $\Psi(x) = \langle v(x) | p - x \rangle$, setting $t = t_n$ in (D.6) yields

$$t_n^{-1} \int_0^{t_n} \Psi(x(s)) ds = t_n^{-1} \int_0^{t_n} \langle v(x(s)) | p - x(s) \rangle ds = \langle v(\bar{x}(t_n)) | p \rangle - \bar{u}(t_n) \quad (\text{D.8})$$

so $\lim_{n \rightarrow \infty} t_n^{-1} \int_0^{t_n} \Psi(x(s)) ds = \langle v(\bar{x}^*) | p \rangle - \bar{u}^* > 0$. This establishes (D.6) for large enough $t = t_n$, completing our proof. \square

References

- Akin, E., 1979. The Geometry of Population Genetics. Lecture Notes in Biomathematics, vol. 31. Springer-Verlag.
- Akin, E., 1980. Domination or equilibrium. *Math. Biosci.* 50, 239–250.
- Aliprantis, C.D., Border, K.C., 1999. *Infinite Dimensional Analysis: A Hitchhiker's Guide*, second edition. Springer, Berlin.
- Alvarez, F., Bolte, J., Brahic, O., 2004. Hessian Riemannian gradient flows in convex programming. *SIAM J. Control Optim.* 43 (2), 477–501.
- Aubin, J.-P., Cellina, A., 1984. *Differential Inclusions*. Springer, Berlin.
- Bayer, D.A., Lagarias, J.C., 1989. The nonlinear geometry of linear programming I. Affine and projective scaling trajectories. *Trans. Am. Math. Soc.* 314, 499–526.
- Benaïm, M., Weibull, J.W., 2003. Deterministic approximation of stochastic evolution in games. *Econometrica* 71 (3), 873–903.
- Björnerstedt, J., Weibull, J.W., 1996. Nash equilibrium and evolution by imitation. In: Arrow, K.J., Colomatto, E., Perlman, M., Schmidt, C. (Eds.), *The Rational Foundations of Economic Behavior*. St. Martin's Press, New York, NY, pp. 155–181.
- Bolte, J., Teboulle, M., 2003. Barrier operators and associated gradient-like dynamical systems for constrained minimization problems. *SIAM J. Control Optim.* 42 (4), 1266–1292.
- Börgers, T., Sarin, R., 1997. Learning through reinforcement and replicator dynamics. *J. Econ. Theory* 77, 1–14.
- Bravo, M., Mertikopoulos, P., 2017. On the robustness of learning in games with stochastically perturbed payoff observations. In: *John Nash Memorial Issue*. *Games Econ. Behav.* 103, 41–66.
- Bregman, L.M., 1967. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. Math. Phys.* 7 (3), 200–217.
- Brown, G.W., 1951. Iterative solutions of games by fictitious play. In: Koopmans, T.C., et al. (Eds.), *Activity Analysis of Production and Allocation*. Wiley, New York, pp. 374–376.
- Coucheny, P., Gaujal, B., Mertikopoulos, P., 2015. Penalty-regulated dynamics and robust learning procedures in games. *Math. Oper. Res.* 40 (3), 611–633.
- Demichelis, S., Ritzberger, K., 2003. From evolutionary to strategic stability. *J. Econ. Theory* 113, 51–75.
- Duistermaat, J.J., 2001. On Hessian Riemannian structures. *Asian J. Math.* 5, 79–91.
- Fisher, R.A., 1930. *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford.
- Friedberg, S.H., Insel, A.J., Spence, L.E., 2002. *Linear Algebra*, fourth edition. Pearson.
- Friedman, D., 1991. Evolutionary games in economics. *Econometrica* 59 (3), 637–666.
- Friesz, T.L., Bernstein, D., Mehta, N.J., Tobin, R.L., Ganjalizadeh, S., 1994. Day-to-day dynamic network disequilibria and idealized traveler information systems. *Oper. Res.* 42, 1120–1136.
- Fudenberg, D., Levine, D.K., 1998. *The Theory of Learning in Games*. Economic Learning and Social Evolution, vol. 2. MIT Press, Cambridge, MA.
- Harper, M., 2011. Escort evolutionary game theory. *Phys. D: Nonlinear Phenom.* 240, 1411–1415.
- Hart, S., Mas-Colell, A., 2001. A general class of adaptive strategies. *J. Econ. Theory* 98, 26–54.
- Helbing, D., 1992. A mathematical model for behavioral changes by pair interactions. In: Haag, G., Mueller, U., Troitzsch, K.G. (Eds.), *Economic Evolution and Demographic Change: Formal Models in Social Sciences*. Springer, Berlin, pp. 330–348.
- Hines, W.G.S., 1980. Three characterizations of population strategy stability. *J. Appl. Probab.* 17, 333–340. Correction, *R. Cressman and W.G.S. Hines*, 21 (1984), 213–214.
- Hiriart-Urruty, J.-B., Lemaréchal, C., 2001. *Fundamentals of Convex Analysis*. Springer, Berlin.
- Hofbauer, J., Sandholm, W.H., 2002. On the global convergence of stochastic fictitious play. *Econometrica* 70 (6), 2265–2294.
- Hofbauer, J., Sandholm, W.H., 2007. Evolution in games with randomly disturbed payoffs. *J. Econ. Theory* 132, 47–69.
- Hofbauer, J., Sandholm, W.H., 2009. Stable games and their dynamics. *J. Econ. Theory* 144, 1710–1725.

