THE ROBUSTNESS OF GAME DYNAMICS UNDER RANDOM PERTURBATIONS

PIERRE-LOUIS CAUVIN

Master Thesis

SUPERVISOR

PANAYOTIS MERTIKOPOULOS — Laboratoire d'Informatique de Grenoble (LIG)

REVIEWERS

SANA LOUHICHI — Laboratoire Jean Kuntzmann (LJK)
PANAYOTIS MERTIKOPOULOS — Laboratoire d'Informatique de Grenoble (LIG)

Master of Sciences in Applied Mathematics (MSIAM)
Internship done at Laboratoire Jean Kuntzmann, Saint-Martin-d'Hères

ABSTRACT

In this work, we examine the robustness of game dynamics in the presence of random shocks and disturbances to the underlying system. Expanding on the microfoundations of evolution by imitation, we propose a class of stochastic imitation dynamics, and we examine the model's long-run behavior as a function of the imitation protocol and the magnitude of the stochastic perturbations affecting the system. In particular, we derive a set of sufficient conditions that guarantee the extinction of dominated strategies and the asymptotic stability of Nash equilibria (with probability 1 and high probability respectively), as well as rates at which convergence occur. In a subsequent part, we also explain how our results can be extended to cover a stochastic variant of the exponential dynamics, which contain several important variants of the stochastic replicator dynamics, now unified under the same general framework.

DISCLAIMER. This research project was carried on during an internship at Laboratoire Jean Kuntzmann supervised by Panayotis Mertikopoulos on the initial topic *The rate of convergence of regularized learning in the presence of noise*. However, we will almost never touch on the field of regularized learning in this manuscript, as the author's research shifted toward a more game theoretic analysis during his internship.

ACKNOWLEDGMENTS

First and foremost, I would like to express my sincere gratitude to my supervisor Panayotis. Always present and cheerful, he has been an infinite source of good ideas and interesting research directions during my whole internship. He offered me the great opportunity to work on such amazing mathematical topics, giving me full trust and freedom in my work even if my scholar formation was not one of a theoretician. Without him this work would have never emerged the same, be it the content or the writing style for which his masterfully written papers inspire me every line, and I am more than eager to continue working with him for my upcoming PhD.

Then, I would like to address my thankful feelings to the whole DAO team at Laboratoire Jean Kuntzmann, an amazing research team with even more amazing people. I had never felt that welcomed in a place as this one, being able to talk about mathematical stuff freely and being *most of the time* understood is a precious feeling that I am happy to finally have.

I also need to devote a special thanks to my parents, who always gave me the freedom to pursue my professional dreams, even if they did not understand any part of what I was working on.

My final, and not least, gratitude goes to the most important person in my life, my girlfriend Ambre, who has always been here for me whatever the difficulties. Without her I would have never felt that happy and fulfilled in life, and I dedicate all of my work for her.

CONTENTS

1

INTRODUCTION

GAME theory is the mathematical study of strategic interactions between agents participating in a competitive environment. One of the main question in game theory is whether it leads to rational behavior of the agents, i.e., if actions that are uniformly worse than another are less and less used along the game, and if the players' choices converge to some equilibrium states where no individual has an incentive to change their strategy (Nash, 1950).

Since the early works of Von Neumann on the equilibrium of zero-sum games (von Neumann, 1928; von Neumann and Morgenstern, 1947), many research has been done to study these rationality questions in different kinds of game and dynamics, such as non-zero, symmetric, asymmetric and congestion games. Furthermore, game theory was found to have various important applications in biology (Smith, 1973; Sandholm, 2010), economy (Bajari et al., 2013), machine learning (Bloembergen et al., 2015; Hazra and Anjaria, 2022) and other social sciences.

Most studies rely on the underlying assumption that interactions are influenced solely by each individual's actions and that the payoff for their strategy is fully known. However, the scarcity of accurate payoff feedback and the effect of an underlying environment are inevitable in numerous modern applications of game theory.

This dichotomy leads to an increasing need for robust game procedures whose rational properties remain true even in the presence of incertitude. In other words, we are seeking answers to the following important question :

Do game dynamics still lead to rational behaviors even under random perturbations ?

This manuscript is mainly interested about the study of robustness for continuous-time dynamics generated by a general class of games called evolutionary games. As their underlying systems are governed by ordinary differential equations, we inject random noise in the dynamics by extending the mathematical framework to work instead with stochastic differential equations.

1.1 RELATED WORKS

An important evolutionary game dynamics is the so-called replicator dynamics, introduced by Taylor and Jonker (1978) to model reproduction of species in an evolutionary setting. It has been shown that this dynamics also arises naturally from models of imitations (cf. Björnerstedt and Weibull, 1996; Schlag, 1998; Weibull, 1995) and continuous-time exponential learning (cf. Hofbauer et al., 2009; Mertikopoulos and Moustakas, 2010).

In particular, dominated strategies become extinct under the replicator dynamics (Akin, 1980; Nachbar, 1990), and the "folk theorem" of evolutionary games (Hofbauer and Sigmund, 2003) gives stability of Nash equilibria.

Random perturbations for the replicator dynamics were first introduced formally by Young and Foster (1991), and were extended rigorously by Fudenberg and Harris (1992) to cover aggregate shocks of nature on the reproducibility of species. It has been then refined and studied by Cabrales (2000), Imhof (2005) and Hofbauer and Imhof (2009), who have shown that the rational behavior of replicator dynamics still hold under random perturbations whenever the noise is small enough. Khasminskii and Potsepun (2006) have also studied the case where the noise is considered of Stratonovich-kind, and Vlasic (2018) for the robustness under discontinuous shocks due to abrupt natural disasters.

Taking a different approach from the point of view of continuous-time exponential learning, Mertikopoulos and Moustakas (2010) have introduced a different stochastic version to the replicator dynamics, which has the surprisingly property that the rationality holds irrespective of the noise level. This result has also been generalized to cover stochastic perturbations in arbitrary regularized learning on the simplex by Bravo and Mertikopoulos (2017).

A last variant to add noise in replicator dynamics was carried on by Mertikopoulos and Viossat (2016) from the starting point of micro-foundations and revision protocols. Similarly to the study in biological setting, they have shown that rational properties stay true only under a mild condition on the noise.

1.2 OUTLINE

Stochastic analysis on semimartingales

CHAPTER 2. In this first chapter, we provide a short introduction to the topic of stochastic analysis on continuous semimartingales, that will be used throughout the text. The content presented in Sections 2.1 and 2.2 is standard and can be found in any theoretical textbook, such as Le Gall (2016) or Revuz and Yor (1999). Less known results such as the existence of correlated Brownian motions and growth rates of stochastic integrals are also given in Section 2.3.

Imitation dynamics

CHAPTER 3. Starting with a gentle exposition of population games and their associated game theoretic notions such as dominated strategies and Nash equilbria (Section 3.1), we then introduce in Section 3.2 the classical framework of deterministic imitation dynamics. After that, we propose a way to consider random perturbations in these dynamics, leading to the definition of so-called stochastic imitation dynamics (Section 3.3). In particular, we show that they have a unique solution staying in the simplex, and that they contain previously studied stochastic dynamics such as the replicator dynamics with payoff shocks of Mertikopoulos and Viossat (2016).

Extinction of strategies

CHAPTER 4. In Chapter 4, we begin our exploration of the long-run behavior of stochastic imitation dynamics by studying the extinction of dominated strategies. We first provide a general theorem which states that the noise needs to be "kind enough" in order to observe elimination of mixed strategies (Section 4.1). To answer more qualitative considerations, we also provide rates of extinction in Section 4.2. In particular, we prove an asymptotic rate, a concentration inequality and an hitting time for the extinction of mixed strategies.

Stability of equilibria

CHAPTER 5. Continuing the study of long-run behaviors of stochastic imitation dynamics, Chapter 5 is focused on the stochastic stability of equilibria. In Section 5.1, we show in particular that limits of interior solution orbits and stochastically stable states verify a perturbed equilibrium condition. On the other hand, we also prove that states verifying strictly this equilibrium condition are stochastically asymptotically stable under stochastic imitation dynamics. Similarly to Chapter 4, Section 5.2 provides rates of convergence to asymptotically stable states. To streamline the presentation of this chapter, the proofs of main results are delayed to Sections 5.3–5.5.

