Stochastic Approximations with Constant Step Size and Differential Inclusions*

Grégory Roth† and William H. Sandholm‡

July 18, 2012

Abstract

We consider stochastic approximation processes with constant step size whose associated deterministic system is an upper semicontinous differential inclusion. We prove that over any finite time span, the sample paths of the stochastic process are closely approximated by a solution of the differential inclusion with high probability. We then analyze infinite horizon behavior, showing that if the process is Markov, its stationary measures must become concentrated on the Birkhoff center of the deterministic system. Our results extend those of Benaïm for settings in which the deterministic system is Lipschitz continuous, and build on the work of Benaïm, Hofbauer, and Sorin for the case of decreasing step sizes. We apply our results to models of population dynamics in games, obtaining new conclusions about the medium and long run behavior of myopic optimizing agents.

1. Introduction

Stochastic approximation theory, whose origins lie in the analysis of stochastic optimization algorithms, links the behavior of certain recursively-defined stochastic processes with that of associated deterministic dynamical systems. Much of the theory addresses the behavior of processes with decreasing step sizes $\{\varepsilon_k\}_{k=1}^{\infty}$, which arise naturally when the state represents a time-averaged quantity. Early work focused on cases in which the approximating dynamical system is linear or a gradient system (Ljung (1977); Duflo (1996); Kushner and Yin (1997)). A series of papers of Benaïm (1998, 1999) and Benaïm

*We thank Michel Benaïm, Mathieu Faure, Tom Kurtz, Mathias Staudigl, and two anonymous referees for helpful comments and discussions. Financial support from Swiss NSF Grant 200020-130574 and U.S. NSF Grants SES-0851580 and SES-1155135 are gratefully acknowledged.

†Institut de Mathématiques, Université de Neuchâtel, Rue Emile-Argand 11, CH-2007 Neuchâtel, Switzerland. e-mail: gregory.roth@unine.ch.

‡Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu.

and Hirsch (1996, 1999) considered the case of Lipschitz continuous dynamical systems, including a model from game theory known as stochastic fictitious play (Fudenberg and Kreps (1993), Fudenberg and Levine (1998), Hofbauer and Sandholm (2002)). Their main results show that the stochastic process converges almost surely to an internally chain transitive set of the deterministic flow.

More recently, Benaïm et al. (2005) develop the theory for the case in which the deterministic dynamic is not a differential equation, but a upper semicontinuous differential inclusion. This level of generality is important in game-theoretic applications in which choices are determined by exact maximization, as the maximizer correspondence is neither single-valued nor continuous, but rather multi-valued and upper semicontinuous. An important application of these results is to the original fictitious play process, defined in the early days of game theory by Brown (1949, 1951). Indeed, Benaïm et al.'s (2005) analysis provides a simple proof that the fictitious play process must converge to a connected set of Nash equilibria in any two-player zero-sum game and in any potential game.[1]

While the discussion above concerns stochastic approximation processes with decreasing step size, a distinct branch of the theory has considered processes with constant step size. Such processes are parameterized by a constant $\varepsilon > 0$, which represents both the step size of the process and the inverse of the step rate of the process, so that the expected increment per time unit remains fixed as ε varies. The analysis concerns the behavior of the processes as ε approaches zero.

While the decreasing step size theory leads naturally to one basic limit theorem, the constant step size theory leads to two. Kurtz (1970) (also see Ethier and Kurtz (1986), Benveniste et al. (1990), and Benaïm (1998)) shows that over finite time horizons, the sample paths of the stochastic process are well-approximated by a solution trajectory of the deterministic system with probability approaching one. To describe infinite horizon behavior, Benaïm (1998) shows that the stationary measures of the stochastic process must become concentrated on the Birkhoff center of the deterministic system.

Existing analyses of stochastic approximation with constant step size focus on cases in which the deterministic system is Lipschitz continuous. Building on the work of Benaïm et al. (2005) for the decreasing step size case, we develop the dynamical systems approach to stochastic approximation with constant step size when the deterministic system is an upper semicontinuous differential inclusion. Our two main results parallel those of Benaïm (1998) for the Lipschitz case. We prove a finite-horizon deterministic approximation theorem, and we prove that stationary measures of the stochastic process must become concentrated on the Birkhoff center of the deterministic system.

[1] A similar but more direct proof for the zero-sum case is provided by Hofbauer and Sorin (2006).

Stochastic approximation theory with constant step size is of basic importance in evolutionary game theory, where it is used to understand the dynamics of behavior in large populations of strategically interacting agents whose decisions are based on simple myopic rules (Benaïm and Weibull (2003), Sandholm (2003, 2010)). Much research in this field focuses on deterministic continuous-time dynamics, which provide an idealized description of the population's aggregate behavior. Stochastic approximation methods connect the properties of these ideal systems with those of processes describing the stochastic evolution of play in large but finite populations. Existing results have drawn these connections for cases in which the deterministic system is Lipschitz continuous. But many interesting deterministic dynamics from game theory, particularly those reflecting exact optimization by individual agents, take the form of differential inclusions. These include the best response dynamics of Gilboa and Matsui (1991) and Hofbauer (1995), which (along with the replicator dynamic of Taylor and Jonker (1978)) represent one of the two leading dynamics studied in the field. More recent optimization-based models, including the refined best response dynamics of Balkenborg et al. (2011), the tempered best response dynamics of Zusai (2011), and the sampling best response dynamics of Oyama et al. (2012), also yield differential inclusions. By combining the main results in this paper with existing analyses of the relevant deterministic dynamics, we obtain a variety of new results on medium and long run behavior under optimization-based decision protocols.

A few recent papers have obtained results for special cases of the model considered here. Gorodeisky (2008, 2009) supposes that the deterministic dynamic is piecewise Lipschitz continuous with unique solution from each initial condition, and uses his analysis to study the behavior of myopic optimizers in the Matching Pennies game. Gast and Gaujal (2010) consider a Markovian model whose mean field is arbitrary, and obtain finite-horizon deterministic approximation results stated in terms of a regularized mean field that takes the form of a differential inclusion.

The rest of the paper proceeds as follows. Section 2 provides background on differential inclusions and introduces the stochastic approximation processes. Section 3 presents our main results. Section 4 offers applications to evolutionary game dynamics. Section 5 contains the proofs of most results.

2. Definitions

2.1 Differential inclusions

A correspondence $V\colon \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ defines a differential inclusion on $\mathbb{R}^n$ via

$$\text{(DI)} \qquad \dot{x} \in V(x).$$

Definition 2.1. We call (DI) a *good upper semicontinuous* (or *good USC*) *differential inclusion* if V is

(i) Nonempty: $V(x) \neq \varnothing$ for all $x \in \mathbb{R}^n$;
(ii) Convex valued: $V(x)$ is convex for all $x \in \mathbb{R}^n$;
(iii) Bounded: There exists a $K \in \mathbb{R}$ such that $\sup\{|y| : y \in V(x)\} \leq K$ for all $x \in \mathbb{R}^n$;
(iv) Upper semicontinuous: The *graph* of V, $\mathrm{gr}(V) = \{(x, y) : y \in V(x)\}$, is closed.

Let X be a compact convex subset of $\mathbb{R}^n$, and let

$$TX(x) = \mathrm{cl}\left(\left\{z \in \mathbb{R}^n : z = \alpha\left(y - x\right) \text{ for some } y \in X \text{ and some } \alpha \geq 0\right\}\right)$$

be the tangent cone of X at $x \in \mathbb{R}^n$. Suppose that

$$(1) \qquad V(x) \in TX(x) \text{ for all } x \in X.$$

Then standard results (see Aubin and Cellina (1984) or Smirnov (2002)) imply that from any initial point $x \in X$, the differential inclusion (DI) admits at least one *positive Carathéodory solution*: that is, an absolutely continuous mapping $\mathbf{x}\colon \mathbb{R}_+ \to X$ with $\mathbf{x}(0) = x$ satisfying $\dot{\mathbf{x}}(t) \in V(\mathbf{x}(t))$ for almost every $t \in \mathbb{R}_+$. In the following, we will suppose that (DI) is a good USC differential inclusion that satisfies (1).

Let $C(\mathbb{R}_+, X)$ denote the space of continuous X-valued maps, endowed with the topology of uniform convergence on compact sets. This topological space is metrizable with the distance D given by

$$(2) \qquad D(\mathbf{x}, \mathbf{y}) := \sum_{k \in \mathbb{N}} \tfrac{1}{2^k} \min\{1, \sup\nolimits_{t \in [0,k]} \left|\mathbf{x}(t) - \mathbf{y}(t)\right|\},$$

under which the space is complete.

Denote by $S_x \subset C(\mathbb{R}_+, X)$ the set of solutions of (DI) with initial condition $\mathbf{x}(0) = x$. The set-valued dynamical system induced by the differential inclusion will be denoted

$\Phi\colon \mathbb{R}_+ \times X \rightrightarrows X$, and defined by

$$\Phi_t(x) := \{\mathbf{x}(t)\colon \mathbf{x} \in S_x\}.$$

Finally, $S_\Phi := \cup_{x \in X} S_x$ is the set of all solution curves of (DI). Since X is compact, Barbashin's (1948) theorem implies that S_Φ is a compact subset of $C(\mathbb{R}_+, X)$.

The *translation semi-flow* $\Theta\colon \mathbb{R}_+ \times C(\mathbb{R}_+, X) \to C(\mathbb{R}_+, X)$ assigns each time $t \in \mathbb{R}_+$ and trajectory $\mathbf{x} \in C(\mathbb{R}_+, X)$ the translated trajectory $\Theta^t(\mathbf{x}) \in C(\mathbb{R}_+, X)$ defined by

$$\Theta^t(\mathbf{x})(s) = \mathbf{x}(t+s).$$

S_Φ is *invariant* under Θ, meaning that $\Theta^t(S_\Phi) = S_\Phi$ for all $t \geq 0$.

Finally, let the map $\pi_0\colon C(\mathbb{R}_+, X) \to X$ be the time 0 projection, defined by $\pi_0(\mathbf{z}) = \mathbf{z}(0)$. Evidently, π_0 is Lipschitz continuous.

2.2 Stochastic approximation processes with constant step size

2.2.1 The discrete-time case

We now introduce a class of stochastic approximation processes with constant step size defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The processes we define here arise naturally in the context of stochastic evolutionary game dynamics. We present this application in detail in Section 4. Some readers may prefer to read Section 4.1 along with the present section before proceeding to our main results in Section 3.

Let $\delta > 0$ be a positive real number. Then $\tilde{V}^\delta\colon \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is the set-valued map defined by

$$\tilde{V}^\delta(x) := \left\{z \in \mathbb{R}^n \mid \exists y \in B(x,\delta) \text{ such that } \operatorname{dist}(z, V(y)) < \delta\right\}.$$

Definition 2.2. For a sequence of values of ε approaching 0, let $\mathbf{U}^\varepsilon = \{\mathbf{U}_k^\varepsilon\}_{k=0}^\infty$ be a sequence of $\mathbb{R}^n$-valued random variables and $\{V_\varepsilon\}$ be a family of set-valued maps on $\mathbb{R}^n$. Let $\dot{x} \in V(x)$ be a good USC differential inclusion. We say that $\{\{\mathbf{X}_k^\varepsilon\}_{k=0}^\infty\}_{\varepsilon>0}$ is a *family of generalized stochastic approximation processes* (or a *family of GSAPs*) for the differential inclusion $\dot{x} \in V(x)$ if the following assumptions are satisfied:

(i) for all $k \geq 0$, we have $\mathbf{X}_k^\varepsilon \in X$,
(ii) we have the recursive formula

$$\mathbf{X}_{k+1}^\varepsilon - \mathbf{X}_k^\varepsilon - \varepsilon \mathbf{U}_{k+1}^\varepsilon \in \varepsilon V_\varepsilon(\mathbf{X}_k^\varepsilon),$$

(*iii*) for any $\delta > 0$, there exists an $\varepsilon_0 > 0$ such that for all $\varepsilon \leq \varepsilon_0$ and $x \in X$,

$$V_\varepsilon(x) \subset \tilde{V}^\delta(x),$$

(*iv*) for all $T > 0$ and for all $\alpha > 0$, we have

$$\lim_{\varepsilon \to 0} \mathbb{P}\left(\max_{k \leq \left[\frac{T}{\varepsilon}\right]} \left\| \sum_{i=0}^{k-1} \varepsilon \mathbf{U}_{i+1}^\varepsilon \right\| > \alpha \,\middle|\, \mathbf{X}_0^\varepsilon = x \right) = 0$$

uniformly in $x \in X$, where $[s]$ denotes the integer part of s.

The first three conditions of Definition 2.2 are basic, and the following propositions, which provide versions of known estimates for stochastic approximation processes,[2] describe two standard settings in which the fourth condition holds. The propositions are proved in Sections 5.1 and 5.2.

Proposition 2.3. *For a sequence of values of ε approaching 0, let $\mathbf{U}^\varepsilon = \{\mathbf{U}_k^\varepsilon\}_{k=0}^\infty$ be a martingale difference sequence, i.e. a sequence of $\mathbb{R}^n$-valued random variables such that $\mathbb{E}(\mathbf{U}_{k+1}^\varepsilon | \mathcal{F}_k^\varepsilon) = 0$ for all $k \geq 0$, where $\{\mathcal{F}_k^\varepsilon\}_{k=0}^\infty$ is the filtration induced by $\mathbf{U}^\varepsilon$. Suppose that for some $q \geq 2$,*

$$\sup_{\varepsilon > 0} \sup_k \mathbb{E}(\|\mathbf{U}_{k+1}^\varepsilon\|^q) < \infty.$$

Then $\mathbf{U}^\varepsilon$ satisfies condition (iv) of Definition 2.2.

Proposition 2.4. *For a sequence of values of ε approaching 0, let $\mathbf{U}^\varepsilon = \{\mathbf{U}_k^\varepsilon\}_{k=0}^\infty$ be a sequence of $\mathbb{R}^n$-valued random variables. Suppose there exist constants $\varepsilon_0 > 0$ and $\Gamma > 0$ such that for all $\varepsilon \leq \varepsilon_0$, $\mathbf{U}^\varepsilon$ is sub-Gaussian, in the sense that for all $\theta \in \mathbb{R}^n$,*

$$\mathbb{E}\left(\exp(\langle \theta, \mathbf{U}_{k+1}^\varepsilon \rangle) \,|\, \mathcal{F}_k\right) \leq \exp(\tfrac{1}{2}\Gamma \|\theta\|^2). \tag{3}$$

Then $\mathbf{U}^\varepsilon$ satisfies condition (iv) of Definition 2.2.

Remark 2.5. One can show that the sub-Gaussian condition (3) implies that $\mathbf{U}^\varepsilon$ is a martingale difference sequence. Thus the requirements of Proposition 2.4 are stronger than those of Proposition 2.3. The sub-Gaussian condition provides stronger control of the rate of convergence of the probability toward zero in condition (*iv*) of Definition 2.2; see Section 5.2 for details.

[2] See Métivier and Priouret (1987), Kushner and Yin (1997), and Benaïm (1999) for the classical setting with decreasing step sizes, Benaïm and Weibull (2003) for the classical setting with constant step sizes, and Benaïm et al. (2005) for the set-valued setting with decreasing step sizes.