- Hofbauer, J., Schuster, P., Sigmund, K., 1979. A note on evolutionarily stable strategies and game dynamics. *J. Theor. Biol.* 81, 609–612.
- Hofbauer, J., Sigmund, K., 1990. Adaptive dynamics and evolutionary stability. *Appl. Math. Lett.* 3, 75–79.
- Hofbauer, J., Sigmund, K., 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.
- Hofbauer, J., Sorin, S., Viossat, Y., 2009. Time average replicator and best reply dynamics. *Math. Oper. Res.* 34 (2), 263–269.
- Hopkins, E., 1999. A note on best response dynamics. *Games Econ. Behav.* 29, 138–150.
- Hopkins, E., 2002. Two competing models of how people learn in games. *Econometrica* 70, 2141–2166.
- Izquierdo, L.R., Izquierdo, S.S., Sandholm, W.H., 2018. An introduction to ABED: agent-based simulation of evolutionary game dynamics. Universidad de Burgos, Universidad de Valladolid, and University of Wisconsin. Unpublished manuscript.
- Kimura, M., 1958. On the change of population fitness by natural selection. *Heredity* 12, 145–167.
- Lahkar, R., Sandholm, W.H., 2008. The projection dynamic and the geometry of population games. *Games Econ. Behav.* 64, 565–590.
- Laraki, R., Mertikopoulos, P., 2015. Inertial game dynamics and applications to constrained optimization. *SIAM J. Control Optim.* 53 (5), 3141–3170.
- Lee, J.M., 1997. *Riemannian Manifolds: An Introduction to Curvature*. Graduate Texts in Mathematics, vol. 176. Springer.
- Lee, J.M., 2003. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics, vol. 218. Springer-Verlag, New York, NY.
- Maynard Smith, J., Price, G.R., 1973. The logic of animal conflict. *Nature* 246, 15–18.
- Mertikopoulos, P., Moustakas, A.L., 2010. The emergence of rational behavior in the presence of stochastic perturbations. *Ann. Appl. Probab.* 20 (4), 1359–1388.
- Mertikopoulos, P., Sandholm, W.H., 2016. Learning in games via reinforcement and regularization. *Math. Oper. Res.* 41 (4), 1297–1324.
- Mertikopoulos, P., Sandholm, W.H., 2018. Discrete choice and game dynamics with lattice-structured traits. CNRS and University of Wisconsin. Unpublished manuscript.
- Mertikopoulos, P., Staudigl, M., 2018. On the convergence of gradient-like flows with noisy gradient input. *SIAM J. Optim.* 28 (1), 163–197.
- Mertikopoulos, P., Zhou, Z., 2018. Learning in games with continuous action sets and unknown payoff functions. *Math. Program.* <https://doi.org/10.1007/s10107-018-1254-8>.
- Monderer, D., Shapley, L.S., 1996. Potential games. *Games Econ. Behav.* 14 (1), 124–143.
- Nagurny, A., Zhang, D., 1997. Projected dynamical systems in the formulation, stability analysis, and computation of fixed demand traffic network equilibria. *Transp. Sci.* 31, 147–158.
- Nemirovski, A.S., Yudin, D.B., 1983. *Problem Complexity and Method Efficiency in Optimization*. Wiley, New York, NY.
- Posch, M., 1997. Cycling in a stochastic learning algorithm for normal form games. *J. Evol. Econ.* 7, 193–207.
- Rockafellar, R.T., 1970. *Convex Analysis*. Princeton University Press, Princeton, NJ.
- Roth, G., Sandholm, W.H., 2013. Stochastic approximations with constant step size and differential inclusions. *SIAM J. Control Optim.* 51 (1), 525–555.
- Rustichini, A., 1999. Optimal properties of stimulus-response learning models. *Games Econ. Behav.* 29, 230–244.
- Sandholm, W.H., 2001. Potential games with continuous player sets. *J. Econ. Theory* 97, 81–108.
- Sandholm, W.H., 2005. Excess payoff dynamics and other well-behaved evolutionary dynamics. *J. Econ. Theory* 124, 149–170.
- Sandholm, W.H., 2010a. Local stability under evolutionary game dynamics. *Theor. Econ.* 5, 27–50.
- Sandholm, W.H., 2010b. *Population Games and Evolutionary Dynamics*. Economic Learning and Social Evolution. MIT Press, Cambridge, MA.
- Sandholm, W.H., 2014. Probabilistic interpretations of integrability for game dynamics. *Dyn. Games Appl.* 4, 95–106.
- Sandholm, W.H., 2015. Population games and deterministic evolutionary dynamics. In: Young, H.P., Zamir, S. (Eds.), *Handbook of Game Theory*, vol. 4. Elsevier, pp. 703–778.
- Sandholm, W.H., Dokumaci, E., Lahkar, R., 2008. The projection dynamic and the replicator dynamic. *Games Econ. Behav.* 64, 666–683.
- Schlag, K.H., 1998. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econ. Theory* 78, 130–156.
- Schuster, P., Sigmund, K., 1983. Replicator dynamics. *J. Theor. Biol.* 100 (3), 533–538.