CHAPTER 6. In this chapter, we generalize our dynamical framework and introduce accordingly the class of stochastic exponential dynamics (Section 6.1). This new class includes many stochastic variants of replicator dynamics (see Section 6.2), and so constitutes a great framework to unify their study. We explore the long-run behavior of such dynamics in Section 6.3 and obtain similar results as those of Chapters 4 and 5, i.e., conditions on the noise to observe extinction of mixed strategies and characterization of stability according to some perturbed equilibria.

Stochastic exponential dynamics

CHAPTER 7. In order to unify the analysis of extinction of strategies and stability of equilibria, we introduce the notion of variationally aligned sets and present important special cases of such sets in Section 7.1. We then study their stability properties with respect to both deterministic (Section 7.2) and stochastic (Section 7.3) exponential dynamics. In particular, we obtain general stability results for closed under better-replies faces when the vector field verifies some monotone condition.

Variationally aligned sets

CHAPTER 8. This last chapter presents research perspectives that the author considers in his future works. Those include a more thorough study of variationally aligned sets and their link with evolutionarily stable sets, the robustness of game dynamics under non continuous perturbations, and the influence of random noise in mirror descent dynamics on general spaces.

Perspectives

CONTRIBUTIONS. We also mention that starting from Section 3.3 of Chapter 3, all results presented are new contributions mainly due to the author (and fruitful discussions with his supervisor).

1.3 NOTATION AND TERMINOLOGY

NOTATIONAL CONVENTIONS. Let $\mathcal{S} = \{s_\alpha\}_{0\le\alpha\le n}$ be a finite space. We denote by $\mathbb{R}^{\mathcal{S}}$ the vector space generated by $\mathcal{S}$, and write $\{e_s\}_{s\in\mathcal{S}}$ its canonical basis. Whenever the context is clear, we will indistinguishably use the notation α to refer either to the element s_α of $\mathcal{S}$ or to the unit basis vector e_{s_α}. Accordingly, the set of probability measure on $\mathcal{S}$ will be identified with the n-dimensional simplex of $\mathbb{R}^{n+1}$ defined by $\Delta(\mathcal{S}) = \left\{x \in \mathbb{R}^{n+1} : \sum_{\alpha=0}^{n} x_\alpha = 1,\ x_\alpha \ge 0\ \forall\alpha\right\}$. For notational convenience, if $\mathcal{S}' = \{\mathcal{S}_k\}_{k\in\mathcal{N}}$ is a finite family of finite sets, we will write the point $(\dots,\alpha_{k-1},\alpha_k,\alpha_{k+1},\dots) \in \mathcal{S}'$ as $(\alpha_k;\alpha_{-k})$, and use $\sum_\alpha^k$ in the place of $\sum_{\alpha\in\mathcal{S}_k}$. When f is a function from $\mathcal{S}^2$ to $\mathbb{R}$, we will sometimes use the shorthand $f(\cdot,\beta)$ to denote the vector $[f(\alpha,\beta)]_{\alpha\in\mathcal{S}}$, and similarly use $g(\cdot,\cdot,\alpha',\beta')$ to denote the matrix $[g(\alpha,\beta,\alpha',\beta')]_{\alpha,\beta\in\mathcal{S}}$ when g is a function from $\mathcal{S}^4$ to $\mathbb{R}$.

CONVENTIONS OF STYLE. Throughout this manuscript, we consider genderless agents and individuals. When an individual is to be singled out, we will consistently employ the pronoun "they" and its inflected or derivative forms.

2

STOCHASTIC ANALYSIS ON SEMIMARTINGALES

This chapter constitutes a quick overview to stochastic analysis on continuous semimartingales and the important results that we will be using throughout the manuscript. We assume that the reader is already familiar with continuous stochastic processes[1] and has already been exposed to the theory of stochastic analysis on Brownian motions in an applied setting (e.g., from finance-oriented textbooks such as Øksendal, 2007).

Sections 2.1 and 2.2 are standard in the litterature of stochastic analysis, see e.g., Le Gall (2016) or Revuz and Yor (1999), so they can be skipped at first hand by an experimented reader. On the other hand, Section 2.3 contains less known results that will be used throughout the manuscript.

2.1 CONTINUOUS SEMIMARTINGALES

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A *filtration* on $(\Omega, \mathcal{F}, \mathbb{P})$ is a collection $(\mathcal{F}_t)_{t\geq 0}$ of sub σ-algebras of $\mathcal{F}$ such that $\mathcal{F}_s \subseteq \mathcal{F}_t$ for all $s \leq t$. We say that the filtration is *complete* if for every $A \in \mathcal{F}$ such that $\mathbb{P}(A) = 0$, all subset $A' \subseteq A$ are in $\mathcal{F}_t$ (i.e., the filtration contain all *null sets*). Furthermore, $(\mathcal{F}_t)$ is called *right-continuous* if $\bigwedge_{\varepsilon>0} \mathcal{F}_{t+\varepsilon} = \mathcal{F}_t$.

In what follows, we always place ourselves in a *filtered* probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ where the filtration $(\mathcal{F}_t)_{t\geq 0}$ is assumed to be complete and right-continuous[2]. *filtered probability space*

A *stochastic process* $X\colon [0,\infty) \times \Omega \longrightarrow \mathbb{R}$ is $(\mathcal{F}_t)$-*adapted* if $\omega \mapsto X(t,\omega)$ is $(\mathcal{F}_t)$-measurable for all $t \geq 0$. A stronger assumption is to require X to be $(\mathcal{F}_t)$-*progressive* , meaning that X is $(\mathcal{B}([0,t]) \otimes \mathcal{F}_t)$-measurable for all $t \geq 0$; where $\mathcal{B}([0,t])$ is the Borel σ-algebra of $[0,t]$. In particular, if X is $(\mathcal{F}_t)$-adapted and has *left-continuous* (resp. *right-continuous*) paths, then it is also $(\mathcal{F}_t)$-progressive. *$(\mathcal{F}_t)$-adapted process* *$(\mathcal{F}_t)$-progressive process*

We recall that X is said to be a $(\mathcal{F}_t)$-*martingale* if *a*) X is $(\mathcal{F}_t)$-adapted; *b*) $X(t)$ is integrable for all $t \geq 0$; and *c*) $\mathbb{E}[X(t) \,|\, \mathcal{F}_s] = X(s)$ whenever $s \leq t$. *martingale*

Definition 2.1.1 (Local martingale)**.** Let X be a $(\mathcal{F}_t)$-adapted stochastic process with values in $\mathbb{R}$. Then X is called a $(\mathcal{F}_t)$-*local martingale* if there exists an increasing sequence τ_n of $(\mathcal{F}_t)$-stopping times[3] such that : *local martingale*

1. $\tau_n \underset{n\to\infty}{\longrightarrow} +\infty$ (a.s.);

2. $t \longrightarrow X(t \wedge \tau_n)$ is a $(\mathcal{F}_t)$-martingale for every integer n.

An $(\mathcal{F}_t)$-adapted stochastic process A is said to be of *finite variation* if it verifies *process of finite variation*

1 Up to Brownian motions and martingales in a measure theoretic framework.

2 Considering complete right-continuous filtrations simplifies the manipulation of stopping times and yields regular modifications of martingales. The interested reader is refereed to Section 5.7 of Çınlar (2011) for an extensive treatment of these assumptions.

3 A $(\mathcal{F}_t)$-stopping time is a random variable $\tau\colon \Omega \longrightarrow [0,\infty)$ such that $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$.

$\sup \sum_k |A(t_k) - A(t_{k-1})| < \infty$ (a.s.) where the sup ranges over all partitions (t_k) of $[0,t]$. In particular, if Y is a $(\mathcal{F}_t)$-progressive process such that $\int_0^t |Y(s,\omega)|\, ds < \infty$ (a.s.) for every $t \geq 0$, then the process $\int_0^t Y(s,\omega)\, ds$ is a finite variation process (see Proposition 4.5 and the following remarks of Le Gall, 2016).

With this knowledge in hand, we can now define what is a *semimartingale*, the main mathematical object of this chapter.

semimartingale

Definition 2.1.2 (Semimartingale)**.** Let X be an $(\mathcal{F}_t)$-adapted stochastic process with values in $\mathbb{R}$. Then X is called a $(\mathcal{F}_t)$-*semimartingale* if it can be decomposed as

$$X(t) = X(0) + M(t) + A(t) \tag{2.1.1}$$

where M is a $(\mathcal{F}_t)$-local martingale and A is a càdlàg[4] $(\mathcal{F}_t)$-adapted process with finite variation.

local martingale and semimartingale in R^d

If the process $X(t)$ has values in $\mathbb{R}^d$, it is called a local martingale (resp. a semimartingale) if each of its component is itself a local martingale (resp. a semimartingale).