Remark 2.6. Our results can be extended to stochastic approximation processes in which the step sizes vary over time, with the maximal length of a step approaching zero. Specifically, we can replace condition (ii) from Definition 2.2 with

$$\mathbf{X}^{\varepsilon}_{k+1} - \mathbf{X}^{\varepsilon}_{k} - \gamma^{\varepsilon}_{k+1}\mathbf{U}^{\varepsilon}_{k+1} \in \gamma^{\varepsilon}_{k+1} V_{\varepsilon}(\mathbf{X}^{\varepsilon}_{k}),$$

where $\boldsymbol{\gamma}^{\varepsilon} = \{\gamma^{\varepsilon}_{k}\}_{k\geq 1}$ is a sequence of positive real numbers such that $\sum_{k\geq 1}\gamma^{\varepsilon}_{k} = \infty$ and $\gamma^{\varepsilon}_{k} \leq \varepsilon$ for all k, and condition (iv) with

$$\lim_{\varepsilon\to 0}\mathbb{P}\left(\max_{k\leq \ell(T,\varepsilon)}\left\|\sum_{i=0}^{k-1}\gamma^{\varepsilon}_{i+1}\mathbf{U}^{\varepsilon}_{i+1}\right\| > \alpha \,\middle|\, \mathbf{X}^{\varepsilon}_{0} = x\right) = 0 \tag{4}$$

uniformly in $x \in X$, where $\ell(T,\varepsilon) = \max\{k \in \mathbb{N}\colon \sum_{i=1}^{k}\gamma^{\varepsilon}_{i} \leq T\}$. Proposition 2.3 and 2.4 remain true in this setting, but with the conclusion being that $\mathbf{U}^{\varepsilon}$ satisfies condition (4) above. Moreover, the family of processes $\{\{\mathbf{X}^{\varepsilon}_{k}\}_{k=0}^{\infty}\}_{\varepsilon>0}$ is still a perturbed solution of the differential inclusion (DI) (see Sections 3.1 and 5.3), and thus the conclusions of Theorems 3.1 and 3.5 continue to hold.

In applications to game theory and other population models, it is natural to suppose that the process $\mathbf{X}^{\varepsilon}$ undergoes ε^{-1} increments per unit of clock time. This feature is captured in the following definition of the interpolated process $\bar{\mathbf{X}}^{\varepsilon}$, which runs in continuous time. Let $\mathcal{A}^{\varepsilon}\colon X^{\mathbb{N}} \to C(\mathbb{R}_{+}, X)$ be a map that assigns each sequence $\mathbf{x} \in X^{\mathbb{N}}$ its time-rescaled affine interpolation $\mathcal{A}^{\varepsilon}(\mathbf{x}) \in C(\mathbb{R}_{+}, X)$, defined by

$$\mathcal{A}^{\varepsilon}(\mathbf{x})(t) = \bar{\mathbf{x}}(t) = \mathbf{x}_{\ell(t)} + (t - \varepsilon\ell(t))\frac{\mathbf{x}_{\ell(t)+1} - \mathbf{x}_{\ell(t)}}{\varepsilon}, \text{ where } \ell(t) = \left\lfloor \frac{t}{\varepsilon}\right\rfloor. \tag{5}$$

Applying this map to each sample path of the discrete-time process $\mathbf{X}^{\varepsilon}$ generates the interpolated process $\bar{\mathbf{X}}^{\varepsilon}$, whose sample paths are in $C(\mathbb{R}_{+}, X)$.

2.2.2 The Markovian continuous-time case

We now introduce continuous-time versions of generalized stochastic approximation processes. We define these as Markov processes on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

Definition 2.7. For a sequence of values of ε approaching 0, let $\{L^{\varepsilon}\}_{\varepsilon>0}$ be a family of operators acting on bounded functions $f\colon X \to \mathbb{R}$ according to the formula

$$L^{\varepsilon}f(x) = \frac{1}{\varepsilon}\int_{\mathbb{R}^n}\big(f(x+\varepsilon z) - f(x)\big)\,\mu^{\varepsilon}_{x}(\mathrm{d}z), \tag{6}$$

where $\{\mu_x^\varepsilon\}_{x\in X}^{\varepsilon>0}$ is a family of positive measures on $\mathbb{R}^n$ such that

(i) The function $x \mapsto \mu_x^\varepsilon(A)$ is measurable for each Borel set $A \subset \mathbb{R}^n$;

(ii) The support of μ_x^ε is contained in the set $\{z \in \mathbb{R}^n \colon x + \varepsilon z \in X\}$ as well as in some compact set independent of x and ε.

(iii) For any $\delta > 0$, there exists $\varepsilon_0 > 0$ such that for all $\varepsilon \leq \varepsilon_0$ and for all $x \in X$,

$$v^\varepsilon(x) := \int_{\mathbb{R}^n} z\, \mu_x^\varepsilon(\mathrm{d}z) \in \tilde{V}^\delta(x).$$

Let the Markov processes $\{\mathbf{Y}^\varepsilon(t)\}_{t\geq 0}^{\varepsilon>0}$ solve the martingale problems for $\{L^\varepsilon\}$ (see Ethier and Kurtz (1986)). We call this collection of processes a *family of Markov continuous-time generalized stochastic approximation processes* for the differential inclusion $\dot{x} \in V(x)$.

Remark 2.8. The following standard construction of the process $\{\mathbf{Y}^\varepsilon\}$ will prove useful in the proof of Theorem 3.2 and in Example 4.4. Let $\{\tau_i\}_{i\geq 1}$ be a sequence of independent exponential random variables with rate ε^{-1}. Define $T_0 = 0$ and $T_k = \tau_1 + \cdots + \tau_k$ for $k \geq 1$, and let $\kappa(s) = \max\{k : T_k \leq s\}$ for $s \geq 0$. Then define the Markov chain $\{\mathbf{X}_k^\varepsilon\}_{k\geq 1}$ on X with transition function

$$\mathbb{P}\left(\mathbf{X}_{k+1}^\varepsilon - \mathbf{X}_k^\varepsilon \in \varepsilon A \,|\, \mathbf{X}_k^\varepsilon = x\right) = \mu_x^\varepsilon(A)$$

for each Borel set $A \subset \mathbb{R}^n$. Then the Markov process $\{\mathbf{Y}^\varepsilon(t)\}_{t\geq 0}$ defined by

$$\mathbf{Y}^\varepsilon(t) = \mathbf{X}_{\kappa(t)}^\varepsilon$$

is a solution to the martingale problem for L^ε.

3. Results

3.1 Finite Horizon Deterministic Approximation

Our first main result shows that over any finite-time span, the interpolated process $\bar{\mathbf{X}}^\varepsilon$ closely mirrors a solution trajectory of the differential inclusion $\dot{x} \in V(x)$ with high probability.

Theorem 3.1. *Suppose that* $\{\mathbf{X}^\varepsilon\}_{\varepsilon>0}$ *is a family of GASPs. Then for any* $T > 0$ *and any* $\alpha > 0$*, we have*

$$\lim_{\varepsilon\to 0} \mathbb{P}\left(\inf_{\mathbf{x}\in S_\Phi} \sup_{0\le s\le T} \left\| \bar{\mathbf{X}}^\varepsilon(s) - \mathbf{x}(s) \right\| \ge \alpha \,\middle|\, \bar{\mathbf{X}}^\varepsilon(0) = x \right) = 0$$

uniformly in $x \in X$.

Kurtz (1970) and Benaïm and Weibull (2003) prove similar results for classical (Lipschitz) dynamical systems. In their analyses, the key tool used to control the distance between the process and the flow is Grönwall's lemma. This approach cannot be adapted to our setting. Here the mean set-valued map V may provide a multitude of directions in which the flow can proceed at each point in time. In particular, two solutions from the same initial condition can move away from each other. Instead of relying on Grönwall's lemma, we instead make use of Lemma 5.3, due to Faure and Roth (2010), which approximates solutions of the differential inclusion by approximate fixed points of a certain set-valued operator; see Section 5.3 for details.

In the continuous-time framework, the direct analogue of Theorem 3.1 is true:

Theorem 3.2. *Suppose that* $\{\mathbf{Y}^\varepsilon\}_{\varepsilon>0}$ *is a family of Markov continuous-time GASPs. Then for any* $T > 0$ *and any* $\alpha > 0$*, we have*

$$\lim_{\varepsilon\to 0} \mathbb{P}\left(\inf_{\mathbf{z}\in S_\Phi} \sup_{0\le s\le T} \|\mathbf{Y}^\varepsilon(s) - \mathbf{z}(s)\| \ge \alpha \,\middle|\, \mathbf{Y}^\varepsilon(0) = x \right) = 0$$

uniformly in $x \in X$.

Theorems 3.1 and 3.2, proved in Sections 5.3 and 5.4, are both derived from a more general result, Proposition 5.2, that establishes finite horizon deterministic approximation for so-called *perturbed solutions* of the differential inclusion (DI) (see Definition 5.1). Our analyses show that both the family of interpolated processes $\{\bar{\mathbf{X}}^\varepsilon\}_{\varepsilon>0}$ and the family of interpolated versions of the processes $\{\mathbf{Y}^\varepsilon\}_{\varepsilon>0}$ are perturbed solutions of (DI).

3.2 Limiting Stationary Distributions and Invariant Measures

In this section, we assume not only that $\{\{\mathbf{X}_k^\varepsilon\}_{k=0}^\infty\}_{\varepsilon>0}$ is a family of GSAPs for the differential inclusion $\dot{x} \in V(x)$, but also that for each $\varepsilon > 0$, $\{\mathbf{X}_k^\varepsilon\}_{k=0}$ is a Markov chain defined on probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We study the asymptotic behavior of a continuous family of probability measures associated with the resulting family of Markov chains. The main result of this section shows us that any weak limit as ε approaches zero of a collection of invariant probability measures $\{\mu^\varepsilon\}_{\varepsilon>0}$ of the Markov chains $\{\mathbf{X}^\varepsilon\}_{\varepsilon>0}$ is itself an semi-invariant measure of the set-valued dynamical system Φ induced by $\dot{x} \in V(x)$.

We recall some useful definitions and notations. Let M be a metric space, and let $\mathcal{P}(M)$ be the space of Borel probability measures on M. The support of a measure ν in $\mathcal{P}(M)$, denoted support(ν), is the smallest closed subset of measure 1. If M' is also a metric space and $f\colon M \to M'$ is Borel measurable, then the induced map $f^*\colon \mathcal{P}(M) \to \mathcal{P}(M')$ associates with $\nu \in \mathcal{P}(M)$ the measure $f^*(\nu) \in \mathcal{P}(M')$ defined by

$$f^*(\nu)(B) = \nu(f^{-1}(B))$$

for all Borel sets B in M'. If $\theta\colon \mathbb{R}_+ \times M \to M$ is a semi-flow on M, a measure $\nu \in \mathcal{P}(M)$ is called *θ-invariant* if $\nu((\theta_t)^{-1}(A)) = \nu(A)$ for all Borel sets $A \in M$ and $t \in \mathbb{R}_+$. In the following we will use this equivalent definition (see e.g. Theorem 6.8 in Walters (1982)): ν is θ-invariant if for all bounded continuous functions $f\colon M \to \mathbb{R}$ and all $t > 0$,

$$\int_M f(x)\mathrm{d}\nu(x) = \int_M f \circ \theta_t(x)\mathrm{d}\nu(x). \tag{7}$$

We now state the definition of semi-invariant measures for set-valued dynamical systems. The definition uses the translation semi-flow Θ introduced in Section 2.1.

Definition 3.3. A probability measure μ on X is an *semi-invariant measure* for the set-valued dynamical system Φ if there exists a probability measure ν on $C(\mathbb{R}_+, X)$ such that

(i) support(ν) $\subset S_\Phi$,
(ii) ν is Θ-invariant, and
(iii) $\pi_0^*(\nu) = \mu$.

In words, condition (iii) requires that ν have time 0 marginal distribution μ. We use the prefix "semi" because the measure ν places its mass on the set S_Φ of forward solutions, rather than on the smaller set of entire solutions (i.e., solutions defined on all $\mathbb{R}$) as in the definition of *invariant measure* from Faure and Roth (2011). Every invariant measure is an semi-invariant measure, but whether the converse statement is also true is an open question (see Remark 2.8 in Faure and Roth (2011)).

While the definition of a semi-invariant measure on X for a set-valued dynamical system Φ is specified in terms of shift-invariant measures on the space $C(\mathbb{R}_+, X)$ of continuous-time trajectories through X, the stochastic approximation processes introduced in Section 2.2.1 run in discrete time. It is therefore useful to have a sufficient condition for invariance on X that is stated in terms of shift-invariant measures on the sequence space $X^{\mathbb{N}}$. Proposition 3.4 provides such a condition. To state it, we let $\tilde{\Theta}$ be the left shift on $X^{\mathbb{N}}$, and let $\tilde{\pi}_0$

be the map which associates with each sequence in $X^{\mathbb{N}}$ its first coordinate. In addition, we let $\mathcal{N}^{\delta}(A)$ denote the δ-neighborhood of a subset $A \subset C(\mathbb{R}_+, X)$ for the distance D defined in (2).

Proposition 3.4. *Let μ be a probability measure on X. If there exist a vanishing sequence $\{\varepsilon_n\}_{n\geq 0}$ of real numbers and a sequence $\{\tilde{\nu}_n\}_{n\geq 0}$ of probability measures on $X^{\mathbb{N}}$ satisfying*

- (IM1) *for every $\delta > 0$, $\lim_{n\to\infty} \nu_n(\mathcal{N}^{\delta}(S_\Phi)) = 1$, where $\nu_n = (\mathcal{A}^{\varepsilon_n})^*(\tilde{\nu}_n)$,*
- (IM2) *$\tilde{\nu}_n$ is $\tilde{\Theta}$-invariant, and*
- (IM3) *$\lim_{n\to\infty} \tilde{\pi}_0^*(\tilde{\nu}_n) = \mu$ (in the sense of weak convergence),*

then μ is semi-invariant for the set-valued dynamical system Φ.

We present the proof of Proposition 3.4 in Section 5.5.

The following result shows us that the invariant measures of the Markov chains $\{\mathbf{X}^{\varepsilon}\}_{\varepsilon>0}$ tend to a semi-invariant measure of the set-valued dynamical system Φ. The corresponding result was proved by Benaïm (1998) for the case of classical dynamical systems. We follow a different line of proof here, since the definition of an semi-invariant measure of a set-valued dynamical system forces us to work in a larger space.

Theorem 3.5. *For each $\varepsilon > 0$, let μ^{ε} be an invariant probability measure of the Markov chain $\mathbf{X}^{\varepsilon}$. Let μ be a limit point of $\{\mu^{\varepsilon}\}_{\varepsilon>0}$ in the topology of weak convergence, and let Φ be the set-valued dynamical system induced by* (DI)*. Then μ is an semi-invariant measure of the set-valued dynamical system Φ.*

We present the proof of Theorem 3.2 in Section 5.6.

In the continuous-time framework, the invariant measures of the Markov jump process $\mathbf{Y}^{\varepsilon}$ are the invariant measures of the Markov chain $\mathbf{X}^{\varepsilon}$ defined in Remark 2.8. Since this Markov chain satisfies the conditions of Theorem 3.5, the following result is immediate.

Theorem 3.6. *Let $\{\mathbf{Y}^{\varepsilon}\}_{\varepsilon>0}$ be a family of Markov continuous-time GSAPs. Let μ^{ε} be an invariant probability measure of $\mathbf{Y}^{\varepsilon}$, and let μ be a limit point of $\{\mu^{\varepsilon}\}_{\varepsilon>0}$ in the topology of weak convergence. Then μ is an semi-invariant measure of the set-valued dynamical system Φ.*

It is generally difficult to compute the collection of semi-invariant measures of a set-valued dynamical system Φ. It is therefore of interest to derive restrictions on the limit measure μ that are easier to evaluate than those from Theorems 3.5 and 3.6. In order to do so, we introduce the notion of the Birkhoff center of Φ.