- Schuster, P., Sigmund, K., Hofbauer, J., Wolff, R., 1981. Self-regulation of behaviour in animal societies I: symmetric contests. *Biol. Cybern.* 40, 1–8.
- Shahshahani, S., 1979. A New Mathematical Framework for the Study of Linkage and Selection. *Memoirs of the American Mathematical Society*, vol. 211.
- Shalev-Shwartz, S., 2011. Online learning and online convex optimization. *Found. Trends Mach. Learn.* 4 (2), 107–194.
- Swinkels, J.M., 1993. Adjustment dynamics and rational play in games. *Games Econ. Behav.* 5, 455–484.
- Taylor, P.D., Jonker, L.B., 1978. Evolutionary stable strategies and game dynamics. *Math. Biosci.* 40 (1–2), 145–156.
- Tsakas, E., Voorneveld, M., 2009. The target projection dynamic. *Games Econ. Behav.* 67, 708–719.
- Tsallis, C., 1988. Possible generalization of Boltzmann–Gibbs statistics. *J. Stat. Phys.* 52, 479–487.
- Weibull, J.W., 1995. *Evolutionary Game Theory*. MIT Press, Cambridge, MA.
- Zusai, D., 2018. Gains in evolutionary dynamics: a unifying approach to stability for contractive games and ESS. Temple University. Unpublished manuscript.