Example 2.1.3. Consider the process X given by

$$X(t) = X(0) + \int_0^t Y(s)\, ds + W(t) \tag{2.1.2}$$

where W is a $(\mathcal{F}_t)$-Brownian motion[5] and Y is a càdlàg $(\mathcal{F}_t)$-adapted process such that $\int_0^t |Y(s,\omega)|\, ds < \infty$ (a.s.) for every $t \geq 0$. Then X is a $(\mathcal{F}_t)$-semimartingale. ♦

Remark 2.1.4. Whenever the context is clear, we will drop the mention of the underlying filtration $(\mathcal{F}_t)$ and only write, for instance, "$X(t)$ is a semimartingale".

The following notion of *quadratic variation* of a semimartingale will also be of utmost importance to define stochastic integrals and the stochastic chain rule in the next section.

quadratic variation $[X]_t$

Definition 2.1.5 (Quadratic variation)**.** Let X be a continuous semimartingale with decomposition $X(t) = X(0) + M(t) + A(t)$. Then its *quadratic variation* is defined as the unique almost-surely right-continuous and increasing process $[X]_t$ such that $M^2(t) - [X]_t$ is a local martingale.

The existence and uniqueness of the process $[X]_t$ in Definition 2.1.5 is ensured by Theorem 4.9 of Le Gall, 2016, which also gives another equivalent characterization whenever X is continuous.

An important special case is the quadratic variation of a Brownian motion. Indeed, if W is a $(\mathcal{F}_t)$-Brownian motion, then it is in particular a continuous $(\mathcal{F}_t)$-martingale so its quadratic variation exists, and is readily given by $[W]_t = t$[6].

quadratic covariation $[X,Y]_t$

quadratic variation in R^d

Given two continuous semimartingales $X(t)$ and $Y(t)$, we also define their *quadratic covariation* or *bracket* as the process $[X,Y]_t = \frac{1}{4}([X+Y]_t - [X-Y]_t)$. This binary operator is in particular bilinear and symmetric, and enjoys many interesting properties (see for instance Subsection 4.4 of Le Gall, 2016). For a d-dimensional continuous semimartingale X, we then define its quadratic variation as the process $[X]_t$ with values in $R^{d\times d}$ given by $[X]_t = \left[[X_i, X_j]_t \right]_{1\leq i,j\leq d}$.

4 Right-continuous with left-limits.

5 A $(\mathcal{F}_t)$-Brownian motion is the unique continuous stochastic process W such that *a)* W is $(\mathcal{F}_t)$-adapted; *b)* $W(t) - W(s)$ is independent of $\mathcal{F}_s$ whenever $t \geq s \geq 0$; and *c)* $W(t) - W(s) \sim \mathcal{N}(0, t-s)$.

6 Follows from the fact that $W(t)^2 - t$ is a martingale, which is shown using standard computations.

2.2 STOCHASTIC INTEGRAL

The main goal of stochastic analysis is to give a sense to the *stochastic integral* $\int_0^t f(s,\omega)\,dX(s,\omega)$ of a stochastic process $(t,\omega) \mapsto f(t,\omega)$ with respect to some continuous[7] semimartingale X. However, in many cases the integrator X is differentiable almost nowhere and has infinite variation over every time interval (take for instance a Brownian motion), and so the classical techniques of Riemann-Stieltjes integration cannot be used to define such an integral pathwise.

Following the development in Karatzas and Shreve (1998) or Le Gall (2016), it is possible to extend the Riemann-Stieltjes framework to construct such stochastic integrals, as long as the integrand process follows some regularity conditions with respect to the semimartingale.

More precisely, if X is a $(\mathcal{F}_t)$-semimartingale with decomposition $X(t) = X(0) + M(t) + A(t)$, this integral construction is well-defined for integrands $f : [0,\infty) \times \Omega \to \mathbb{R}$ such that

1. the process $f(t,\cdot)$ is $(\mathcal{F}_t)$-progressive;
2. $\mathbb{E}\left[\int_0^t |f(s)|^2\, d[M]_s\right] < \infty$[8] for every $t \geq 0$.

We denote $L_{ad}(X)$ this class of functions, and write the stochastic integral of any process $f \in L_{ad}(X)$ with respect to X as $\int_0^t f(t,\omega)\,dX(s)$ or simply $\int_0^t f\,dX$.

stochastic integral for $f \in L_{ad}(X)$

This construction can also be generalized to integrate $\left(\mathbb{R}^{r\times d}\right)$-valued functions with respect to a d-dimensional semimartingale. Indeed, if X is a d-dimensional semimartingale and if $f : [0,\infty) \times \Omega \to \mathbb{R}^{r\times d}$ is a function such that $f_{ij} \in L_{ad}(X_j)$ for all i and j, then we define its stochastic integral with respect to X as the r-dimensional random vector $\int_0^t f(s,\omega)\,dX(s) \equiv \left[\sum_{j=1}^d \int_0^t f_{ij}(s,\omega)\,dX_j(s)\right]_{1\leq i\leq r}$. By a small abuse of notation, we will also write $f \in L_{ad}(X)$ for integrands verifying the multidimensional integrability condition.

We provide below some fundamental consistency properties of stochastic integrals.

Proposition 2.2.1 (Properties of stochastic integral)**.** *Let X, Y be continuous $(\mathcal{F}_t)$-semimartingales and let $f, g \in L_{ad}(X)$. Then :*

1. *(Consistency) $(t,\omega) \mapsto \int_0^t f(s,\omega)\,dX(s,\omega)$ is a continuous $(\mathcal{F}_t)$-semimartingale.*
2. *(Linearity) If $a, b \in \mathbb{R}$, then $\int_0^t (af + bg)\,dX = a\int_0^t f\,dX + b\int_0^t g\,dX$ (a.s.) and $\int_0^t f\,d(aX + bY) = a\int_0^t f\,dX + b\int_0^t f\,dY$ (a.s.).*

Proof. See Definition 5.7 and following properties of Le Gall (2016). □

Extending classical conventions, a process of the form $X(t) = X(0) + \int_0^t f(s,\omega)\,ds + \int_0^t g(s,\omega)\,dZ$ where Z is a continuous semimartingale will also be often written in the *stochastic differential form* $dX = f(t,\omega)dt + g(t,\omega)dZ$.

stochastic differential form

As stochastic integrals are themselves continuous semimartingales, they also admit a quadratic variation given by the following proposition.

Proposition 2.2.2. *Let X be a continuous d-dimensional semimartingale and let*

quadratic variation of stochastic integral

7 In fact, stochastic integration can also be extended to any semimartingale even if not continuous, but we do not delve into this generalization to avoid having to decompose in every formula the quadratic variation as a continuous and a purely discontinuous part. The interested reader is refereed to Kuo (2006) or Applebaum (2009) for a discussion of such extensions.

8 Here the integral is taken pathwise in a Riemann-Stieltjes sense, which is possible because the quadratic variation $[M]_t$ is increasing by definition.

$f \in L_{ad}(X)$ a function with values in $\mathbb{R}^{r\times d}$. Then

$$\left[\int_0^t f(s,\omega)dX(s)\right]_t = \left[\sum_{p=1}^d \sum_{q=1}^d \int_0^t f_{ip}(s,\omega) f_{jq}(s,\omega)\, d\left[X_p, X_q\right]_s\right]_{1\le i,j\le r} \tag{2.2.1}$$

Proof. See Proposition 3.2.24 of Karatzas and Shreve (1998) for continuous 1-dimensional local martingales, which extends easily to 1-dimensional semimartingales. To treat the case when X is a d-dimensional semimartingale and f a function with values in $\mathbb{R}^{r\times d}$, we use the bilinearity of the quadratic covariation and the result in one dimension to write

$$\left[\int_0^t f(s,\omega)dX(s)\right]_t = \left[\left[\sum_{p=1}^d \int_0^t f_{ip}(s,\omega)\, dX_p(s), \sum_{q=1}^d \int_0^t f_{jq}(s,\omega)\, dX_q(s)\right]_t\right]_{1\le i,j\le r} \tag{2.2.2}$$

$$= \left[\sum_{p=1}^d \sum_{q=1}^d \left[\int_0^t f_{ip}(s,\omega)\, dX_p(s), \int_0^t f_{jq}(s,\omega)\, dX_q(s)\right]_t\right]_{1\le i,j\le r} \tag{2.2.3}$$

$$= \left[\sum_{p=1}^d \sum_{q=1}^d \int_0^t f_{ip}(s,\omega) f_{jq}(s,\omega)\, d\left[X_p, X_q\right]_s\right]_{1\le i,j\le r} \tag{2.2.4}$$

which proves the statement. □

As for stochastic integration with respect to Brownian motion, we can show that the stochastic integral can even be a martingale whenever the underlying integrating process is a martingale itself; as it is stated in the following theorem.