Recall that the *limit set* $L(\mathbf{z})$ of a solution $\mathbf{z}$ of (DI) is defined by

$$L(\mathbf{z}) = \bigcap_{t\geq 0} \mathrm{cl}(\mathbf{z}([t,\infty[)),$$

and the limit set $L(x)$ of a point $x \in X$ is defined by

$$L(x) = \bigcup_{\mathbf{z} \in S_x} L(\mathbf{z}).$$

We define the set of *recurrent points* of Φ by

$$R^L_\Phi = \{x \in X \colon x \in L(x)\}.$$

The closure $\mathrm{cl}(R^L_\Phi)$ of the set of recurrent points is called the *Birkhoff center* of Φ and is denoted by $BC(\Phi)$.

Remark 3.7. The limit set $L(x)$ may be equal to or strictly contained in the *ω-limit set* of the point x, defined by

$$\begin{aligned}
\omega(x) &= \bigcap_{t \geq 0} \mathrm{cl}(\Phi_{[t,\infty[}) \\
&= \{y \in X \mid \lim_{n\to\infty} \mathbf{z}^n(t_n) = y \text{ for some } \{t_n\} \uparrow \infty \text{ and } \{\mathbf{z}^n\} \subset S_\Phi \text{ with } \mathbf{z}^n(0) = x\}.
\end{aligned}$$

Unlike the definition of $L(x)$, the definition of $\omega(x)$ allows one to move between different solution trajectories starting from x as one considers later moments in time. For instance, if $V \colon [0,1] \rightrightarrows \mathbb{R}$ is defined by $V(0) = [0,1]$ and $V(x) = 1 - x$ for $x \in (0,1]$, then $L(0) = \{0,1\}$ but $\omega(0) = [0,1]$. Under our definition of recurrence based on $L(x)$, point x is recurrent if there exists at least one solution curve starting from x whose limit set contains x. ◆

The following theorem is a version of the Poincaré recurrence theorem for standard set-valued dynamical systems for semi-invariant measures.

Theorem 3.8. *Let μ be an semi-invariant measure for Φ, then*

$$\mu(BC(\Phi)) = 1.$$

The proof of Theorem 3.8, which is directly inspired by the one of Faure and Roth (2011) for invariant measures, is presented in Section 5.7.

The following result follows immediately from the previous theorems.

Corollary 3.9. *Under the conditions of Theorem 3.5 or Theorem 3.6, the support of μ is contained in the Birkhoff center of Φ.*

4. Applications to Game Theory

We now define classes of Markov process from game theory that fit within the framework introduced above. Convergence results for these stochastic evolutionary game dynamics are presented in Section 4.4. For background on the material presented here, see Sandholm (2010).

4.1 Population Games and Stochastic Evolutionary Game Dynamics

4.1.1 Population games

Consider a population of $N \in \mathbb{N}$ agents, each of whom chooses a strategy from the set $S = \{1, \ldots, n\}$. The distribution of strategies chosen by the members of this population is represented by a point in the simplex $X = \{x \in \mathbb{R}^n_+ \colon \sum_i x_i = 1\}$, or, more precisely, in the uniform grid $\mathcal{X}^N = X \cap \frac{1}{N}\mathbb{Z}^n$. We typically refer to such a point as a *population state*.

A *population game* is defined by a continuous function $F \colon X \to \mathbb{R}^n$. Here $F_i(x)$ represents the payoff to strategy $i \in S$ when the population state is x. One could allow payoffs to depend directly on the population size N, but to keep the notation manageable we do not do so here.

Since we are concerned with behavior in large populations, the relevant notion of equilibrium for F is that for the large N limit. We therefore call state $x \in X$ a *Nash equilibrium* of F if $x_i > 0$ implies that $F_i(x) \geq F_j(x)$ for all $j \in S$. In other words, at a Nash equilibrium, no (infinitesimal) agent could improve his payoffs by switching strategies.

4.1.2 Revision protocols

In our model of game dynamics, each of the N agents occasionally receives opportunities to switch strategies. The arrivals of these opportunities across agents are such that the expected number of revision opportunities that each agent receives in a unit of clock time is 1. Three specifications of the arrival process are introduced below. Each variant of the model generates a Markov process on the grid $\mathcal{X}^N$; for a more general extension, see Remark 4.8.

The behavior of an agent playing game F after receiving a revision opportunity is governed by a *revision protocol*, a pair of functions $(r^{F,N}, \sigma^{F,N})$ with $r^{F,N} \colon \mathcal{X}^N \to [0,1]^n$ and $\sigma^{F,N} \colon \mathcal{X}^N \to X^n$. When a current i player receives a revision opportunity, he considers switching strategies with probability $r^{F,N}_i(x)$; if he does consider a switch, he chooses strategy j with probability $\sigma^{F,N}_{i,j}(x)$. (To be clear, the choice probabilities of an i player at state x are given by the probability vector $\sigma^{F,N}_i(x) \in X$.) Evidently, an agent's decisions only depend on the revision protocol used in the game, the current population state, and his own current strategy, and do not otherwise depend on the agent's identity.

The dependence of $r^{F,N}$ and $\sigma^{F,N}$ on N allows for vanishing finite-population effects—for instance, the effects of sampling without replacement. We ensure that any such effects are small by requiring the existence of a limit protocol (r^F, σ^F) with $r^F\colon X \to [0,1]^n$ and $\sigma^F\colon X \to X^n$ that satisfies

$$\lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| r^{F,N}(x) - r^F(x) \right| = 0 \quad \text{and} \quad \lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| \sigma^{F,N}(x) - \sigma^F(x) \right| = 0.$$

4.1.3 Mean dynamics

The *mean dynamic* (or *mean field*) $v\colon X \to \mathbb{R}^n$ induced by the population game F and revision protocol (r^F, σ^F) represents the expected increment per time unit in the proportions of agents playing each strategy. The mean dynamic is defined by

$$(8) \qquad v_i(x) = \sum_{j\in S} x_j\, r_j^F(x)\, \sigma_{j,i}^F(x) - x_i\, r_i^F(x).$$

The initial sum in (8) represents the expected inflow to strategy i from other strategies, while the second term represents the expected outflow from strategy i to other strategies.

It is easy to verify that $v(x) \in TX(x)$ for all $x \in X$ for any choice of game and revision protocol. While earlier papers (e.g., Benaïm and Weibull (2003), Sandholm (2003)) focus on revision protocols under which the mean dynamic v is Lipschitz continuous, the examples introduced in Section 4.2 lead to mean dynamics that are discontinuous selections from a good USC differential inclusion.

4.1.4 Arrivals of revision opportunities and stochastic approximation processes

We now introduce three distinct models of the arrivals of revision opportunities:

I. A simple discrete-time model: Each unit of clock time is divided into subintervals of length $\frac{1}{N}$. At the end of each subinterval, one of the N agents is chosen at random to receive a revision opportunity, with these choices being independent over time.

II. An alternative discrete-time model: Each unit of clock time is again divided into subintervals of length $\frac{1}{N}$. At the end of each subinterval, each of the N agents receives a revision opportunity with probability $\frac{1}{N}$, with these draws being independent across agents and over time. Thus, the number of opportunities arriving at the end of each subinterval follows a *binomial*$(N, \frac{1}{N})$ distribution.

III. A continuous-time model: Each agent receives revision opportunities according to independent, rate 1 Poisson processes.

Each of these models leads to a distinct collection of stochastic approximation processes.

Example 4.1. Discrete-time process, one revision opportunity per period. To construct the Markov chain $\{\mathbf{X}^N_k\}_{k=0}^{\infty}$ on $\mathcal{X}^N$ corresponding to Model I, let e_i denote the ith standard basis vector in $\mathbb{R}^n$, and define independent auxiliary random variables $\zeta^{N,x}$ with distributions

$$\mathbb{P}(\zeta^{N,x} = z) = \begin{cases} x_i\, r_i^{F,N}(x)\, \sigma_{i,j}^{F,N}(x) & \text{if } z = \frac{1}{N}(e_j - e_i), \\ \sum_{j\in S} x_j \left(1 - r_j^{F,N}(x) + r_j^{F,N}(x)\, \sigma_{j,j}^{F,N}(x)\right) & \text{if } z = \mathbf{0}, \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

Then define the transition law of the Markov chain $\mathbf{X}^N$ as

$$\mathbb{P}(\mathbf{X}^N_{k+1} = \mathbf{X}^N_k + z | \, \mathbf{X}^N_k = x) = \mathbb{P}(\zeta^{N,x} = z).$$

Since there are N periods per unit of clock time, the expected increment per time unit of Markov chain $\mathbf{X}^N$ is given by

$$\begin{aligned} v^N(x) &= N\, \mathbb{E}(\mathbf{X}^N_{k+1} - \mathbf{X}^N_k | \, \mathbf{X}^N_k = x) \\ &= N\, \mathbb{E}\zeta^{N,x} \\ &= \sum_{i\in S} \sum_{j\in S} x_i\, r_i^{F,N}(x)\, \sigma_{i,j}^{F,N}(x)\, (e_j - e_i) \\ &= \sum_{i\in S} e_i \left(\sum_{j\in S} x_j\, r_j^{F,N}(x)\, \sigma_{j,i}^{F,N}(x) - x_i\, r_i^{F,N}(x) \right). \end{aligned} \tag{10}$$

Thus

$$\lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| v^N(x) - v(x) \right| = 0, \tag{11}$$

where $v\colon X \to \mathbb{R}^n$ is defined by (8).

Now suppose that v is a selection from $V\colon X \rightrightarrows \mathbb{R}^n$, a good USC differential inclusion. If we let $\varepsilon = \frac{1}{N}$, then conditions (*i*) and (*ii*) of Definition 2.2 are clearly satisfied with

$$\mathbf{U}^\varepsilon_{k+1} = \mathbf{U}^N_{k+1} = N\left(\mathbf{X}^N_{k+1} - \mathbf{X}^N_k - \mathbb{E}(\mathbf{X}^N_{k+1} - \mathbf{X}^N_k | \, \mathcal{F}^N_k)\right), \tag{12}$$

and the uniform convergence in (11) implies condition (*iii*). Finally, since each $\mathbf{U}^N$ is a martingale difference sequence, and since the collection $\{\mathbf{U}^N\}_{N=N_0}^{\infty}$ is uniformly bounded, Propsition 2.3 implies that condition (*iv*) of Definition 2.2 is satisfied as well. Thus, the collection $\{\{\mathbf{X}_k^N\}_{k=0}^{\infty}\}_{N=N_0}^{\infty}$ is a family of GSAPs. ◆

Example 4.2. Discrete-time process, random number of revision opportunities per period. To construct the Markov chain corresponding to Model II, first define the auxiliary random variables $\xi_{i,\ell}^{N,x}$, $\ell \in \{1, \dots, Nx_i\}$, with distributions

$$\mathbb{P}(\xi_{i,\ell}^{N,x} = z) = \begin{cases} r_i^{F,N}(x)\, \sigma_{i,j}^{F,N}(x) & \text{if } z = \frac{1}{N}(e_j - e_i), \\ 1 - r_i^{F,N}(x) + r_i^{F,N}(x)\, \sigma_{i,i}^{F,N}(x) & \text{if } z = \mathbf{0}, \\ 0 & \text{otherwise,} \end{cases}$$

and define $R_i^{N,x}$ to be a *binomial*$(Nx_i, \frac{1}{N})$ random variable, specifying all of the random variables to be mutually independent. Then define the random variables $\xi^{N,x}$ by

$$\xi^{N,x} = \sum_{i \in S} \sum_{\ell=1}^{R_i^{N,x}} \xi_{i,\ell}^{N,x}.$$

Now define the transition law of the Markov chain $\mathbf{X}^N$ by

$$\mathbb{P}(\mathbf{X}_{k+1}^N = \mathbf{X}_k^N + z | \, \mathbf{X}_k^N = x) = \mathbb{P}(\xi^{N,x} = z).$$

The expected increment per time unit of this Markov chain is

$$\begin{aligned} w^N(x) &= N\, \mathbb{E}(\mathbf{X}_{k+1}^N - \mathbf{X}_k^N | \, \mathbf{X}_k^N = x) \\ &= N\, \mathbb{E}\xi^{N,x} \\ &= \sum_{i \in S} \mathbb{E}R_i^{N,x} \mathbb{E}\xi_{i,1}^{N,x} \\ &= \sum_{i \in S} x_i \sum_{j \in S} r_i^{F,N}(x)\, \sigma_{i,j}^{F,N}(x)\, (e_j - e_i) \\ &= v^N(x). \end{aligned} \tag{13}$$

So as before,

$$\lim_{N \to \infty} \max_{x \in \mathcal{X}^N} \left| w^N(x) - v(x) \right| = 0,$$

where $v \colon X \to \mathbb{R}^n$ is again defined by (8).

Suppose again that v is a selection from a good USC differential inclusion. If we define $\mathbf{U}^N$ as in (12), then conditions (*i*), (*ii*), and (*iii*) of Definition 2.2 are clearly satisfied, so to prove that the collection $\{\{\mathbf{X}_k^N\}_{k=0}^{\infty}\}_{N=N_0}^{\infty}$ is a family of GSAPs, we need only verify condition (*iv*). Doing so requires the following lemma, which we prove in Section 5.8:

Lemma 4.3. *There exists a constant* $M > 0$ *such that* $\sup_{k\in\mathbb{N}} \mathbb{E}(\left\|\mathbf{U}_{k+1}^N\right\|^2) \leq M < \infty$.

It follows from this lemma and Proposition 2.3 that the collection $\{\{\mathbf{X}_k^N\}_{k=0}^{\infty}\}\}_{N=N_0}^{\infty}$ is a family of GSAPs. ◆

Example 4.4. Continuous-time process. We now construct the continuous-time Markov chains $\{\mathbf{Y}_t^N\}_{t\geq 0}$ on $\mathcal{X}^N$ corresponding to Model III. To do so, we let the Markov chain $\{\mathbf{X}_k^N\}_{k=0}^{\infty}$ be defined as in Example 4.1, and we define $\mathbf{Y}^N(t) = \mathbf{X}_{\kappa(t)}^N$ as in Remark 2.8. If we let $Z = \{e_i - e_j : 0 \leq i, j \leq n\}$, then the generator L^N of the Markov process $\mathbf{Y}^N$ acts on bounded functions $f : X \to \mathbb{R}$ according to

$$L^N f(x) = N \sum_{z\in Z} \left(f(x + \tfrac{1}{N}z) - f(x)\right) \mu_x^N(z),$$

where $\mu_x^N(z) = \mathbb{P}(\zeta^{N,x} = \frac{1}{N}z)$ for the random variable $\zeta^{N,x}$ defined in equation (9).