stochastic integral is a martingale

Theorem 2.2.3 (Stochastic integral as a martingale)**.** *Let M be a continuous d-dimensional $(\mathcal{F}_t)$-martingale and let $f \in L_{ad}(M)$ with values in $\mathbb{R}^{r\times d}$. Then $\int_0^t f(s,\omega)\, dM(s)$ is a continuous r-dimensional $(\mathcal{F}_t)$-martingale.*

Proof. See Theorem 6.5.8 of Kuo (2006) for the one-dimensional case, which can be used to prove the result in the multidimensional setting by noticing that the stochastic integral is a sum of one-dimensional stochastic integrals, and so is a sum of martingales with respect to the same filtration. □

Example 2.2.4. Suppose that X is a r-dimensional stochastic process given by

$$X(t) = X(0) + \int_0^t f(s,\omega)\, ds + \int_0^t g(s,\omega)\, dY(s) \tag{2.2.5}$$

where Y is a continuous d-dimensional martingale, $g \in L_{ad}(Y)$ with values in $R^{r\times d}$, and f is a continuous r-dimensional adapted process such that $\int_0^t |f(s,\omega)|\, ds < \infty$ (a.s.) for every $t \ge 0$. Then Theorem 2.2.3 yields that X is a continuous semimartingale with martingale part $M(t) = \int_0^t g(s,\omega)\, dY(s)$ and finite variation part $A(t) = \int_0^t f(s,\omega)\,ds$. In particular, Proposition 2.2.2 also tells us that its quadratic variation is given by

$$[X]_t = [M]_t = \left[\sum_{p=1}^{d}\sum_{q=1}^{d}\int_0^t g_{ip}(s,\omega)g_{jq}(s,\omega)\,d\left[Y_p, Y_q\right]_s\right]_{1\le i,j\le r}. \tag{2.2.6}$$

♦

We are now ready to state the generalization of *Itô's formula*, one of the most important result of stochastic analysis, for continuous semimartingales. In particular, it can be seen as the stochastic version of the classical chain rule for differentiating composition of functions.

Itô's formula

Theorem 2.2.5 (Itô's formula for continuous semimartingale)**.** *Let $X(t)$ be a continuous d-dimensional semimartingale and let $f\colon [0,\infty)\times\mathbb{R}^d \longrightarrow \mathbb{R}$ be a function of class C^2. Then*

$$df(t,X) = \frac{\partial f}{\partial t}(t,X)dt + \sum_{i=1}^{d}\frac{\partial f}{\partial x_i}(t,X)dX_i + \frac{1}{2}\sum_{i=1}^{d}\sum_{j=1}^{d}\frac{\partial^2 f}{\partial x_i \partial x_j}(t,X)d\left[X_i, X_j\right]. \tag{2.2.7}$$

Proof. See Theorems 3.3.3 and 3.3.6 of Karatzas and Shreve (1998). □

Remark 2.2.6. To simplify the manipulations and to make the link with the formula for Itô processes that is mostly used in applied texts, we will often use the abuse of notation $d\left[X_i, X_j\right] = dX_i \cdot dX_j^T$. This can be done without harm as the quadratic variation is also symmetric bilinear and verifies the usual formal product rule

$$\begin{cases} dt \cdot dt & = 0 \\ dt \cdot dW(t) & = 0 \\ dW(t)\cdot dW(t) & = dt \end{cases} \tag{2.2.8}$$

whenever $W(t)$ is a standard Brownian motion. Moreover, the result of Proposition 2.2.2 can then be written with this notation as

$$d\left[\int_0^t f(s,\omega)\,dX(s)\right]_t = f(t,\omega)dX\cdot dX^T f(t,\omega)^T \tag{2.2.9}$$

which often leads to easier computations of quadratic variations.

An important use of stochastic analysis is to allow the study of random continuous perturbations in ordinary differential equations. To be more precise, if Z is a continuous semimartingale, an equation of the form $dX(t) = f(t, X(t))dZ(t)$ is called a *stochastic differential equation* on X. This can be seen as a perturbed generalization of a differential equation of the form $\dot{x} = f(t,x)$. The following theorem shows when the above stochastic differential has a meaning (i.e., when the right-hand side is well-defined) and admits a unique solution.

stochastic differential equation

solutions to stochastic differential equations

Theorem 2.2.7 (Strong solution)**.** *Let Z be a continuous d-dimensional semimartingale and let $f\colon [0,\infty)\times\mathbb{R}^r \longrightarrow \mathbb{R}^{r\times d}$ be a function satisfying :*

1. *f is Lipschitz in its second variable;*
2. *for every $y \in \mathbb{R}^d$, $f(\cdot, y)$ is locally bounded.*

Then for all $x \in \mathbb{R}^r$, there exists a unique (up to indistinguishability[9]) continuous process X such that

$$X(t) = x + \int_0^t f(s, X(s))dZ(s). \tag{2.2.10}$$

Moreover, X is adapted to the filtration $(\mathcal{F}_t^Z)$ generated by Z.

Proof. See Theorem IX.2.1 of Revuz and Yor (1999). □

strong solution

If a stochastic process X verifies the conclusion of Theorem 2.2.7, we say that it is a *strong solution* of the stochastic differential equation (2.2.10).

Example 2.2.8. Let M be a d-dimensional continuous martingale and let $a\colon [0,\infty) \times \mathbb{R}^r \to \mathbb{R}$, $b\colon [0,\infty) \times \mathbb{R}^r \to \mathbb{R}^d$ be two arbitrary functions. Assume that we want to prove that the stochastic differential equation

$$dX = a(t,X)dt + b(t,X)^T dM \tag{2.2.11}$$

admits a unique strong solution. First, we notice that t and $M(t)$ are both continuous semimartingales, and so $Z(t) = \begin{bmatrix} t \\ M(t) \end{bmatrix}$ is a $(d+1)$-dimensional continuous semimartingale. Introducing the function $f(t,x) = \begin{bmatrix} a(t,x) & b(t,x) \end{bmatrix}$, we can therefore write $f(t,X(t))dZ(t) = a(t,X(t))dt + b(t,X(t))^T dM(t)$. In other words, as long as we have

$$|a(t,x) - a(t,x')| + \|b(t,x) - b(t,x)\| \le K\|x - x'\| \tag{2.2.12}$$

and the local boundeness of a and b, Theorem 2.2.7 yields that there exists a unique strong solution to (2.2.11). ♦

To finish this section, we state another result that will be used extensively in the rest of our work, accordingly that all martingales can be expressed as time-changed Brownian motions.

time-change theorem

Theorem 2.2.9 (Time-change theorem for continuous local martingales)**.** *Let M be a continuous local martingale. Then there exists a one-dimensional Brownian motion W (defined on a possibly enlarged probability space) such that $M(t) = W([M]_t)$ (a.s.) for all $t \geq 0$.*

Proof. See Theorem 3.4.6 and Problem 3.4.7 of Karatzas and Shreve (1998). □

An important consequence of Theorems 2.2.3 and 2.2.9 is that any stochastic integral with respect to a continuous martingale can be expressed as a time-changed Brownian motion. In Corollary 2.3.4, we will see how this fact can be used to prove growth rates of stochastic integrals.

2.3 CORRELATED BROWNIAN MOTIONS & GROWTH RATES

In the rest of this manuscript, we will be mainly interested in stochastic differential equations for which the noise term is of the form $Z(t,\omega)^T dW(t)$ where W is a d-dimensional random vector whose components are $(\mathcal{F}_t)$-Brownian motions with pairwise correlation matrix C, and Z is a bounded continuous process in $\mathbb{R}^d$ adapted to the filtration $(\mathcal{F}_t)$. In particular, Z is therefore in $L_{ad}(W)$.

9 Meaning that if Y is another such process, then $\mathbb{P}(X(t) = Y(t) \text{ for all } t \geq 0) = 1$.

In the following proposition, we show that such a random vector W of correlated Brownian motions indeed exists. This kind of result is well-known for generating correlated Gaussian random variables, but a proof for correlated Brownian motions as it is stated here is not that easy to find in the literature, so we provide one that is customized for our needs.

Proposition 2.3.1. *Let C be a symmetric definite $(d \times d)$-matrix with values in $[-1, 1]$ and with diagonal made of 1's. Then there exists a filtration $(\mathcal{F}_t)$ and a random vector $W = (W_1, \ldots, W_d)$ such that*

existence of correlated Brownian motions

1. *$Cor(W_i(t), W_j(t)) = C_{ij}$ for all $t \geq 0$;*
2. *W_i is a $(\mathcal{F}_t)$-Brownian motion for all $i = 1, \ldots, d$.*

Moreover, the quadratic variation of W is given by $[W]_t = Ct$.

Proof. As C is symmetric definite, we can use Cholesky's factorization to obtain an inferior triangular $d \times d$ matrix L such that $C = LL^T$. Let $W' = (W'_1, \ldots, W'_d)$ be a d-dimensional Brownian motion. We will show that the vector $W = LW'$ and the filtration $(\mathcal{F}_t)$ generated by W' verify the conditions of the proposition.

First, we have immediately that $W_i(0) = 0$ and that W_i is continuous for any $i = 1, \ldots, d$. Furthermore, $W(t) - W(s) = L[W'(t) - W'(s)]$, so $W(t) - W(s)$ is independent of $\mathcal{F}_s$ whenever $s < t$, and $W(t) - W(s) \sim \mathcal{N}(0, (t - s)C)$ by linear transformation of a Gaussian process. It follows that each component of W is a $(\mathcal{F}_t)$-Brownian motion verifying the correlations described by C.

As W is a linear transform of the d-dimensional Brownian motion W' which has quadratic variation $[W']_t = tI_n$, we obtain immediately from the bilinearity of the quadratic covariation that

$$[W]_t = [LW']_t = [LW', LW']_t = L[W', W']L^T = LtI_d L^T = Ct, \tag{2.3.1}$$

hence proving the second statement of the proposition. □

Remark 2.3.2. From now on, whenever we talk about correlated Brownians motions with prescribed pairwise correlation matrix, we will always assume that they are generated as in Proposition 2.3.1, i.e., are Brownian motions under the same underlying filtration. In this case, we will also commonly drop the reference to the filtration whenever the context is clear of ambiguity.

Now that we have proved that such a vector of correlated Brownian motions exist, we may wonder how fast a stochastic integral of the form $\int_0^t Z^T dW$ grows. This will be useful in future chapters to show that the amplitude of the random perturbations are negligible with respect to the deterministic drift in our stochastic differential equations.

But before proving this rate, we first provide for completeness a classical result on the growth of Brownian motions : the so-called *law of iterated logarithm*.

Lemma 2.3.3 (Law of iterated logarithm)**.** *If W is a Brownian motion, then*

law of iterated logarithm

$$\limsup_{t \to \infty} \frac{|W(t)|}{\sqrt{2t \log \log t}} = 1 \quad (a.s.). \tag{2.3.2}$$

Proof. See Corollary 3.2 of Baldi (2017), who also proves many other versions of this law. □

We are now ready to prove our main result on the growth rate of stochastic integrals with respect to correlated Brownian motions.

growth of stochastic integral

Corollary 2.3.4. *Let W be a d-dimensional random vector whose components are standard $(\mathcal{F}_t)$-Brownian motions with pairwise correlation matrix C, and let Z be a bounded continuous process in $\mathbb{R}^d$ adapted to the filtration $(\mathcal{F}_t)$. Then,*

$$f(t) + \int_0^t Z(s)^T dW(s) \sim f(t) \quad (a.s.) \tag{2.3.3}$$

for all function $f\colon [0,\infty) \to \mathbb{R}$ such that $\dfrac{f(t)}{\sqrt{t \log\log t}} \to \infty$.

Proof. Let ξ be the stochastic process defined as $\xi(t) = \int_0^t Z^T dW$. By Proposition 2.2.2 and using notations of Remark 2.2.6, its quadratic variation $[\xi]_t$ is then given by

$$d[\xi]_t = d\xi \cdot d\xi^T = Z(t)^T dW \Big(Z(t)^T dW\Big)^T = Z(t)^T dW dW^T Z(t) = Z(t)^T C Z(t) dt. \tag{2.3.4}$$

where we have used the quadratic variation of W provided by Proposition Proposition 2.3.1. From the assumption that Z is a bounded process and that C is a stationary deterministic matrix, we further have

$$\rho(t) \equiv [\xi]_t = \int_0^t Z(s)^T C Z(s) ds \leq M^2 t \tag{2.3.5}$$

where $M^2 = \sup\limits_{t \geq 0} Z(t)^T C Z(t) < \infty$. As W is a continuous $(\mathcal{F}_t)$-martingale and as $Z \in L_{ad}(W)$, Theorem 2.2.3 then yields that ξ is itself also a continuous 1-dimensional $(\mathcal{F}_t)$-martingale. As such, we can use the time-change theorem (Theorem 2.2.9) to conclude that there exists a standard Brownian motion $\widetilde{W}$ (defined on a possibly enlarged probability space) such that $\xi(t) = \widetilde{W}(\rho(t))$ for all $t \geq 0$. If $\rho(t) \to \rho(\infty) < \infty$, then $\widetilde{W}(\rho(t))$ is bounded, and so in particular $\left|\widetilde{W}(\rho(t))\right| \leq \dfrac{f(t)}{\sqrt{t \log\log t}}$ for t big enough. This yields the limit

$$\frac{f(t) + \xi(t)}{f(t)} = 1 + \frac{\widetilde{W}(\rho(t))}{f(t)} \to 1. \tag{2.3.6}$$

On the other hand, if $\rho(t) \to \infty$, then the law of iterated logarithm (cf. Lemma 2.3.3) states that

$$\limsup_{t\to\infty} \frac{\left|\widetilde{W}(\rho(t))\right|}{\sqrt{2\rho(t)\log\log\rho(t)}} = 1 \quad \text{(a.s.)}, \tag{2.3.7}$$

from which we deduce the limit

$$\begin{aligned} \frac{|\widetilde{W}(\rho(t))|}{f(t)} &= \frac{\left|\widetilde{W}(\rho(t))\right|}{\sqrt{2\rho(t)\log\log\rho(t)}} \times \frac{\sqrt{2\rho(t)\log\log\rho(t)}}{f(t)} & (2.3.8) \\ &\leq \frac{\left|\widetilde{W}(\rho(t))\right|}{\sqrt{2\rho(t)\log\log\rho(t)}} \times 2M \frac{\sqrt{t\log\log t}}{f(t)} \to 0 \quad \text{(a.s.)}, & (2.3.9) \end{aligned}$$

where we have used the fact that $2\rho(t)\log\log\rho(t) \leq 4Mt\log\log t$ for t big enough. $\square$

Part I

STUDY OF STOCHASTIC IMITATION DYNAMICS

3

IMITATION DYNAMICS

After recalling in Section 3.1 the mathematical framework of population games, Section 3.2 introduces a classical way to derive deterministic game dynamics by the use of *imitative protocols*. We then propose in Section 3.3 a general way to consider random perturbations in these game dynamics.

3.1 POPULATION GAMES

In this work, we are interested in the framework of population games with *nonatomic players*. The content of this section is standard in the game theoretic literature, see for instance Sandholm (2010) or Weibull (1995).

More precisely, we are considering a finite set of player *populations* $\mathcal{N} = \{1, \dots, N\}$, for each of which we associate a finite set of *pure actions* or genotypes $\mathcal{A}_k = \{\alpha_{k,1}, \dots, \alpha_{k,A_k}\}$ when $k \in \mathcal{N}$. *action set $\mathcal{A}_k$*

By *nonatomic players*, we then mean that each population $k \in \mathcal{N}$ is made of a continuum of individuals, e.g., can be considered as an interval of $\mathbb{R}$ endowed with Lebesgue measure. This is a way to model populations made of a large number of agents, where each of them taken individually has only a very small impact on the whole game behavior. For instance, it is usually a good approximation of what happens with buyers in an economical setting, or with the reproduction of species in evolutionary biology. *nonatomic players*

During the game, each member of each population chooses a strategy in their action space, so that the *strategic state* of a population $k \in \mathcal{N}$ is given by the distribution $x_k = (x_{k,\alpha})_{\alpha \in \mathcal{A}_k}$ where $x_{k,\alpha}$ is the mass[1] of individuals playing action α. Assuming for simplicity that each population has unit total mass, the strategic states of the k-th population all lie into the simplex $\mathcal{X}_k \equiv \Delta(\mathcal{A}_k)$, that we name accordingly the *state space* of the k-th population. Similarly, we define the state space of the game as the product $\mathcal{X} \equiv \prod_{k \in \mathcal{N}} \mathcal{X}_k$. *state space $\mathcal{X}_k$* *state space of the game $\mathcal{X}$*