Conditions (*i*) and (*ii*) of Definition 2.7 follow immediately from the definition of μ_x^N. Moreover, it is easy to see that $v^N(x) = \int_Z z\, \mu_x^N(\mathrm{d}z)$ is given by the last expression in display (10), and so (11) implies that v^N converges uniformly to the function v defined in (8). Thus, if v is a selection from a good USC differential inclusion, then condition (*iii*) of Definition 2.7 holds, and so the collection $\{\mathbf{Y}^N\}_{N\geq 1}$ is a family of Markov continuous-time GSAPs. ◆

4.2 Examples of Revision Protocols and their Mean Dynamics

We now present three examples of revision protocols whose mean fields are described by differential inclusions. In all of the examples, a central role is played by the *maximizer correspondence* $M : \mathbb{R}^n \rightrightarrows X$ and the *best response correspondence* $B^F : X \rightrightarrows X$, defined by

$$M(\pi) = \operatorname*{argmax}_{y\in X} \langle y, \pi\rangle \quad \text{and} \quad B^F = M \circ F.$$

Because M is convex-valued and upper semicontinuous, the mean dynamics in the next three examples are all selections from good USC differential inclusions.

To simplify the presentation, we suppose below that the revision protocol $(r^{F,N}, \sigma^{F,N}) \equiv (r^F, \sigma^F)$ does not depend on the population size. Examples 4.1, 4.2, and 4.4 show that our results for the examples coming next remain valid without this assumption, so long as the protocol converges uniformly to its limit.

Example 4.5. The best response dynamic. Suppose that an agent who receives a revision opportunity only chooses a strategy with positive probability if it is optimal at the current population state. This corresponds to revision protocols with $r_i^F \equiv 1$ and $\sigma_i^F(x) \in B^F(x)$. Note that by allowing revision probabilities to depend on the agent's current strategy i, one can allow players of different strategies to break ties in different ways; for instance, one can suppose that an agent currently playing a best response does not switch to an alternate best response. For a further generalization, see Remark 4.8 below, and for the implications for stochastic approximation, see the examples in Section 4.3.

For any revision protocol of the form described above, the mean field v is a selection from the differential inclusion

$$\text{(14)} \qquad \dot{x} \in V(x) = B^F(x) - x.$$

Equation (14) is the *best response dynamic* of Gilboa and Matsui (1991) and Hofbauer (1995).[3] ◆

Example 4.6. Tempered best response dynamics. It may be more realistic to assume that an agent who receives a revision opportunity becomes less likely to bother to revise as his current payoff becomes closer to the optimal payoff. This is achieved by letting $r_i^F(x) = \rho(\widehat{F}_i(x))$, where $\widehat{F}_i(x) = \max_j F_j(x) - F_i(x)$ is the *payoff deficit* of strategy i at population state x, and $\rho\colon \mathbb{R}_+ \to [0,1]$ is a nondecreasing, Lipschitz continuous function satisfying $\rho(0) = 0$ and $\rho(s) > 0$ when $s > 0$. If we again let $\sigma_i^F(x) \in B^F(x)$, then the mean field v is a selection from the differential inclusion

$$\text{(15)} \qquad \dot{x} \in V(x) = \sum_{i \in S} x_i\, \rho(\widehat{F}_i(x))\, (B^F(x) - e_i),$$

Equation (15) defines the *tempered best response dynamic* of Zusai (2011). ◆

Example 4.7. Sampling best response dynamics. If agents cannot directly observe the population state, they can estimate this state by taking a finite sample from the population, and using the empirical distribution of strategies in the sample as an estimate of the actual population state. Dynamics based on this idea are introduced by Oyama et al. (2012).

Let $\mathbb{Z}_+^{n,k} = \{z \in \mathbb{Z}_+^n \colon \sum_{i \in S} z_i = k\}$ be the set of possible outcomes of samples of size k,

[3] In some games, the best response protocol may be seen as too permissive, as it allows agents to switch to strategies that are optimal at the current state, but that fail to be optimal at any nearby states. To remedy this, Balkenborg et al. (2011) suppose that players only switch to strategies that are not only best responses to the current state, but also unique best responses at certain states arbitrarily close to the current state. This refinement of the best response correspondence leads to a refinement of (14) called the *refined best response dynamic*. See Balkenborg et al. (2011) for further discussion and analysis.

and define the *k-sampling best response correspondence* $B^{F,k}\colon X \rightrightarrows X$ by

$$B^{F,k}(x) = \sum_{z \in \mathbb{Z}^{n,k}_+} \binom{k}{z_1 \; \ldots \; z_n} \left(x_1^{z_1} \cdots x_n^{z_n}\right) B^F\left(\frac{1}{k}z\right). \tag{16}$$

If an agent who receives a revision opportunity takes a sample of size k from the population and switches to a best response to its empirical distribution, we obtain a revision protocol with $r_i^F \equiv 1$ and $\sigma_i^F(x) \in B^{F,k}(x)$. The resulting mean field is a selection from the differential inclusion

$$\dot{x} \in V(x) \in B^{F,k}(x) - x, \tag{17}$$

which is known as the *k-sampling best response dynamic*.

More generally, one can allow the size of a revising agent's sample to itself be random. If the sample size is determined by the probability distribution $\lambda = \{\lambda_k\}_{k=1}^{\infty}$ on $\mathbb{N}$, so that $r_i^F \equiv 1$ and $\sigma_i^F(x) \in \sum_{k\in\mathbb{N}} \lambda_k B^{F,k}(x)$, the mean dynamic is a selection from the *λ-sampling response dynamic*,

$$\dot{x} \in V(x) = \sum_{k\in\mathbb{N}} \lambda_k B^{F,k}(x) - x. \quad \blacklozenge \tag{18}$$

Remark 4.8. Since the processes introduced in Section 4.1 are Markov processes, the maps (10) and (13) describing the expected increment per time unit as a function of the state x are single-valued. Thus the mean fields in Examples 4.5–4.7 are (single-valued) selections from good USC differential inclusions.

Since our discrete-time deterministic approximation result, Theorem 3.1, does not require the underlying processes to be Markov, it can also be applied to more general processes. Let $(R^{F,N}, \Sigma^{F,N})$ with $R^{F,N}\colon \mathcal{X}^N \rightrightarrows [0,1]^n$ and $\Sigma^{F,N}\colon \mathcal{X}^N \rightrightarrows X^n$ be nonempty, convex-valued, and upper semicontinuous correspondences. When a current i player receives a revision opportunity, he considers switching strategies with probability given by a random variable $r_i^{F,N}(x)$ that takes values in $R^{F,N}(x)$; if he does consider a switch, his choice probabilities are given by a random variable $\sigma_i^{F,N}(x)$ that takes values in $\Sigma^{F,N}(x)$. We make no assumptions about the dependence of these random variables on the history of the process.

Let $\mathbf{X}^N$ be the resulting stochastic process, with revision opportunities arriving as in Example 4.1 or 4.2, and suppose that $R^{F,N}$ and $\Sigma^{F,N}$ converge to $R^F\colon X \rightrightarrows [0,1]^n$ and

$\Sigma^F\colon X \rightrightarrows X^n$ uniformly in the Hausdorff topology. If

$$\dot{x} \in V(x) = \sum_{i \in S} e_i \left(\sum_{j \in S} x_j \, R_j^F(x) \, \Sigma_{j,i}^F(x) - x_i \, R_i^F(x) \right). \tag{19}$$

is a good USC differential inclusion, as is true under the protocols from Examples 4.5–4.7, then it is easy to verify that $\{\{\mathbf{X}_k^N\}_{k=0}^\infty\}_{N=N_0}^\infty$ is a family of GSAPs for $\dot{x} = V(x)$. ◆

4.3 Which Solution Trajectories Approximate Sample Paths?

When the mean field is Lipschitz, it admits a unique solution trajectory from every initial condition, but when it is a differential inclusion, there may be multiple solutions. While sometimes more than one such solution can approximate sample paths of the stochastic process, it need not be the case that all solutions do so. Thus, if the differential inclusion is understood as an approximation of the underlying stochastic process, not all of its solutions need be relevant. We show in the following examples that solutions sometimes can be ruled out by basic considerations about the stochastic processes. For definiteness, we assume below that revision opportunities arrive in discrete time as specified in Example 4.1 or 4.2.

Example 4.9. When considering two-strategy games, it is convenient to let the strategy set be $S = \{0, 1\}$, to let $\chi = x_1$ denote the mass of agents using strategy 1, and to abuse notation by using this $\chi \in \mathcal{X} = [0, 1]$ as our state variable. Doing so now, we define the coordination game $F\colon \mathcal{X} \to \mathbb{R}^2$ by

$$F(\chi) = \begin{pmatrix} \chi^* & 0 \\ 0 & 1-\chi^* \end{pmatrix} \begin{pmatrix} 1-\chi \\ \chi \end{pmatrix} = \begin{pmatrix} (1-\chi)\,\chi^* \\ \chi\,(1-\chi^*) \end{pmatrix}.$$

where $\chi^* \in (0, 1)$. This game has three Nash equilibria, $\chi = 0$, $\chi = 1$, and $\chi = \chi^*$.

Suppose that agents follow a best response protocol as introduced in Example 4.5, and let $\mathbf{X}^N$ denote the resulting N-agent stochastic process. By Theorem 3.1, finite horizon behavior of sample paths must be approximated by solutions to the best response dynamic for F, here given by the differential inclusion

$$\dot{\chi} = B^F(\chi) - \chi = \begin{cases} \{-\chi\} & \text{if } \chi < \chi^*, \\ [-\chi, 1-\chi] & \text{if } \chi = \chi^*, \\ \{1-\chi\} & \text{if } \chi > \chi^*. \end{cases} \tag{20}$$

From each initial condition $\chi \neq \chi^*$, this differential inclusion admits a unique solution, which converges to state 0 or state 1 according to whether χ is less than or greater than χ^*:

$$\chi < \chi^* \text{ implies that } \boldsymbol{x}(t) = \mathrm{e}^{-t}\chi, \tag{21}$$

$$\chi > \chi^* \text{ implies that } \boldsymbol{x}(t) = 1 - \mathrm{e}^{-t}(1-\chi). \tag{22}$$

From initial condition χ^* itself, (20) admits infinitely many solutions: there is the stationary solution $\boldsymbol{x}(t) \equiv \chi^*$, solutions of forms (21) and (22) that proceed directly to a pure equilibrium, and solutions that stay at χ^* for a finite amount of time before proceeding to a pure equilibrium.

Which solutions of (20) approximate sample paths of the process $\mathbf{X}^N$? If χ^* is not a state in the set $\mathcal{X}^N = \{0, \frac{1}{N}, \ldots, 1\}$ where the process $\mathbf{X}^N$ runs, the only solutions of (20) are relevant for stochastic approximation are those of forms (21) and (22). Of course, this is necessarily the case whenever χ^* is irrational.

If instead χ^* is in $\mathcal{X}$ and is the initial condition of $\mathbf{X}^N$, then which solutions of (20) are relevant depends on just how $\mathbf{X}^N$ is specified. If one assumes that indifferent agents do not switch strategies, then $\mathbf{X}^N$ never leaves state χ^*, so only the stationary solution of (20) is relevant. If any other Markovian specification is used, so that there is a positive probability that the initial increment of $\mathbf{X}^N$ is not null, then for large N only the solutions of (20) that immediately leave χ^* can approximate sample paths of $\mathbf{X}^N$. Thus, the solutions of (20) that leave χ^* after some delay are only relevant if a non-Markovian specification of $\mathbf{X}^N$ is followed, as described in Remark 4.8.[4] ◆

Example 4.10. Now suppose that agents follow a best response protocol when playing the following game (Zeeman (1980), Hofbauer (1995)):

$$F(x) = \begin{pmatrix} 0 & 6 & -4 \\ -3 & 0 & 5 \\ -1 & 3 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}.$$

The phase diagram for the best response dynamic in this game is presented in Figure Figure 1(i). Solution trajectories from most initial conditions are unique and converge to the Nash equilibrium e_1. However, from Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, there is a

[4] If the population game is intended to model random matching in a normal form game, then adjustments should be made to F and σ to account for the absence of self-matching. If this is done, then the only solutions of (20) starting from χ^* that are relevant for stochastic approximation are those that immediately leave χ^*. See Sandholm (2010, Sec. 11.4 and 12.5.3) for related observations.

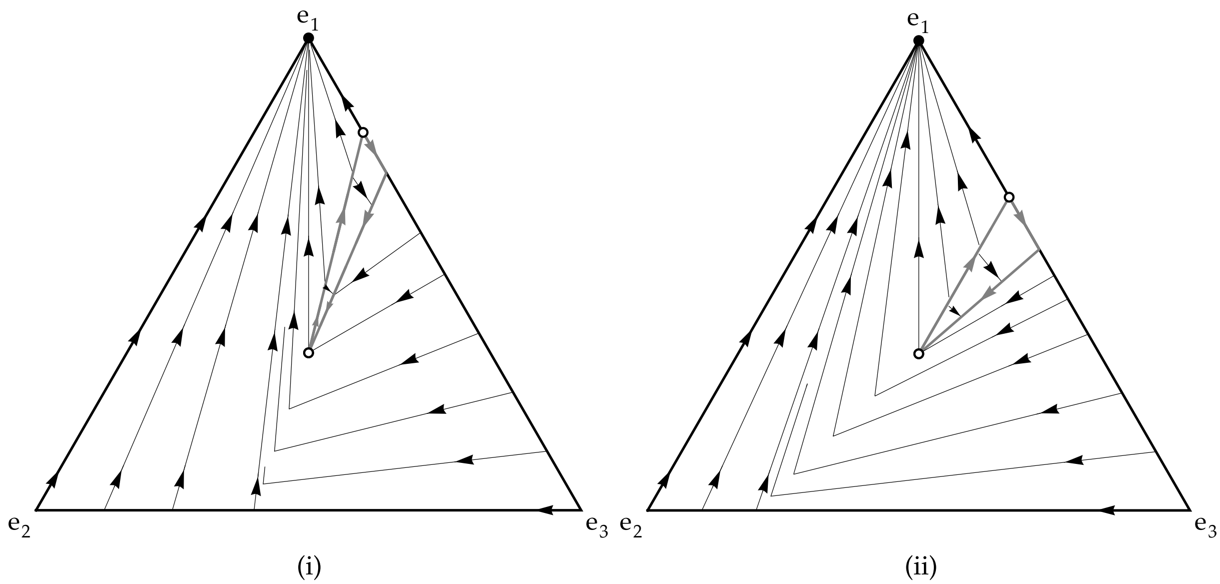

Figure 1: The best response dynamic in two variations on Zeeman's game.

stationary solution at x^*, as well as solutions that head toward e_1, possibly after some delay. Other solutions head toward the Nash equilibrium $y^* = (\frac{4}{5}, 0, \frac{1}{5})$. Some of these converge to y^*; others leave segment x^*y^* before reaching y^*. Of those that leave, some head to e_1, while others head toward e_3 and then return to x^*. If x^* is revisited, any of the behaviors just described can occur again.

Suppose that the initial state x of the process $\mathbf{X}^N$ lies in the interior of segment x^*y^*. The possible increments in the state are of the form $\frac{1}{N}(e_j - e_i)$. In terms of the figure, the direction of any such increment is parallel to a face of the simplex, and in particular is not parallel to segment x^*y^*. Therefore, any non-null transition must take the state off of the segment. This implies that the only solutions to the best response dynamic starting in the interior of segment x^*y^* that can be relevant for stochastic approximation are those that immediately leave the segment.

The game

$$G(x) = \begin{pmatrix} 0 & 3 & -2 \\ -9 & 0 & 10 \\ -1 & 2 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}$$

has a similar incentive structure to F, but in this game the segment between the Nash equilibria $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and $y^* = (\frac{2}{3}, 0, \frac{1}{3})$ is parallel to the e_2e_1 face of the simplex (Figure 1(ii)). Now suppose that the process $\mathbf{X}^N$ begins at some state x in the interior this segment.