A strategic state $x_k \in \mathcal{X}_k$ will also be called a *mixed strategy* for population k, and if such a distribution gives all mass to a pure action $\alpha_k \in \mathcal{A}_k$ (i.e., if x_k is a vertex of $\mathcal{X}_k$), then it will be called a *pure strategy* and will also be denoted as α_k for simplicity. *mixed strategy $x_k \in \mathcal{X}_k$* *pure strategy $\alpha_k \in \mathcal{A}_k$*

Given a strategic state $x \in \mathcal{X}$ of the game, the total payoff of the individuals in a population $k \in \mathcal{N}$ playing the pure action $\alpha \in \mathcal{A}_k$ is determined by $v_{k\alpha}(x)$, where $v_{k\alpha} \colon \mathcal{X} \to \mathbb{R}$ is named the *payoff function* and is assumed to be Lipschitz. Accordingly, the average payoff to population k can be written as $\sum_{\alpha}^{k} x_{k\alpha} v_{k\alpha}(x) = \langle v_k(x), x_k \rangle$ where $v_k(x) \equiv (v_{k\alpha}(x))_{\alpha \in \mathcal{A}_k}$ is the k-th population's *payoff vector* . *payoff function $v_{k\alpha}$* *payoff vector v_k*

1 Here *mass* should be understood as the Lebesgue measure of the continuum of individuals playing action α.

All in all, the player populations $\mathcal{N}$, their associated pure action spaces $\mathcal{A} \equiv \prod_k \mathcal{A}_k$ and the payoff functions $v \equiv (v_{k\alpha})_{k,\alpha}$ constitute what we define as the *population game* $\Gamma(\mathcal{N}, \mathcal{A}, v)$.

population game $\Gamma(\mathcal{N}, \mathcal{A}, v)$

dominated strategies

Fixing a population $k \in \mathcal{N}$, we say that the mixed strategy $p_k \in \mathcal{X}_k$ is *dominated* by $p'_k \in \mathcal{X}_k$ if

$$\langle v_k(x), p_k \rangle < \langle v_k(x), p'_k \rangle \quad \text{for all } x \in \mathcal{X}, \tag{3.1.1}$$

i.e., the strategy profile p_k always achieves a lesser average payoff than p'_k, whatever the strategies of the other player populations are. In particular, a pure strategy $\alpha \in \mathcal{A}_k$ is dominated by another pure strategy $\beta \in \mathcal{A}_k$ if

$$v_{k\alpha}(x) < v_{k\beta}(x) \quad \text{for all } x \in \mathcal{X}. \tag{3.1.2}$$

Nash equilibrium

A strategic state $x^* \in \mathcal{X}$ is said to be a *Nash equilibrum* of the game if it is unilaterally stable for all populations, i.e.,

$$v_{k\alpha}(x^*) \geq v_{k\beta}(x^*) \quad \text{for all } \alpha \in \operatorname{supp}(x^*_k) \text{ and for all } \beta \in \mathcal{A}_k,\ k \in \mathcal{N}; \tag{EQ}$$

support $\operatorname{supp}(x)$)

where $\operatorname{supp}(x^*_k) \equiv \{\alpha \in \mathcal{A}_k : x^*_{k\alpha} > 0\}$ is the support of x^*_k in the simplex $\mathcal{X}_k$. Equivalently, we can rewrite this condition as a Stampacchia variational inequality of the form

$$\langle v_k(x^*), x^*_k \rangle \geq \langle v_k(x^*), x_k \rangle \quad \text{for all } x_k \in \mathcal{X}_k,\ k \in \mathcal{N}. \tag{3.1.3}$$

strict Nash equilibrium

When the inequality in (EQ) (resp. in (3.1.3)) is strict for all actions $\beta \neq \alpha$ (resp. for all states $x_k \neq x^*_k$), then x^* will be called a *strict Nash equilibrium*. In particular, it will then be *pure*; in the sense that x^*_k is a pure strategy for every population $k \in \mathcal{N}$, or in other words that x^* is a vertex of $\mathcal{X}$.

Remark 3.1.1. To avoid dealing with double subindices, in the remaining of the manuscript we will drop completely the population index k, hence focusing ourselves on single-population games. This can be done without any loss of generality, as all of our results still apply as stated in a multi-population setting.

3.2 THE DETERMINISTIC MODEL

Following the development in Sandholm (2010), a general and meaningful way to derive game dynamics from micro-foundations is by the use of a so called *revision protocol*. The underlying idea behind this construction is to assume that each individual can revise their strategies at random times, based on the strategies and payoffs of all the population.

game dynamics from revision protocol

To be more precise, let Γ be a single-population game, and assume that each (nonatomic) individual has a clock attached to them. Independently from an individual to the other, this clock will ring at random times based on a Poisson process (i.e., the intertimes are distributed as exponential random variables of constant rate), at which the player will choose (or not) to revise their strategy. If a nonatomic individual was playing strategy $\alpha \in \mathcal{A}$ in a population strategic state $x \in \mathcal{X}$, they will then decide to play strategy $\beta \in \mathcal{A}$ instead with probability proportional to $\rho_{\alpha\beta}(v(x), x)$, where $\rho_{\alpha\beta}$ is called the *conditional switch rate* from α to β.

conditional switch rate $\rho_{\alpha\beta}$

Assuming the whole state of the population do not change "too much" in an infinitesimal time interval dt, the relative mass $dx_{\alpha\beta}$ of individuals revising their play from strategy α to β over dt can then be approximated by

$$dx_{\alpha\beta} = x_\alpha(t)\rho_{\alpha\beta}(v(x(t)), x(t))dt. \tag{3.2.1}$$

For conciseness and notational clarity, we will commonly drop the dependence on time and on the variables when writing the game state and the conditional switch rate, i.e., writing $\rho_{\alpha\beta}$ instead of the heavy full expression $\rho_{\alpha\beta}(v(x(t)), x(t))$.

Accordingly, we define the *revision protocol dynamics* as the mean dynamics driving the evolution of the relative mass of individuals playing strategy $\alpha \in \mathcal{A}$, given by

revision protocol dynamics

$$\dot{x}_\alpha = \sum_\beta \frac{dx_{\beta\alpha}}{dt} - \sum_\beta \frac{dx_{\alpha\beta}}{dt} = \sum_\beta x_\beta \rho_{\beta\alpha} - x_\alpha \sum_\beta \rho_{\alpha\beta}. \tag{3.2.2}$$

In this work, we are interested in a particular form of the dynamics that is generated by so called *imitative* revision protocols. They are specified by the conditional switch rates

imitative revision protocols

$$\rho_{\alpha\beta} = x_\beta r_{\alpha\beta}(v, x), \tag{3.2.3}$$

where an individual imitates the strategy of a randomly chosen opponent with probability proportional to the *conditional imitation rate* $r_{\alpha\beta}$ (assumed to be Lipschitz). In this case, the associated revision dynamics can be written in the compact form

conditional imitation rate $r_{\alpha\beta}$

imitation dynamics (ID)

$$\dot{x}_\alpha = x_\alpha \sum_\beta x_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right], \tag{ID}$$

which is then called an *imitation dynamics*.

We further assume for now that the conditional imitation rates $r_{\alpha\beta}$ are *monotone* in the sense of Sandholm (2010), i.e., that for all states $x \in \mathcal{X}$ and for all pure strategies $\alpha, \beta, \gamma \in \mathcal{A}$,

monotone condition

$$v_\alpha(x) < v_\beta(x) \Leftrightarrow r_{\gamma\alpha} - r_{\alpha\gamma} < r_{\gamma\beta} - r_{\beta\gamma}. \tag{3.2.4}$$

Under this condition, dominated strategies are sure to become extinct along the trajectories of (ID) (see e.g., Akin, 1980 and Nachbar, 1990). This is not necessarily the case in non monotone imitation dynamics where such strategies can survive indefinitely in some well-chosen games, as was shown by Mertikopoulos and Viossat (2022).

Example 3.2.1 (Replicator Dynamics)**.** An important example of a well-know dynamics generated by imitation is the *replicator dynamics*

$$\dot{x}_\alpha = x_\alpha\Big(v_\alpha - \sum_\beta x_\beta v_\beta\Big), \tag{RD}$$

introduced originally by Taylor and Jonker (1978) as a model of natural selection in biology[2]. It can be derived from at least three imitation protocols :

- (imitation of success) $\rho_{\alpha\beta} = x_\beta(K + v_\beta)$ where K is big enough so that $\rho_{\alpha\beta}$ stays positive;
- (imitation driven by dissatisfaction) $\rho_{\alpha\beta} = x_\beta(K - v_\alpha)$ where K is again big enough;
- (pairwise imitation) $\rho_{\alpha\beta} = x_\beta[v_\beta - v_\alpha]_+$ where $[\cdot]_+$ denotes the positive part.

♦

The long-run behavior of monotone imitation dynamics is quite well understood in the literature, for instance the "folk theorem" of evolutionary game

2 We will see in Chapter 6 how this dynamics can be derived from an evolutionary biology point of view as well as from a learning approach.

theory (cf. Hofbauer and Sigmund, 2003) states that *a*) Lyapunov stable states are Nash; *b*) limits of interior trajectories are Nash; and *c*) strict Nash equilibriums are asymptotically stable under (ID).

3.3 A STOCHASTIC APPROACH

incertitude in payoffs

In the approach adopted in the previous subsection, the construction of the dynamics lays on a strong assumption that is rarely met in practice : each individual knows exactly their and others' payoffs at every time, without any influence of exogenous random perturbations. However, in many applications this is far too much of a stretch, as in the study of populations in biology where the reproductive fitness of a species can be heavily influenced by the random shocks of nature (Fudenberg and Harris, 1992), or in telecommunication networks for which noise is an intrinsic physical phenomenon due to interference.

Following the stochastic construction initiated in Mertikopoulos and Viossat (2016) for replicator dynamics governed by imitation protocols, a first idea would then be to assume that the players' payoffs at time t are perturbed by *random shocks*, i.e., $\hat{v}_\alpha(t) = v_\alpha(x(t)) + \xi_\alpha(t)$ where ξ_α is a formal "white noise" process. In the case of imitations protocols leading to the replicator dynamics, for instance the imitation of success $r_{\alpha\beta} = K + v_\beta$ (see Example 3.2.1), we then obtain immediately a linearly perturbed version $\hat{r}_{\alpha\beta} = r_{\alpha\beta} + \xi_\beta$ of the imitation rates when introducing these payoff shocks. Unfortunately, in general (monotone) imitation protocols, we do not make any assumption on the form of the conditional imitation rates $r_{\alpha\beta}$ nor on their dependence on the payoff vector, and as such we cannot easily deduce the explicit expression of the perturbed rates $\hat{r}_{\alpha\beta}$ with respect to noise.

incertitude on conditional rates

However, we notice that the payoffs' influence on the whole imitation dynamics (ID) comes only through these conditional imitations rates, and as such we can reasonably assume that the incertitude on the payoffs is well approximated by an incertitude about the rates themselves. In other words, instead of considering payoff shocks $\hat{v}$, we assume that the rates at time t are perturbed as $\hat{r}_{\alpha\beta} = r_{\alpha\beta} + \xi_{\alpha\beta}$, where the white noise $\xi_{\alpha\beta}$ now depends possibly on both the source and the target strategies in order to estimate the involvement of $\hat{v}_\alpha$ and $\hat{v}_\beta$.

In formal Langevin notation[3], the imitation dynamics (ID) then becomes

$$\dot{x}_\alpha = x_\alpha \sum\nolimits_\beta x_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right] + x_\alpha \sum\nolimits_\beta x_\beta \left[\xi_{\beta\alpha} - \xi_{\alpha\beta}\right] \tag{3.3.1}$$

stochastic imitation dynamics (SID)

or, in stochastic differential equation form,

$$dX_\alpha = X_\alpha \sum\nolimits_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right] dt + X_\alpha \sum\nolimits_\beta X_\beta \left[\sigma_{\beta\alpha}(X) dW_{\beta\alpha} - \sigma_{\alpha\beta}(X) dW_{\alpha\beta}\right] \tag{SID}$$

where $\sigma_{\alpha\beta} \colon \mathcal{X} \to \mathbb{R}$ are diffusion coefficients assumed to be Lipschitz, and $W_{\alpha\beta}$ are standard Brownian motions with pairwise correlations given by the function $C \colon \mathcal{A}^4 \to [-1, 1]$, i.e., such that

$$Cor\left(W_{\alpha\beta}(t), W_{\alpha'\beta'}(t)\right) = C(\alpha, \beta; \alpha'\beta') \quad \text{for all } \alpha, \beta, \alpha', \beta' \in \mathcal{A}. \tag{3.3.2}$$

3 Langevin notation is an abuse of notation usually used in physical fields, where a "white noise" process $\xi(t)$ should be understood as the time derivative $\frac{dW}{dt}$ of a Brownian motion. However Brownian motions are nowhere differentiable, so this derivative does not exist in the classical sense. In fact, it can be shown that it exists in some weak sense using the theory of distributions, but we do not enter into such depths here.

As the function C defines correlations, we obviously also need that C is symmetric in the sense that $C(\alpha,\beta;\alpha',\beta') = C(\alpha',\beta';\alpha,\beta)$ for all $\alpha,\beta,\alpha',\beta' \in \mathcal{A}$.

The stochastic imitation dynamics (SID) will be our main focus in this part. As such, we first provide some important examples of dynamics that we can express in this framework.

Example 3.3.1 (Source vs. target noise)**.** In the construction of (SID), we have assumed that the conditional imitations rates are perturbed by a white noise $\xi_{\alpha\beta}$ depending both on the source and the target strategies. However, there may be some cases when the rates $r_{\alpha\beta}$ depend either only on the payoff of the source strategy (e.g., imitation driven by dissatisfaction) or only on the target strategy (e.g., imitation of success), and so the same should be true for its perturbation noise.

We can also think about strategic games where we know with certitude our payoff but not about ones of the other players, so that the random shocks should only come from the target payoffs but not from source strategies.

We show below that such situations can also be expressed in our general stochastic imitation model.

(i) (independent source noise) Let $\sigma_{\alpha\beta} = \sigma^s_\alpha$ and $W_{\alpha\beta} = W^s_\alpha$ for all $\alpha,\beta \in \mathcal{A}$, where $\sigma^s_\alpha\colon \mathcal{X} \to \mathbb{R}$ are Lipschitz and W^s_α are independent Brownian motions, i.e., $C(\alpha,\beta;\alpha',\beta') = \delta_{\alpha,\alpha'}$. Then, we have in Langevin notation that $\hat{r}_{\alpha\beta} = r_{\alpha\beta} + \xi^s_\alpha$ where ξ^s_α are independent white noises depending only on the source strategy, and the associated stochastic imitation dynamics is *independent source noise*

$$dX_\alpha = X_\alpha \sum\nolimits_\beta X_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right] dt - X_\alpha \left[\sigma^s_\alpha dW^s_\alpha - \sum\nolimits_\beta X_\beta \sigma^s_\beta dW^s_\beta\right]. \quad (3.3.3)$$

(ii) (independent target noise) Let $\sigma_{\alpha\beta} = \sigma^t_\beta$ and $W_{\alpha\beta} = W^t_\beta$ for all $\alpha,\beta \in \mathcal{A}$, where $\sigma^t_\beta\colon \mathcal{X} \to \mathbb{R}$ are Lipschitz and W^t_β are independent Brownian motions, i.e., $C(\alpha,\beta;\alpha',\beta') = \delta_{\beta,\beta'}$. Then, we have in Langevin notation that $\hat{r}_{\alpha\beta} = r_{\alpha\beta} + \xi^t_\beta$ where ξ^t_β are independent white noises depending only on the target strategy, and the associated stochastic imitation dynamics is *independent target noise*

$$dX_\alpha = X_\alpha \sum\nolimits_\beta X_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right] dt + X_\alpha \left[\sigma^t_\alpha dW^t_\alpha - \sum\nolimits_\beta X_\beta \sigma^t_\beta dW^t_\beta\right]. \quad (3.3.4)$$

We notice that the source and target noises both generate the same stochastic dynamics up to a change of noise magnitude $\sigma^s \leftrightarrow -\sigma^t$, and so should behave in a similar manner. ♦