(Such a state is in $\mathcal{X}^N$ if and only if N is a multiple of 3). If one assumes that agents playing an optimal strategy do not switch, and that strategy 2 players who are indifferent between strategies 1 and 3 always switch to strategy 1, then the process $\mathbf{X}^N$ will stay on segment x^*y^*, sometimes remaining in place and sometimes moving toward y^*. So in this case, the only approximating solution of the best response dynamic for G is the unique one that starts at state x and converges to equilibrium y^*. ◆

4.4 Convergence of Stochastic Evolutionary Game Dynamics

We now use our main results to understand the finite-horizon and infinite-horizon behavior of the processes introduced in Section 4.2 in certain classes of well-behaved games. Analyses of this sort were performed by Hofbauer and Sandholm (2007) for processes based on perturbed best responses, under which agents choose optimally after the payoffs to each of their strategies is subject to a random perturbation. These random perturbations ensure that the relevant mean dynamics are smooth, allowing results on stochastic approximation in classical contexts to be applied. The main results in the present paper allow us to understand the behavior of processes based on exact best responses.

4.4.1 Notions of convergence

Following Hofbauer and Sandholm (2007), we introduce two notions of convergence for stochastic evolutionary game dynamics, focusing for convenience on the discrete time processes $\mathbf{X}^N$ from Examples 4.1 and 4.2. We say that the processes $\mathbf{X}^N$ *converge in the medium run from initial conditions in A* to the closed set $C \subset X$ if for each $x \in A$ and $\alpha > 0$, there is a time $T = T(x)$ such that for all $U \geq T$,

$$\lim_{N\to\infty} \mathbb{P}\left(\sup_{t\in[T,U]} \operatorname{dist}(\mathbf{X}^N_{[Nt]}, C) \geq \alpha \,\middle|\, \mathbf{X}^N_0 = x \right) = 0.$$

The time index in $\mathbf{X}^N_{[Nt]}$ accounts for the fact that each period in the discrete time process lasts $\frac{1}{N}$ units of clock time. If the set of initial conditions is not specified, it is understood to be the entire simplex X.

We say that the processes $\mathbf{X}^N$ *converge in the long run* to the closed set $C \subset X$ if for any sequence of invariant measures $\{\mu^N\}_{N=N_0}^{\infty}$ of these processes, and for each open set O containing C, we have that $\lim_{N\to\infty} \mu^N(O) = 1$.

4.4.2 Best response protocols in potential games

The population game $F\colon X \to \mathbb{R}^n$ is a *potential game* (Monderer and Shapley (1996), Sandholm (2001b, 2009)) if it admits a C^1 *potential function* $f\colon X \to \mathbb{R}$, meaning that

$$\mathbf{\Phi} F(x) = \nabla f(x) \text{ for all } x \in X,$$

where $\mathbf{\Phi} = I - \frac{1}{n}\mathbf{1}\mathbf{1}' \in \mathbb{R}^{n\times n}$ is the orthogonal projection of $\mathbb{R}^n$ onto $TX = \{z \in \mathbb{R}^n : z'\mathbf{1} = 0\}$, the tangent space of the simplex X. Applications of potential games include models of genetic competition, firm competition, network congestion, externality pricing, and evolutionary implementation.

The Nash equilibria of a potential game are the states that (along with appropriate multipliers) satisfy the Kuhn-Tucker first order conditions for maximizing the potential function f on the set X. Moreover, all solutions of a wide range of deterministic evolutionary dynamics ascend the potential function and converge to connected sets of Nash equilibria (Sandholm (2001b, 2010)). Combining these ideas with the analysis above, we obtain

Theorem 4.11. *Consider evolution under any best response protocol for the potential game F. Then the processes $\{\mathbf{X}^N\}_{N=N_0}^{\infty}$ converge in the medium run and in the long run to the set of Nash equilibria of F.*

The proof of this result proceeds as follows. Sandholm (2010, Theorem 7.1.3) (also see Hofbauer (1995)) shows that the potential function f is a strict Lyapunov function for the best response dynamic (14): along every solution $\mathbf{x}$ of (14), we have $\frac{d}{dt} f(\mathbf{x}(t)) \geq 0$ at all points of differentiability of $\mathbf{x}$, with equality holding only at Nash equilibria of F. Since (14) is a good USC differential inclusion, it follows that the limit set of every solution trajectory of (14) is contained in the set of Nash equilibria of F (cf. Sandholm (2010, Theorem 7.B.4)). Since the mean field of any best response protocol is a selection from (14), convergence in the medium run follows from Theorem 3.1, and convergence in the long run from Theorem 3.5 and Corollary 3.9.

4.4.3 Best response protocols in contractive games

The population game $F\colon X \to \mathbb{R}^n$ is a *contractive game*[5] (Hofbauer and Sandholm (2009)) if

$$(y - x)'(F(y) - F(x)) \leq 0 \text{ for all } x, y \in X. \tag{23}$$

[5] Hofbauer and Sandholm (2009) and Sandholm (2010) call these games *stable games*. The term *negative semidefinite game* is sometimes used in the literature as well.

Contractive games include zero-sum games, games with an interior ESS (Maynard Smith and Price (1973)), wars of attrition, and perturbed concave potential games as instances.

It can be shown that the set of Nash equilibria of any contractive game is convex. The term "contractive game" refers to the game's out-of-equilibrium properties: if one considers an adjustment process from any two initial states, assuming that the trajectories "follow the payoff vectors" to the greatest possible extent while remaining in the simplex, then the two trajectories move (weakly) closer together over time.[6] In fact, it can be shown that in stable games, solutions of a range of evolutionary dynamics converge to Nash equilibrium from all initial conditions. However, unlike in the case of potential games, there is no single function serves as a Lyapunov function for all dynamics; instead, a distinct Lyapunov function must be constructed for each dynamic under consideration.

In the case of the best response dynamic, Hofbauer (1995, 2000) and Hofbauer and Sandholm (2009) show that the function

$$\Lambda(x) = \max_{y \in X} (y - x)' F(x)$$

serves as a strict Lyapunov function. Combining this fact with our earlier analyses yields

Theorem 4.12. *Consider evolution under any best response protocol for game F, and assume that F is a stable game. Then the processes $\{\mathbf{X}^N\}_{N=N_0}^{\infty}$ converge in the medium run and in the long run to the set of Nash equilibria of F.*

Remark 4.13. Zusai (2011) shows that in potential games, the potential function serves as a strict Lyapunov function for any tempered best response dynamic. Zusai (2011) also constructs Lyapunov functions for stable games for all such dynamics. By combining these results with our earlier analyses, one can establish versions of Theorems 4.11 and 4.12 for evolutionary processes derived from tempered best response protocols.

4.4.4 Sampling best response protocols in games with iterated p-dominant equilibria

We conclude the paper by presenting a convergence result for sampling best response processes. We say that strategy i is *p-dominant* in game F (Morris et al. (1995)) if i is the unique best response whenever it is used by at least fraction p of the population:

$$B^F(x) = \{e_i\} \text{ for all } x \in X \text{ with } x_i \geq p.$$

[6] The dynamic that "follows the payoff vectors" in this sense is the projection dynamic of Nagurney and Zhang (1997) (see also Lahkar and Sandholm (2008)). For a detailed presentation of the contraction property, see Sandholm (2012).

Evidently, lowering the value of p strengthens the p-dominance criterion, and in any population game F, at most one strategy can be p-dominant with $p \geq \frac{1}{2}$.

To define stronger solution concepts, we call a nonempty set of strategies $S^* \subset S$ a *p-best response set* of F if

$$B^F(x) \subset S^* \text{ for all } x \in X \text{ with } \sum_{i \in S^*} x_i \geq p.$$

Thus S^* is a p-best response set if whenever at least fraction p of the population plays actions in S^*, all best responses are themselves in S^*. We call a nonempty set of strategies $S^* \subset S$ an *iterated p-best response set* of F (Tercieux (2006)) if there exists a sequence $S^0, S^1, \ldots, S^m$ with $S = S^0 \supset S^1 \supset \cdots \supset S^m = S^*$ such that S^ℓ is a p-best response set in $F|_{S^{\ell-1}}$, the restricted version of F in which only strategies in $S^{\ell-1}$ are allowed, for each $\ell = 1, \ldots, m$. Strategy $i \in S$ is an *iterated p-dominant equilibrium* of F if $\{i\}$ is an iterated p-best response set of F.

To link these concepts with sampling best response processes, we say that the probability distribution $\lambda = \{\lambda_k\}_{k=1}^\infty$ on $\mathbb{N}$ is *k-good* if

$$\sum_{\ell=1}^{k} \lambda_\ell \left(1 - (\tfrac{k-1}{k})^\ell\right) > \tfrac{1}{k}.$$

Oyama et al. (2012) show that if distribution λ is k-good, and if strategy i is an iterated $\frac{1}{k}$-dominant equilibrium that is initially played by a positive mass of agents, then the λ-sampling best response dynamic converges to state e_i.[7] Combining this fact with our results from Section 3, we obtain

Theorem 4.14. *Suppose that strategy i is an iterated $\frac{1}{k}$-dominant equilibrium of F, and that λ is k-good. Consider evolution under any λ-sampling best response protocol in F. Then the processes $\{\mathbf{X}^N\}_{N=N_0}^\infty$ converge in the medium run from the set $X_i^+ = \{x \in X \colon x_i > 0\}$ to the singleton $\{e_i\}$.*

Theorem 4.14 ensures that starting from most initial conditions, any λ-sampling best response process proceeds to a neighborhood of the equilibrium e_i and remains in this neighborhood for a long period of time. Unlike many of the processes studied in stochastic evolutionary game theory (e.g., those in Hofbauer and Sandholm (2007)), sampling best response processes are not irreducible: if the process reaches a state at which all agents choose the same strategy, it remains in that state forever. This suggests the possibility of strengthening the conclusion of Theorem 4.14 to that of absorption at state e_i with probability approaching 1. For work in this direction, see Sandholm (2001a).

[7] Analogous results hold for iterated k-best response sets that are not singletons.

5. Proofs

5.1 Proof of Proposition 2.3

To simplify the proof, we suppose that $q = 2$. Using the familiar Hölder inequality, the proof of the general case is almost the same.

Fix $T > 0$. By the assumption on the second moment of $\mathbf{U}^\varepsilon$, there exists a constant $M > 0$ such that $\mathbb{E}\left(\|\mathbf{U}^\varepsilon_{i+1}\|^2\right) < M$, for all $i \geq 1$. Since $\mathbf{U}^\varepsilon$ is a martingale difference sequence, Burkholder's inequality (see, e.g., Stroock (1993)) yields

$$\begin{aligned}
\mathbb{E}\left(\max_{0\leq k\leq\left[\frac{T}{\varepsilon}\right]}\left\|\sum_{i=0}^{k-1}\varepsilon\mathbf{U}^\varepsilon_{i+1}\right\|^2\right) &\leq \mathbb{E}\left(\sum_{i=0}^{\left[\frac{T}{\varepsilon}\right]-1}\varepsilon^2\left\|\mathbf{U}^\varepsilon_{i+1}\right\|^2\right)\\
&\leq \left[\frac{T}{\varepsilon}\right]Mc\varepsilon^2\\
&\leq C(T)\varepsilon,
\end{aligned}$$

where $c > 0$ is a constant given by Burkholder's inequality and $C(T) = T(cM+1)$. Markov's inequality then implies that

$$\mathbb{P}\left(\max_{k\leq\left[\frac{T}{\varepsilon}\right]}\left\|\sum_{i=0}^{k-1}\varepsilon\mathbf{U}^\varepsilon_{i+1}\right\| > \alpha\right) \leq \frac{C(T)\varepsilon}{\alpha^2},$$

so we conclude the proof by taking the limit as ε goes to zero. ■

5.2 Proof of Proposition 2.4

For $k \in \mathbb{N}$ and $\theta \in \mathbb{R}^n$, let

$$Z_k(\theta) = \exp\left(\sum_{i=0}^{k-1}\langle\theta, \varepsilon\mathbf{U}^\varepsilon_{i+1}\rangle - \tfrac{\Gamma}{2}k\varepsilon^2\|\theta\|^2\right).$$

According to (3), $\{Z_k(\theta)\}_{k=0}^\infty$ is a supermartingale. Thus, for any $\beta > 0$ and any $n \in \mathbb{N}$, Doob's supermartingale inequality implies that

$$\begin{aligned}
\mathbb{P}\left(\max_{1\leq k\leq n}\langle\theta, \sum_{i=0}^{k-1}\varepsilon\mathbf{U}^\varepsilon_{i+1}\rangle \geq \beta\right) &= \mathbb{P}\left(\max_{1\leq k\leq n}Z_k(\theta) \geq \exp(\beta - \tfrac{\Gamma}{2}\|\theta\|^2 n\varepsilon^2)\right)\\
&\leq \exp\left(\tfrac{\Gamma}{2}\|\theta\|^2 n\varepsilon^2 - \beta\right).
\end{aligned}$$

Let $e_1, \ldots, e_m$ be the canonical basis of $\mathbb{R}^n$, $\delta > 0$, and $e = \pm e_i$ for some i. Set $\beta = \frac{\delta^2}{\Gamma n \varepsilon^2}$ and $\theta = \frac{\beta}{\delta} e$. Then

$$\begin{aligned}\mathbb{P}\left(\max_{1\leq k\leq n}\langle e, \sum_{i=0}^{k-1} \varepsilon \mathbf{U}_{i+1}^{\varepsilon}\rangle \geq \delta\right) &= \mathbb{P}\left(\max_{1\leq k\leq n}\langle \theta, \sum_{i=0}^{k-1} \varepsilon \mathbf{U}_{i+1}^{\varepsilon}\rangle \geq \beta\right) \\ &\leq \exp\left(\frac{-\delta^2}{2\Gamma n \varepsilon^2}\right).\end{aligned}$$

Fixing $\gamma > 0$, and using the previous inequality with $\delta = \frac{\gamma}{m}$ and $n = \left[\frac{T}{\varepsilon}\right]$, we conclude that

$$\begin{aligned}\mathbb{P}\left(\max_{1\leq k\leq \left[\frac{T}{\varepsilon}\right]}\left\|\sum_{i=0}^{k-1} \varepsilon \mathbf{U}_{i+1}^{\varepsilon}\right\| \geq \gamma\right) &\leq \mathbb{P}\left(\bigcup_{e=\pm e_i}\left\{\max_{1\leq k\leq n}\langle e, \sum_{i=0}^{k-1} \varepsilon \mathbf{U}_{i+1}^{\varepsilon}\rangle \geq \frac{\gamma}{m}\right\}\right) \\ &\leq 2m \exp\left(\frac{-\gamma^2}{2\Gamma \varepsilon T m^2}\right). \blacksquare\end{aligned}$$

5.3 Proof of Theorem 3.1

We introduce the notations $\|\mathbf{x}\|_{[0,T]} := \sup_{s\in[0,T]} \|\mathbf{x}(s)\|$, $d_{[0,T]}(\mathbf{x}, \mathbf{y}) := \left\|\mathbf{x} - \mathbf{y}\right\|_{[0,T]}$ and $d_{[0,T]}(\mathbf{x}, A) := \inf_{\mathbf{a}\in A} d_{[0,T]}(\mathbf{x}, \mathbf{a})$ for $\mathbf{x}, \mathbf{y} \in C(\mathbb{R}_+, X)$ and $A \subset C(\mathbb{R}_+, X)$. In addition we let $S_{[0,T]}$ denote the set of all solution curves of (DI) restricted to time interval $[0, T]$.

To prove Theorem 3.1, we first demonstrate a similar result for so-called perturbed solutions of the differential inclusion (DI). We then show that a family of GSAPs is a perturbed solution of (DI).

The following definition builds on one of Benaïm et al. (2005) for the decreasing step size case.

Definition 5.1. Let $\boldsymbol{\delta} = \{\boldsymbol{\delta}_T(\varepsilon)\}_{\varepsilon>0,T>0}$ be a family of real random variables and $\underline{\mathbf{U}} = \{\underline{\mathbf{U}}^{\varepsilon}(t)\}_{\varepsilon>0}$ a family of stochastic process $\underline{\mathbf{U}}^{\varepsilon} : \Omega \times \mathbb{R}_+ \to \mathbb{R}^n$. A family $\{\mathbf{y}_{\varepsilon}\}_{\varepsilon>0}$ of continuous processes $\mathbf{y}_{\varepsilon} : \Omega \times \mathbb{R}_+ \to \mathbb{R}^n$ is a $(\boldsymbol{\delta}, \underline{\mathbf{U}})$-perturbed solution of the differential inclusion (DI) if

(i) for almost all $\omega \in \Omega$ and all $\varepsilon > 0$, $\mathbf{y}_{\varepsilon}$ is absolutely continuous in t.

(ii) for all $T > 0$ and $\alpha > 0$, $\lim_{\varepsilon\to 0} \mathbb{P}(\boldsymbol{\delta}_T(\varepsilon) > \alpha) = 0$,

(iii) for almost all $\omega \in \Omega$, all $T > 0$ and almost all $0 < t < T$,

$$\frac{\mathrm{d}\mathbf{y}_{\varepsilon}(t)}{\mathrm{d}t} - \underline{\mathbf{U}}^{\varepsilon}(t) \in \tilde{V}^{\boldsymbol{\delta}_T(\varepsilon)}(\mathbf{y}_{\varepsilon}(t)),$$

(iv) for all $\varepsilon > 0$, $\underline{\mathbf{U}}^{\varepsilon}$ is almost surely locally integrable,

(v) for all $T > 0$ and $\alpha > 0$,

$$\lim_{\varepsilon\to 0} \mathbb{P}\left(\sup_{0\le v\le T} \left\| \int_0^v \underline{\mathbf{U}}^\varepsilon(s)\, \mathrm{d}s \right\| \ge \alpha \right) = 0. \tag{24}$$

Theorem 3.1 is a corollary of the following proposition, which generalizes the theorem to $(\boldsymbol{\delta}, \underline{\mathbf{U}})$-perturbed solutions.

Proposition 5.2. *Let $\{\mathbf{y}_\varepsilon\}_{\varepsilon>0}$ be a $(\boldsymbol{\delta}, \underline{\mathbf{U}})$-perturbed solution of the differential inclusion* (DI)*. Then for any $\alpha > 0$ and any $T > 0$, we have*

$$\lim_{\varepsilon\to 0} \mathbb{P}\left(d_{[0,T]}(\mathbf{y}_\varepsilon, S_\Phi) \ge \alpha\right) = 0,$$

Before starting the proof, we introduce a useful technical lemma due to Faure and Roth (2010). Let $T > 0$, define $\|V\| = \sup_{x\in X} \sup_{y\in V(x)} |y|$, and consider the compact set

$$K_V := \left\{ \mathbf{y} \in Lip([0,T], X) \colon Lip(\mathbf{y}) \le \|V\| + 1 \right\},$$

where $Lip([0,T], X)$ is the set of Lipschitz functions from $[0,T]$ to X and $Lip(\mathbf{y})$ is the Lipschitz constant of the function $\mathbf{y}$. Note that K_V contains every solution curve of (DI), restricted to $[0,T]$.

For $\gamma \in [0,1]$, we define the set-valued map $\Lambda^\gamma \colon K_V \rightrightarrows K_V$ by letting $\mathbf{y} \in \Lambda^\gamma(\mathbf{z})$ if and only if there exists an integrable $\mathbf{h}\colon [0,T] \to \mathbb{R}^n$ such that

$$\mathbf{h}(u) \in \tilde{V}^\gamma(\mathbf{z}(u)) \text{ for all } u \in [0,T], \text{ and}$$
$$\mathbf{y}(\tau) = \mathbf{z}(0) + \int_0^\tau \mathbf{h}(u)\, \mathrm{d}u \text{ for all } \tau \in [0,T].$$

We adopt the convention that $\Lambda^0 = \Lambda$. Observe that $Fix(\Lambda) := \{\mathbf{z} \in K_V \colon \mathbf{z} \in \Lambda(\mathbf{z})\} = S_{[0,T]}$. The following lemma of Faure and Roth (2010) says that if $\mathbf{z}$ is almost a fixed point of Λ^γ for some small enough γ, then it is almost a solution of $\dot{x} \in V(x)$.

Lemma 5.3. *Let $\alpha > 0$. There exist $\beta > 0$ and $\gamma_0 > 0$ such that for any $\gamma < \gamma_0$*

$$d_{[0,T]}(\mathbf{z}, \Lambda^\gamma(\mathbf{z})) < \beta \Rightarrow d_{[0,T]}(\mathbf{z}, S_{[0,T]}) < \alpha$$

for all $\mathbf{z} \in K_V$.

Proof of Proposition 5.2. Let $\{\mathbf{y}_\varepsilon\}$ be a $(\boldsymbol{\delta}, \underline{\mathbf{U}})$-perturbed solution of the differential inclusion (DI). Fix $T > 0$ and $\alpha > 0$. For $0 \le t \le T$, set

$$\mathbf{v}_\varepsilon(t) := \dot{\mathbf{y}}_\varepsilon(t) - \underline{\mathbf{U}}^\varepsilon(t) \in \tilde{V}^{\boldsymbol{\delta}_T(\varepsilon)}(\mathbf{y}_\varepsilon(t)) \text{ and}$$

$$\mathbf{z}_\varepsilon(t) := \mathbf{y}_\varepsilon(0) + \int_0^t \mathbf{v}_\varepsilon(s)\,\mathrm{d}s.$$

Then

$$\mathbf{y}_\varepsilon(t) - \mathbf{y}_\varepsilon(0) = \mathbf{z}_\varepsilon(t) - \mathbf{y}_\varepsilon(0) + \int_0^t \underline{\mathbf{U}}^\varepsilon(s)\,\mathrm{d}s. \tag{25}$$

Thus $\Delta(\varepsilon, T) := d_{[0,T]}\left(\mathbf{z}_\varepsilon, \mathbf{y}_\varepsilon\right) = \sup_{0\le s\le T}\left\|\int_0^s \underline{\mathbf{U}}^\varepsilon(u)\,\mathrm{d}u\right\|$, and in addition, the definitions of $\mathbf{v}_\varepsilon$ and of Λ^γ imply that

$$\mathbf{z}_\varepsilon \in \Lambda^{\Delta(\varepsilon,T)+\boldsymbol{\delta}_T(\varepsilon)}(\mathbf{z}_\varepsilon).$$

By Lemma 5.3, there exists a γ_0 such that for all $\varepsilon > 0$,

$$\Delta(\varepsilon, T) + \boldsymbol{\delta}_T(\varepsilon) < \gamma_0 \Rightarrow d_{[0,T]}(\mathbf{z}_\varepsilon, S_{[0,T]}) < \tfrac{\alpha}{2}.$$

Furthermore, equality (25) implies that

$$d_{[0,T]}(\mathbf{y}_\varepsilon, S_{[0,T]}) \le d_{[0,T]}(\mathbf{z}_\varepsilon, S_{[0,T]}) + \Delta(\varepsilon, T).$$

We therefore have

$$\begin{aligned}\mathbb{P}\left(d_{[0,T]}(\mathbf{y}_\varepsilon, S_{[0,T]}) > \alpha\right) &\le \mathbb{P}\left(d_{[0,T]}(\mathbf{z}_\varepsilon, S_{[0,T]}) + \Delta(\varepsilon, T) > \alpha\right) \\ &\le \mathbb{P}\left(d_{[0,T]}(\mathbf{z}_\varepsilon, S_{[0,T]}) > \tfrac{\alpha}{2}\right) + \mathbb{P}\left(\Delta(\varepsilon, T) > \tfrac{\alpha}{2}\right) \\ &\le \mathbb{P}\left(\Delta(\varepsilon, T) > \tfrac{\gamma_0}{2}\right) + \mathbb{P}\left(\boldsymbol{\delta}_T(\varepsilon) > \tfrac{\gamma_0}{2}\right) + \mathbb{P}\left(\Delta(\varepsilon, T) > \tfrac{\alpha}{2}\right).\end{aligned}$$

We conclude the proof by taking the limit in ε and using hypotheses (*ii*) and (*v*) on $\mathbf{y}_\varepsilon$. ■

The last step in the proof of Theorem 3.1 is to show that the family of interpolated processes $\{\bar{\mathbf{X}}^\varepsilon\}_{\varepsilon>0}$ induced by a family of GSAPs is a $(\boldsymbol{\delta}, \underline{\mathbf{U}})$-perturbed solution of the differential inclusion (DI), where

$$\boldsymbol{\delta}_T(\varepsilon) := \inf\{\delta\colon V_\varepsilon(x) \subset \tilde{V}^\delta(x)\ \forall x \in X\} + \varepsilon \max_{k\le[\frac{T}{\varepsilon}]}\left(\|\mathbf{U}^\varepsilon_{k+1}\| + \|V\| + 1\right),$$

and where $\underline{\mathbf{U}}^\varepsilon$ is the piecewise constant process induced by the sequence of random variables $\mathbf{U}^\varepsilon$ from Definition 2.2:

$$\underline{\mathbf{U}}^\varepsilon(t) = \mathbf{U}^\varepsilon_{\ell(t)+1} \text{ for } t \ge 0, \text{ where } \ell(t) = \left[\frac{t}{\varepsilon}\right].$$

To do this, we must verify the five conditions of Definition 5.1. Conditions (i), (iv) and (v) are direct consequences of the definitions of $\bar{\mathbf{X}}^\varepsilon$ and $\underline{\mathbf{U}}^\varepsilon$, so we need only consider conditions (ii) and (iii).

To show that $\boldsymbol{\delta}_T(\varepsilon)$ satisfies condition (ii), fix $\alpha > 0$. Then for ε small enough we have

$$\mathbb{P}\left(\boldsymbol{\delta}_T(\varepsilon) > \alpha\right) \leq \mathbb{P}\left(\varepsilon \max_{k\leq \ell(T)}\{\|\mathbf{U}_k^\varepsilon\|\} > \frac{\alpha}{2}\right) \leq \mathbb{P}\left(2\varepsilon \max_{k\leq \ell(T)} \left\|\sum_{i=0}^{k-1} \mathbf{U}_{i+1}^\varepsilon\right\| > \frac{\alpha}{2}\right).$$

Condition (ii) of Definition 5.1 then follows from condition (iv) of Definition 2.2.

To establish condition (ii) of Definition 5.1, fix $T > 0$. By the definition of $\bar{\mathbf{X}}^\varepsilon$ and by assumption (ii) of Definition 2.2, we have for any $0 \leq t \leq T$ that

$$\bar{\mathbf{X}}^\varepsilon(t) \in \mathbf{X}^\varepsilon_{\ell(t)} + (t - \varepsilon\ell(t))\left(\underline{\mathbf{U}}^\varepsilon(t) + V_\varepsilon(\mathbf{X}^\varepsilon_{\ell(t)})\right). \tag{26}$$

By assumption (iii) of Definition 2.2, the last inclusion implies that for $\varepsilon > 0$ small enough,

$$\left\|\bar{\mathbf{X}}^\varepsilon(t) - \mathbf{X}^\varepsilon_{\ell(t)}\right\| \leq \varepsilon \max_{k\leq [\frac{T}{\varepsilon}]}(\|\mathbf{U}^\varepsilon_{k+1}\| + \|V\| + 1),$$

and thus, using this assumption once more, that

$$V_\varepsilon(\mathbf{X}^\varepsilon_{\ell(t)}) \subset \tilde{V}^{\boldsymbol{\delta}_T(\varepsilon)}(\bar{\mathbf{X}}^\varepsilon(t)). \tag{27}$$

Condition (iii) of Definition 5.1 follows from expressions (26) and (27). This completes the proof of Theorem 3.1.

5.4 Proof of Theorem 3.2

Denote by $\{\mathbf{X}_k^\varepsilon\}_{k\geq 1}$ the Markov chain associated with the process $\mathbf{Y}^\varepsilon$ as defined in Remark 2.8. The continuous time affine interpolated process $\bar{\mathbf{Y}}^\varepsilon(\cdot)$ induced by $\mathbf{Y}^\varepsilon$ can be defined as

$$\bar{\mathbf{Y}}^\varepsilon(t) = \mathbf{X}^\varepsilon_{\kappa(t)} + (t - T_{\kappa(t)})\frac{\mathbf{X}^\varepsilon_{\kappa(t)+1} - \mathbf{X}^\varepsilon_{\kappa(t)}}{\tau_{\kappa(t)+1}}. \tag{28}$$

Fix $T > 0$. By assumption (ii) on the measure μ_x^ε, there is a constant $M > 0$ such that

$$\|\mathbf{X}^\varepsilon_{n+1} - \mathbf{X}^\varepsilon_n\| \leq M\varepsilon \quad \text{a.s.}, \tag{29}$$

which implies that

$$\sup_{0\le t\le T}\left\|\mathbf{Y}^\varepsilon(t)-\bar{\mathbf{Y}}^\varepsilon(t)\right\|\le M\varepsilon\quad\text{a.s..}\tag{30}$$

For each $n\in\mathbb{N}$, define $\mathbf{U}^\varepsilon_{n+1}$ by

$$\begin{aligned}\mathbf{U}^\varepsilon_{n+1}&=\frac{1}{\tau_{n+1}}\left(\mathbf{Y}^\varepsilon(T_{n+1})-\mathbf{Y}^\varepsilon(T_n)\right)-v^\varepsilon(\mathbf{Y}^\varepsilon(T_n))\\&=\frac{1}{\tau_{n+1}}\left(\mathbf{X}^\varepsilon_{n+1}-\mathbf{X}^\varepsilon_n\right)-v^\varepsilon(\mathbf{X}^\varepsilon_n),\end{aligned}$$

where v^ε is the vector field introduced in assumption (iii) of Definition 2.7. Let $\underline{\mathbf{U}}^\varepsilon(t)$ be the piecewise constant process induced by the sequence of random variables $\{\mathbf{U}^\varepsilon_n\}$:

$$\underline{\mathbf{U}}^\varepsilon(t)=\mathbf{U}^\varepsilon_{\kappa(t)+1}\text{ for }t\ge 0.$$

We claim that $\{\bar{\mathbf{Y}}^\varepsilon\}$ is a $(\boldsymbol{\delta},\underline{\mathbf{U}})$-perturbed solution of the differential inclusion (DI) for

$$\boldsymbol{\delta}_T(\varepsilon):=\inf\{\delta\colon v^\varepsilon(x)\in\tilde{V}^\delta(x)\ \forall x\in X\}+M\varepsilon.$$

If this claim is true, then the theorem follows from Proposition 5.2 and inequality (30).

To prove the claim, we check the five conditions of Definition 5.1. Conditions (i), (ii) and (iv) are direct consequences of the definition of $\{\bar{\mathbf{X}}^\varepsilon\}$. Moreover, by differentiating (28) and using the definition of $\mathbf{U}^\varepsilon_n$, we have almost surely for almost all $t\ge 0$ that

$$\begin{aligned}\tfrac{\mathrm{d}}{\mathrm{d}t}\bar{\mathbf{Y}}^\varepsilon(t)&=\frac{1}{\tau_{\kappa(t)+1}}\left(\mathbf{X}^\varepsilon_{\kappa(t)+1}-\mathbf{X}^\varepsilon_{\kappa(t)}\right)\\&=\underline{\mathbf{U}}^\varepsilon(t)+v^\varepsilon(\mathbf{X}^\varepsilon_{\kappa(t)})\\&\in\underline{\mathbf{U}}^\varepsilon(t)+\tilde{V}^\gamma(\mathbf{X}^\varepsilon_{\kappa(t)})\\&\subset\underline{\mathbf{U}}^\varepsilon(t)+\tilde{V}^{\boldsymbol{\delta}_T(\varepsilon)}(\bar{\mathbf{Y}}^\varepsilon(t)).\end{aligned}$$

where $\gamma=\inf\{\gamma\colon v^\varepsilon(x)\in\tilde{V}^\gamma(x)\text{ for all }x\in X\}$; the second step is a consequence of assumption (iii) of Definition 2.7, and the last inclusion is a consequence of (29) and the definition of $\boldsymbol{\delta}_T(\varepsilon)$. This establishes condition (iii). Finally, condition (v) is a consequence of the following lemma.

Lemma 5.4. *Under the hypotheses of Theorem 3.2, we have*

$$\lim_{\varepsilon\to 0}\mathbb{P}\left(\sup_{0\le t\le T}\left\|\int_0^t\underline{\mathbf{U}}^\varepsilon(s)\,\mathrm{d}s\right\|\ge\alpha\,\middle|\,\mathbf{X}^\varepsilon_0=x\right)=0,$$

uniformly in $x \in X$.

Proof of Lemma 5.4. Note first that for all $t \leq T$,

$$
\begin{aligned}
\left\|\int_0^t \underline{\mathbf{U}}^\varepsilon(s)ds\right\| &= \left\|\sum_{k=0}^{\kappa(t)} \tau_{k+1}\mathbf{U}^\varepsilon_{k+1} - (T_{\kappa(t)+1} - t)\mathbf{U}^\varepsilon_{\kappa(t)+1}\right\| \\
&= \left\|\mathbf{X}^\varepsilon_{\kappa(t)+1} - \mathbf{X}^\varepsilon_0 - \sum_{k=0}^{\kappa(t)} \tau_{k+1} v^\varepsilon(\mathbf{X}^\varepsilon_k) + (T_{\kappa(t)+1} - t)v^\varepsilon(\mathbf{X}_{\kappa(t)}) - \frac{T_{\kappa(t)+1} - t}{\tau_{\kappa(t)+1}}(\mathbf{X}^\varepsilon_{\kappa(t)+1} - \mathbf{X}^\varepsilon_{\kappa(t)})\right\| \\
&= \left\|\mathbf{Y}^\varepsilon(t) - \mathbf{X}^\varepsilon_0 - \int_0^t v^\varepsilon(\mathbf{Y}^\varepsilon(s))\,\mathrm{d}s + \frac{t - T_{\kappa(t)+1}}{\tau_{\kappa(t)+1}}(\mathbf{X}^\varepsilon_{\kappa(t)+1} - \mathbf{X}^\varepsilon_{\kappa(t)})\right\| \\
&\leq \left\|\mathbf{Y}^\varepsilon(t) - \mathbf{X}^\varepsilon_0 - \int_0^t v^\varepsilon(\mathbf{Y}^\varepsilon(s))\,\mathrm{d}s\right\| + \varepsilon M.
\end{aligned}
$$

It remains to estimate the probability that the right hand side of the last inequality is small. We proceed as in the proof of Proposition 4.6 in Benaïm (1999). Fix $\theta \in \mathbb{R}^n$, $x \in X$ and $T > 0$. Let $f : X \to \mathbb{R}$ be the map defined by $f(y) = \exp(\langle \theta, y - x\rangle)$. By standard results (see Lemma 4.3.2 of Ethier and Kurtz (1986)),

$$
f(\mathbf{Y}^\varepsilon(t)) \exp\left(-\int_0^t \frac{L^\varepsilon f(\mathbf{Y}^\varepsilon(s))}{f(\mathbf{Y}^\varepsilon(s))}\,\mathrm{d}s\right)
$$

is a martingale for the natural filtration. Moreover, if we set $g(u) = \mathrm{e}^u - u - 1$, then it follows from the definition of L^ε that

$$
\frac{L^\varepsilon f(y)}{f(y)} = \langle \theta, v^\varepsilon(y)\rangle + \frac{1}{\varepsilon}\int_{\mathbb{R}^n} g(\varepsilon\langle \theta, z\rangle)\,\mu^\varepsilon_y(dz).
$$

From this expression we deduce that there exists a constant $\Gamma > 0$ such that for all $\theta \in \mathbb{R}^n$ with $\|\theta\| \leq \frac{1}{\varepsilon}$,

$$
\frac{L^\varepsilon f(y)}{f(y)} - \langle \theta, v^\varepsilon(y)\rangle \leq \varepsilon\Gamma\,\|\theta\|^2 .
$$

It follows that the process

$$
Z_\theta(t) = \exp\left(\left\langle \theta, \mathbf{Y}^\varepsilon(t) - x - \int_0^t v^\varepsilon(\mathbf{Y}^\varepsilon(s))ds\right\rangle - t\varepsilon\Gamma\,\|\theta\|^2\right)
$$

is a supermartingale.

Now fix $\alpha > 0$ small enough. Using Doob's inequality as in the proof of Proposition

2.4, we obtain

$$\mathbb{P}\left(\sup_{0\le t\le T}\left\|\mathbf{Y}^{\varepsilon}(t)-x-\int_0^t v^{\varepsilon}(\mathbf{Y}^{\varepsilon}(s))ds\right\|\ge\alpha\;\middle|\;\mathbf{Y}^{\varepsilon}(0)=x\right)<C\exp\left(-\frac{\alpha^2}{\varepsilon 4\Gamma T}\right)$$

for some constant $C>0$. This concludes the proof of the lemma. ■

5.5 Proof of Proposition 3.4

Let the probability measure μ be given, and let the vanishing sequence $\{\varepsilon_n\}_{n=1}^{\infty}$ and the sequence of probability measures $\{\tilde{\nu}_n\}_{n=1}^{\infty}$ satisfy (IM1), (IM2) and (IM3). Our aim is to show that μ is semi-invariant for the set-valued dynamical system Φ. We begin by introducing a candidate measure ν on $C(\mathbb{R}_+,X)$ such that μ and ν satisfy the requirements of Definition 3.3.

As in hypothesis (IM1) we write ν_n for $(\mathcal{A}^{t_n})^*(\tilde{\nu}_n)$. We first establish

Lemma 5.5. $\{\nu_n\}_{n=1}^{\infty}$ *admits a tight subsequence.*

Proof. By hypothesis (IM1) there is an increasing sequence $\{N_k\}_{k\ge 0}$ of positive integers such that

$$\nu_n\left(\mathcal{N}^{1/k}(S_\Phi)\right)>1-\tfrac{1}{k}$$

for every $k\ge 1$ and $n\ge N_k$. We now show that the subsequence $\{\nu_{N_k}\}_{k\ge 1}$ is tight, i.e., for all $m\ge 1$, there exists a compact subset K_m of $C(\mathbb{R}_+,X)$ such that $\nu_{N_k}(K_m)>1-\frac{1}{m}$ for all $k\ge 1$. Let K be the subset of $C(\mathbb{R}_+,X)$ defined by

$$K:=S_\Phi\cup\left(\bigcup_{k\ge 1}\left(\mathcal{N}^{1/k}(S_\Phi)\cap\mathcal{A}^{\varepsilon_{N_k}}(X^{\mathbb{N}})\right)\right).$$

One can prove that K is a compact subset of $C(\mathbb{R}_+,X)$ by showing that every sequence in K has a convergent subsequence.

Now fix $m\ge 1$, and let K_m be the compact subset of $C(\mathbb{R}_+,X)$ defined by

$$K_m:=\bigcup_{k=1}^{m-1}\mathcal{A}^{\varepsilon_{N_k}}(X^{\mathbb{N}})\cup K. \tag{31}$$

For all $k\le m$, the definition of ν_{N_k} implies directly that $\nu_{N_k}(K_m)=1$. Moreover, for all $k>m$

we have that

$$
\begin{aligned}
(32) \qquad \nu_{N_k}(K_m) &\geq \nu_{N_k}\left(\mathcal{N}^{1/k}(S_\Phi) \cap \mathcal{A}^{\varepsilon_{N_k}}(X^{\mathbb{N}})\right) \\
&= \nu_{N_k}\left(\mathcal{N}^{1/k}(S_\Phi)\right) \\
&> 1 - \tfrac{1}{k} \geq 1 - \tfrac{1}{m}.
\end{aligned}
$$

This proves that $\{\nu_{N_k}\}_{k\geq 1}$ is tight. ■

Lemma 5.5 implies that the sequence $\{\nu_n\}$ admits a weak limit point ν. Without loss of generality we henceforth write $\nu = \lim_{n\to\infty} \nu_n$.

We now verify conditions (i), (ii), and (iii) of Definition 3.3. By construction, we have that $\pi_0 \circ \mathcal{A}^{\varepsilon_n} = \tilde{\pi}_0$ for all $n \geq 0$, which implies that $\pi_0^*(\nu_n) = \tilde{\pi}_0^*(\tilde{\nu}_n)$. Since π_0 is continuous, we conclude from hypothesis (IM3) that $\pi_0^*(\nu) = \mu$. This is condition (iii). Moreover, hypothesis (IM1) implies that

$$\nu\left(\mathcal{N}^{1/k}(S_\Phi)\right) = 1$$

for all $k \geq 1$. By taking the limit as k goes to infinity, we conclude that support$(\nu) \subset S_\Phi$. This is condition (i).

To complete the proof of the proposition, we show that ν is invariant under the shift operator Θ, which is condition (ii) of Definition 3.3. Let $f\colon S_\Phi \to \mathbb{R}$ be a continuous function and let $T \geq 0$; by equation (7), it is sufficient to prove that

$$(33) \qquad \int_{S_\Phi} f(\mathbf{z})\, \mathrm{d}\nu(\mathbf{z}) = \int_{S_\Phi} f(\Theta^T(\mathbf{z}))\, \mathrm{d}\nu(\mathbf{z}).$$

Since S_Φ is closed, the Tietze extension theorem implies that f can be extended to a continuous function defined on all of $C(\mathbb{R}_+, X)$ with $\left\|f\right\|_\infty = \sup_{\mathbf{z}\in C(\mathbb{R}_+,X)} \left|f(\mathbf{z})\right| < \infty$. We write

$$\xi_n = \int_{C(\mathbb{R}_+,X)} f(\mathbf{z})\, \mathrm{d}\nu_n(\mathbf{z}) \ \text{ and } \ \xi_n^T = \int_{C(\mathbb{R}_+,X)} f \circ \Theta^T(\mathbf{z})\, \mathrm{d}\nu_n(\mathbf{z}),$$

so that the two sides of equation (33) can be expressed as

$$\xi := \lim_{n\to\infty} \xi_n \ \text{ and } \ \xi^T := \lim_{n\to\infty} \xi_n^T.$$

The definition of $\mathcal{A}^t$ implies that $\Theta^t \circ \mathcal{A}^t = \mathcal{A}^t \circ \tilde{\Theta}$ for all $t > 0$. Thus if $\beta_n = T - [\frac{T}{\varepsilon_n}]\varepsilon_n$, then

$$\begin{aligned}\Theta^T \circ \mathcal{A}^{\varepsilon_n} &= \Theta^{\beta_n} \circ (\Theta^{\varepsilon_n})^{[\frac{T}{\varepsilon_n}]} \circ \mathcal{A}^{\varepsilon_n} \\ &= \Theta^{\beta_n} \circ \mathcal{A}^{\varepsilon_n} \circ \tilde{\Theta}^{[\frac{T}{\varepsilon_n}]}.\end{aligned}$$

This observation and the $\tilde{\Theta}$-invariance of ν_n (from hypothesis (IM2)) imply that

$$\begin{aligned}\xi_n^T &= \int_{C(\mathbb{R}_+,X)} f \circ \Theta^T(\mathbf{z})\, \mathrm{d}(\mathcal{A}^{\varepsilon_n})^*(\tilde{\nu}_n)(\mathbf{z}) \\ &= \int_{X^{\mathbb{N}}} f \circ \Theta^T \circ \mathcal{A}^{\varepsilon_n}(\mathbf{x})\, \mathrm{d}\tilde{\nu}_n(\mathbf{x}) \\ &= \int_{X^{\mathbb{N}}} f \circ \Theta^{\beta_n} \circ \mathcal{A}^{\varepsilon_n}(\mathbf{x})\, \mathrm{d}\tilde{\nu}_n(\mathbf{x}) \\ &= \int_{C(\mathbb{R}_+,X)} f \circ \Theta^{\beta_n}(\mathbf{z})\, \mathrm{d}(\mathcal{A}^{\varepsilon_n})^*(\tilde{\nu}_n)(\mathbf{z}) \\ &= \int_{C(\mathbb{R}_+,X)} f \circ \Theta^{\beta_n}(\mathbf{z})\, \mathrm{d}\nu_n(\mathbf{z})\end{aligned}$$

We now show that $\left|\xi_n - \xi_n^T\right|$ converges to zero. Fix $m > 0$, and let K_m be the compact subset of $C(\mathbb{R}_+, X)$ defined in (31). Since $\Theta\colon \mathbb{R}_+ \times C(\mathbb{R}_+, X) \to C(\mathbb{R}_+, X)$ is continuous, one can prove that the set

$$K^* := \bigcup_{n \geq 1} \Theta^{\beta_n}(K_m) \cup K_m$$

is a compact subset of $C(\mathbb{R}_+, X)$ by showing that every sequence in K has a convergent subsequence. Since $\Theta\colon [0,1] \times K_m \to C(\mathbb{R}_+, X)$ is uniformly continuous with $\Theta^0(\cdot)$ the identity function on K_m, β_n vanishes, and f is uniformly continuous on K^*, there exists an $N_0 \in \mathbb{N}$ large enough that

$$\left|f \circ \Theta^{\beta_n}(\mathbf{z}) - f(\mathbf{z})\right| < \tfrac{1}{m}$$

for all $n \geq N_0$ and all $\mathbf{z} \in K_m$. Using this fact and inequality (32), we see that

$$\begin{aligned}\left|\xi_n - \xi_n^T\right| &\leq \int_{K_m} \left|f \circ \Theta^{\beta_n}(\mathbf{z}) - f(\mathbf{z})\right| \mathrm{d}\nu_n(\mathbf{z}) + \int_{K_m^c} \left|f \circ \Theta^{\beta_n}(\mathbf{z}) - f(\mathbf{z})\right| \mathrm{d}\nu_n(\mathbf{z}) \\ &\leq \tfrac{1}{m} + 2\left\|f\right\|_\infty \tfrac{1}{m}\end{aligned}$$

Taking the limit as m grows large yields equation (33). This completes the proof of the theorem.

5.6 Proof of Theorem 3.5

For each $\varepsilon > 0$, let μ^ε be an invariant probability measure of the Markov chain $\mathbf{X}^\varepsilon$, and let μ be a weak limit point of $\{\mu^\varepsilon\}_{\varepsilon>0}$. Then there is a vanishing sequence $\{\alpha_k\}_{k=1}^\infty$ of positive real numbers such that μ^{α_k} converges to $\mu \in \mathcal{P}(X)$. In addition, let ν_k be the measure on the path space $X^\mathbb{N}$ induced by the Markov chain $\{\mathbf{X}^{\alpha_k}\}$ when it is run from initial distribution μ^{α_k}: that is, $\nu_k(A) := \mathbb{P}_{\mu^{\alpha_k}}(\mathbf{X}^{\alpha_k} \in A)$ for each Borel set $A \subset X^\mathbb{N}$.

To show that μ is semi-invariant for the set-valued dynamical system Φ, we apply Proposition 3.4 to the sequences $\{\alpha_k\}_{k=1}^\infty$ and $\{\nu_k\}_{k=1}^\infty$. Because μ^{α_k} is an invariant probability measure for $\{\mathbf{X}^{\alpha_k}\}$, the measure ν_k is $\tilde{\Theta}$-invariant, establishing condition (IM2). Moreover, $\tilde{\pi}_0^*(\nu_k) = \mu^{\alpha_k}$ by construction, and so $\{\tilde{\pi}_0^*(\nu_k)\}_{k=1}^\infty$ converges to μ, establishing condition (IM3).

It remains to verify condition (IM1). Fix $\delta > 0$, and then choose $T > 0$ and $\gamma > 0$ such that $\sum_{i=[T]}^\infty \frac{1}{2^{i+1}} < \frac{1}{2}\delta$ and $\gamma \sum_{i=0}^{[T]} \frac{1}{2^{i+1}} < \frac{1}{2}\delta$. These choices imply that

$$\left\{\mathbf{z} \in C(\mathbb{R}_+, X) \colon \inf_{\mathbf{y}\in S_\Phi} \sup_{0\le s\le T} \left\|\mathbf{z}(s) - \mathbf{y}(s)\right\| \le \gamma\right\} \subset \mathcal{N}^\delta(S_\Phi).$$

It follows that

$$\begin{aligned}
\tilde{\nu}_k(\mathcal{N}^\delta(S_\Phi)) &= \nu_k\big((\mathcal{A}^{\alpha_k})^{-1}(\mathcal{N}^\delta(S_\Phi))\big) \\
&= \mathbb{P}_{\mu^{\alpha_k}}\big(\mathbf{X}^{\alpha_k} \in (\mathcal{A}^{\alpha_k})^{-1}(\mathcal{N}^\delta(S_\Phi))\big) \\
&= \int_X \mathbb{P}_x\big(\mathbf{X}^{\alpha_k} \in (\mathcal{A}^{\alpha_k})^{-1}(\mathcal{N}^\delta(S_\Phi))\big)\, \mathrm{d}\mu^{\alpha_k}(x) \\
&\ge \int_X \mathbb{P}_x\left(\inf_{\mathbf{z}\in S_\Phi} \sup_{0\le s\le T} \|\bar{\mathbf{X}}^{\alpha_k}(s) - \mathbf{z}(s)\| \le \gamma\right) \mathrm{d}\mu^{\alpha_k}(x).
\end{aligned}$$

This inequality and Theorem 3.1 imply condition (IM1), completing the proof of the theorem.

5.7 Proof of Theorem 3.8

In the classical framework of a semi-flow on a separable metric space, the Poincaré recurrence theorem is stated as follows.

Theorem 5.6 (Poincaré)**.** *Let $(\mathcal{X}, d)$ be a separable metric space and $\theta = \{\theta_t\}_{t\in\mathbb{R}}$ a semi-flow on $\mathcal{X}$. Define the ω-limit set of $x \in \mathcal{X}$ by $\omega_\theta(x) = \bigcap_{t>0} \mathrm{cl}(\theta_{[t,\infty)}(x))$, the set of recurrent points by $R_\theta^\omega = \{x \in \mathcal{X} \colon x \in \omega_\theta(x)\}$, and the Birkhoff center of θ and denoted $BC(\theta) = \mathrm{cl}(R_\theta^\omega)$. If μ is an invariant measure for θ, then $\mu(BC(\theta)) = 1$.*

See Mañé (1987) for further discussion.

We now show that Theorem 3.8 is a corollary of Poincaré's recurrence theorem. Let μ be a semi-invariant measure for Φ and ν be an invariant measure for $\Theta_{|_{S_\Phi}}$ (henceforth written as Θ) such that $\pi_0^*(\nu) = \mu$. To begin, notice that

$$\pi_0(BC(\Theta)) \subset BC(\Phi).$$

Indeed, let $\mathbf{z}$ be a recurrent point of Θ. Then there exists a sequence of positive real numbers t_n approaching infinity such that $\lim_{n\to\infty} \Theta_{t_n}(\mathbf{z}) = \mathbf{z}$. In particular, $\pi_0(\mathbf{z}) = \mathbf{z}(0) = \lim_{n\to\infty} \mathbf{z}(t_n)$, which means that $\mathbf{z}(0) \in L(\mathbf{z}(0))$. Since π_0 is continuous, we obtain the inclusion. The inclusion implies that

$$\begin{aligned}\mu(BC(\Phi)) &\geq \mu(\pi_0(BC(\Theta)))\\ &= \nu(\pi_0^{-1} \circ \pi_0(BC(\Theta)))\\ &\geq \nu(BC(\Theta)),\end{aligned}$$

and the last quantity is equal to 1 by Theorem 5.6 applied to semi-flow Θ. ■

5.8 Proof of Lemma 4.3

Observe that

$$\mathbb{E}\left(\left\|\mathbf{U}_{k+1}^N\right\|^2\right) = N^2 \sum_{i\in S} \mathbb{E}\left(\left\|\sum_{r=0}^{R_i^{N,X_k}} \xi_{i,r}^{N,X_k} - \mathbb{E}\left(\sum_{r=0}^{R_i^{N,\mathbf{X}_k^N}} \xi_{i,r}^{N,\mathbf{X}_k^N}\right)\right\|^2\right). \tag{34}$$

We therefore evaluate the conditional expectation

$$\mathbb{E}\left(\left\|\sum_{r=0}^{R_i^{N,\mathbf{X}_k^N}} \xi_{i,r}^{N,\mathbf{X}_k^N} - \mathbb{E}\left(\sum_{r=0}^{R_i^{N,\mathbf{X}_k^N}} \xi_{i,r}^{N,\mathbf{X}_k^N}\right)\right\|^2 \middle| \mathbf{X}_k^N = x\right) = \mathbb{E}\left(\left\|\sum_{r=0}^{R_i^{N,x}} \xi_{i,r}^{N,x} - \mathbb{E}\left(R_i^{N,x}\right)\mathbb{E}\left(\xi_{i,r}^{N,x}\right)\right\|^2\right). \tag{35}$$

To do so, we compute the conditional expectation of the right hand side of (34) given that $R_i^{N,x} = R$, taking advantage of the fact that $\left\|\xi_{i,r}^{N,x}\right\| \leq \frac{\sqrt{2}}{N}$:

$$\mathbb{E}\left(\left\|\sum_{r=0}^{R_i^{N,x}} \xi_{i,r}^{N,x} - \mathbb{E}\left(R_i^{N,x}\right)\mathbb{E}\left(\xi_{i,r}^{N,x}\right)\right\|^2 \middle| R_i^{N,x} = R\right) = \mathbb{E}\left(\left\|\sum_{r=0}^{R} \xi_{i,r}^{N,x} - \mathbb{E}\left(R_i^{N,x}\right)\mathbb{E}\left(\xi_{i,r}^{N,x}\right)\right\|^2\right)$$

$$\leq \mathbb{E}\left(\left(\sum_{r=0}^{R}\left\|\xi_{i,r}^{N,x}\right\| + \mathbb{E}\left(R_i^{N,x}\right)\left\|\mathbb{E}\left(\xi_{i,r}^{N,x}\right)\right\|\right)^2\right)$$
$$\leq \frac{2}{N^2}\left((R + \mathbb{E}R_i^{N,x})^2\right).$$

Since $\mathbb{E}R_i^{N,x} = x_i \leq 1$ and $\mathbb{E}\left(R_i^{N,x}\right)^2 = \left(\mathbb{E}R_i^{N,x}\right)^2 + \operatorname{Var}\left(R_i^{N,x}\right) = x_i^2 + \frac{N-1}{N}x_i \leq 2$, we find that

$$\text{(36)} \qquad \mathbb{E}\left(\left\|\sum_{r=0}^{R_i^{N,\mathbf{x}_k^N}} \xi_{i,r}^{N,\mathbf{X}_k^N} - \mathbb{E}\left(\sum_{r=0}^{R_i^{N,\mathbf{x}_k^N}} \xi_{i,r}^{N,\mathbf{X}_k^N}\right)\right\|^2\right) \leq \frac{2}{N^2}\sum_{R=0}^{Nx_i}\mathbb{P}\left(R_i^{N,x} = R\right)\left((R + \mathbb{E}R_i^{N,x})^2\right)$$
$$\leq \frac{2}{N^2}\left(\mathbb{E}\left(R_i^{N,x}\right)^2 + 2\left(\mathbb{E}R_i^{N,x}\right)^2 + \left(\mathbb{E}R_i^{N,x}\right)^2\right)$$
$$= \frac{10}{N^2}.$$

We therefore conclude from (34), (35), and (36) that $\mathbb{E}\left(\left\|\mathbf{U}_{k+1}^N\right\|^2\right) \leq 10\,|S|$. ■

References

Aubin, J.-P. and Cellina, A. (1984). *Differential Inclusions*. Springer, Berlin.

Balkenborg, D., Hofbauer, J., and Kuzmics, C. (2011). Refined best-response correspondence and dynamics. *Theoretical Economics*, forthcoming.

Barbashin, E. A. (1948). On the theory of generalized dynamical systems. *Moskov. Gos. Ped. Inst. Učen. Zap.*, 2:110–133.

Benaïm, M. (1998). Recursive algorithms, urn processes, and the chaining number of chain recurrent sets. *Ergodic Theory and Dynamical Systems*, 18:53–87.

Benaïm, M. (1999). Dynamics of stochastic approximation algorithms. In Azéma, J., Émery, M., Ledoux, M., and Yor, M., editors, *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pages 1–68. Springer, Berlin.

Benaïm, M. and Hirsch, M. W. (1996). Asymptotic pseudotrajectories and chain recurrent flows, with applications. *Journal of Dynamics and Differential Equations*, 8:141–176.

Benaïm, M. and Hirsch, M. W. (1999). Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior*, 29:36–72.

Benaïm, M., Hofbauer, J., and Sorin, S. (2005). Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44:328–348.

Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.

Benveniste, A., Métivier, M., and Priouret, P. (1990). *Adaptive Algorithms and Stochastic Approximations*. Springer, Berlin.

Brown, G. W. (1949). Some notes on computation of games solutions. Report P-78, The Rand Corporation.

Brown, G. W. (1951). Iterative solutions of games by fictitious play. In Koopmans, T. C. et al., editors, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, New York.

Duflo, M. (1996). *Algorithmes Stochastiques*. Springer, New York.

Ethier, S. N. and Kurtz, T. G. (1986). *Markov Processes: Characterization and Convergence*. Wiley, New York.

Faure, M. and Roth, G. (2010). Stochastic approximations of set-valued dynamical systems: Convergence with positive probability to an attractor. *Mathematics of Operations Research*, 35:624–640.

Faure, M. and Roth, G. (2011). Ergodic properties of weak asymptotic pseudotrajectories for set-valued dynamical systems. *arXiv:1101.2154v3*.

Fudenberg, D. and Kreps, D. M. (1993). Learning mixed equilibria. *Games and Economic Behavior*, 5:320–367.

Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge.

Gast, N. and Gaujal, B. (2010). Mean-field limit of non-smooth systems and differential inclusions. *ACM SIGMETRICS Performance Evaluation Review*, 38:30–32.

Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.

Gorodeisky, Z. (2008). Stochastic approximation of discontinuous dynamics. Unpublished manuscript, Hebrew University.

Gorodeisky, Z. (2009). Deterministic approximation of best-response dynamics for the Matching Pennies game. *Games and Economic Behavior*, 66:191–201.

Hofbauer, J. (1995). Stability for the best response dynamics. Unpublished manuscript, University of Vienna.

Hofbauer, J. (2000). From Nash and Brown to Maynard Smith: Equilibria, dynamics, and ESS. *Selection*, 1:81–88.

Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294.

Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69.

Hofbauer, J. and Sandholm, W. H. (2009). Stable games and their dynamics. *Journal of Economic Theory*, 144:1665–1693.

Hofbauer, J. and Sorin, S. (2006). Best response dynamics for continuous zero-sum games. *Discrete and Continuous Dynamical Systems B*, 6:215–224.

Kurtz, T. G. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58.

Kushner, H. J. and Yin, G. G. (1997). *Stochastic Approximation Algorithms and Applications*. Springer, New York.

Lahkar, R. and Sandholm, W. H. (2008). The projection dynamic and the geometry of population games. *Games and Economic Behavior*, 64:565–590.

Ljung, L. (1977). Analysis of recursive stochastic algorithms. *IEEE Transactions on Automatic Control*, 22:551–575.

Mañé, R. (1987). *Ergodic Theory and Differentiable Dynamics*. Springer-Verlag, Berlin.

Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.

Métivier, M. and Priouret, P. (1987). Théorèmes de convergence presque sure pour une classe d'algorithmes stochastiques à pas décroissant. *Probab. Theory Related Fields*, 74(3):403–428.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

Morris, S., Rob, R., and Shin, H. S. (1995). p-dominance and belief potential. *Econometrica*, 63:145–157.

Nagurney, A. and Zhang, D. (1997). Projected dynamical systems in the formulation, stability analysis, and computation of fixed demand traffic network equilibria. *Transportation Science*, 31:147–158.

Oyama, D., Sandholm, W. H., and Tercieux, O. (2012). Sampling best response dynamics and deterministic equilibrium selection. Unpublished manuscript, University of Tokyo, University of Wisconsin, and Paris School of Economics.

Sandholm, W. H. (2001a). Almost global convergence to p-dominant equilibrium. *International Journal of Game Theory*, 30:107–116.

Sandholm, W. H. (2001b). Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108.

Sandholm, W. H. (2003). Evolution and equilibrium under inexact information. *Games and Economic Behavior*, 44:343–378.

Sandholm, W. H. (2009). Large population potential games. *Journal of Economic Theory*, 144:1710–1725.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.

Sandholm, W. H. (2012). Population games and deterministic evolutionary dynamics. Unpublished manuscript, University of Wisconsin.

Smirnov, G. V. (2002). *Introduction to the Theory of Differential Inclusions*. American Mathematical Society, Providence, RI.

Stroock, D. W. (1993). *Probability Theory: An Analytic View*. Cambridge University Press, Cambridge.

Taylor, P. D. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.

Tercieux, O. (2006). p-best response set. *Journal of Economic Theory*, 131:45–70.

Walters, P. (1982). *An Introduction to Ergodic Theory*. Springer-Verlag, New York.

Zeeman, E. C. (1980). Population dynamics from game theory. In Nitecki, Z. and Robinson, C., editors, *Global Theory of Dynamical Systems (Evanston, 1979)*, number 819 in Lecture Notes in Mathematics, pages 472–497, Berlin. Springer.

Zusai, D. (2011). The tempered best response dynamic. Unpublished manuscript, Temple University.