Example 3.3.2 (Replicator dynamics with payoff shocks)**.** Following in the steps of Example 3.2.1, we know that the imitation of success protocol $r_{\alpha\beta} = K + v_\beta$ leads to the replicator dynamics (RD). As the rates only depend on the target payoff, we can apply the approach of independent target noise explained in Example 3.3.1 to obtain a stochastic version of the replicator dynamics given by *replicator dynamics with payoff shocks*

$$dX_\alpha = X_\alpha \left(v_\alpha - \sum\nolimits_\beta X_\beta v_\beta\right) dt + X_\alpha \left(\sigma^t_\alpha dW^t_\alpha - \sum\nolimits_\beta X_\beta \sigma^t_\beta dW^t_\beta\right). \quad \text{(RDPS)}$$

In fact, this dynamics is exactly the one obtained by Mertikopoulos and Viossat (2016) using a construction based on random payoff shocks. ♦

To study stochastic imitation dynamics, we first need to prove that there are indeed strong solutions to (SID) that stay in $\mathcal{X}$ for all times. To do so, we first provide an important compact form of the stochastic dynamics in the following remark, which will become quite useful in the long-run to avoid dealing with multiple sums.

compact stochastic imitation dynamics (SID′)

Remark 3.3.3. Let $W_\alpha(t) = \begin{bmatrix} W_{\cdot\alpha}(t) \\ W_{\alpha\cdot}(t) \end{bmatrix}$ and $\sigma_\alpha(X) = \begin{bmatrix} X \otimes \sigma_{\cdot\alpha}(X) \\ -X \otimes \sigma_{\alpha\cdot}(X) \end{bmatrix}$ where $\otimes$ denotes the term by term multiplication. Then, the martingale part of stochastic imitation dynamics (SID) can be written into the more compact form

$$dX_\alpha = X_\alpha \sum\nolimits_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right] dt + X_\alpha \sigma_\alpha^T(X) dW_\alpha \tag{SID'}$$

where the diffusion coefficients σ_α are still Lipschitz thanks to the fact that $\mathcal{X}$ is compact. Furthermore, the pairwise correlations between random vectors W_α's are fully characterized by the matrix function $\widehat{C}\colon \mathcal{A}^2 \to \mathcal{M}_{2A\times 2A}(\mathbb{R})$ given by

$$\widehat{C}(\alpha,\beta) = \begin{bmatrix} C(\cdot,\alpha;\cdot,\beta) & C(\cdot,\alpha;\beta,\cdot) \\ C(\alpha,\cdot;\cdot,\beta) & C(\alpha,\cdot;\beta,\cdot) \end{bmatrix}. \tag{3.3.5}$$

For every $\alpha \in \mathcal{A}$, we notice thanks to Proposition Proposition 2.3.1 that the random vector W_α can be seen a $2A$-dimensional vector of correlated Brownian motions under the same underlying filtration, and so is itself a continuous $2A$-dimensional martingale.

interior solution orbits

We can now prove the existence and uniqueness of strong solutions to (SID). In particular, we show that we can even consider *interior solution orbits*, i.e., solutions of (SID) that stay in the relative interior $\mathcal{X}^\circ$ for all times whenever the trajectory also starts from $\mathcal{X}^\circ$.

Proposition 3.3.4. *For all initial condition $X(0) \in \mathcal{X}$, there exists a unique (up to indistinguishability) continuous solution X of* (SID) *such that $X(t) \in \mathcal{X}$ for all $t \geq 0$. Moreover, if $X(0) \in \mathcal{X}^\circ$ then $X(t) \in \mathcal{X}^\circ$ for all $t \geq 0$.*

Proof. First, let us write (SID′) into the more general form

$$dX_\alpha = a_\alpha(X)dt + b_\alpha^T(X)dW_\alpha \tag{3.3.6}$$

where $a_\alpha(x) = x_\alpha \sum_\beta x_\beta[r_{\beta\alpha} - r_{\alpha\beta}]$ and $b_\alpha(x) = x_\alpha\sigma_\alpha$. From the end of Remark 3.3.3 we know that W_α is a martingale, so we are in the setup of Theorem 2.2.7; and more precisely in the one of Example 2.2.8.

As a_α and b_α are constructed from sums and products of Lipschitz functions, they can both be shown to be Lipschitz on any compact subset $\mathcal{K}$ of $\mathbb{R}^A$. However they are not necessarily globally Lipschitz[4] on $\mathbb{R}^A$, which is needed for the existence and uniqueness of a strong solution.

Even if the coefficient do not verify the Lipschitz condition, it still possible to prove the existence of solutions to (3.3.6). To do so, let ϕ be a smooth bump function that is equal to 1 on an open $\mathcal{U} \supset \mathcal{X}$ and equal to 0 outside of a compact $\mathcal{K} \supset \mathcal{U}$. Then, the smoothed stochastic differential equation

$$dX_\alpha = \phi(X)\Big(a_\alpha(X)dt + b_\alpha^T(X)dW_\alpha\Big) \tag{3.3.7}$$

4 In general the product of Lipschitz functions is not itself Lipschitz, unless the functions are bounded on the underlying space. Take for instance $f(x) = x$ which is globally Lipschitz on $\mathbb{R}$ whereas $f^2(x) = x^2$ is not.

has globally Lipschitz coefficients[5] on $\mathbb{R}^A$, and thus admits a unique strong solution $X(t)$ for every initial condition $X(0)$. Now, notice that (3.3.6) coincides with (3.3.7) on $\mathcal{X}$, and that $d(\sum_\alpha X_\alpha) = 0$ whenever X is a solution of (3.3.6). Therefore, if the solution to the smoothed dynamics (3.3.7) starts at initial position $X(0) \in \mathcal{X}$ then it will always stays in the simplex $\mathcal{X}$, and so it is easy to conclude that our former equation (3.3.6) has also a unique solution staying in the simplex $\mathcal{X}$.

Furthermore, if the initial condition $X(0)$ is in $\mathcal{X}^\circ$, then Itô's formula (Theorem 2.2.5) applied to the continuous semimartingale X with $f(x) = \log x$ can be used[6] to show that

$$X_\alpha(t) = X_\alpha(0) \exp\{M_\alpha(t)\} \tag{3.3.8}$$

where M_α is some continuous semimartingale depending on X. In particular, $X_\alpha(t) > 0$ for all $\alpha \in \mathcal{A}$ whenever $X(0) \in \mathcal{X}^\circ$, which proves the second statement. □

We end this chapter by mentioning an important consequence of the existence of solutions, accordingly that the right-hand side of (SID) is a continuous semimartingale, allowing us to use classical stochastic analysis tools that we have presented in Chapter 2.

Remark 3.3.5. As σ_α is Lipschitz continuous on the compact $\mathcal{X}$, it is also bounded so that $(t,\omega) \mapsto X(t,\omega)\sigma_\alpha(X(t,\omega)) \in L_{ad}(W_\alpha)$ whenever X is a strong solution to (SID) in $\mathcal{X}$, which always exists by Proposition 3.3.4. Therefore, Theorem 2.2.3 yields that $\int_0^t X_\alpha(s)\sigma_\alpha^T(X(s))\, dW_\alpha(s)$ is a continuous martingale, and an analogous argument can be carried on to prove that X is in this case a continuous semimartingale as the one shown in Example 2.2.4.

---

5 It follows from the more general fact that if ϕ is a globally Lipschitz function compactly supported on $\mathcal{K} \subset \mathbb{R}^A$ and f is a locally Lipschitz function on $\mathbb{R}^A$, then the product ϕf is globally Lipschitz on $\mathbb{R}^A$.

6 See proof of Theorem 4.1.2 for a precise computation.

4

EXTINCTION OF DOMINATED STRATEGIES

In this chapter, we start the study of the long-run behavior of stochastic imitation dynamics by first focusing on the elimination of dominated strategies. We provide a general theorem for the extinction of strategies along trajectories of (SID), and rates at which they occur.

4.1 EXTINCTION OF STRATEGIES IN STOCHASTIC IMITATION DYNAMICS

rational behavior

A rational behavior in practical applications of game theory would be that strategies that are dominated, i.e., always worse than another, should be less and less used along the game, and so should become extinct in the long-run. This is indeed what we observe in the deterministic model based on the imitiation dynamics (ID), however what can be said about the robustness of such results when being strongly influenced by random perturbations ?

extinction of strategies

To be more precise, we say that a pure strategy $\alpha \in \mathcal{A}$ *becomes extinct* along a trajectory $X(t) \in \mathcal{X}$ of (SID) if $X_\alpha(t) \longrightarrow 0$ (a.s.) when t goes to infinity. Similarly and following Samuelson and Zhang (1992), a mixed strategy $p \in \mathcal{X}$ is said to become extinct along $X(t)$ if $\min\{X_\alpha(t) : \alpha \in \operatorname{supp} p\} \longrightarrow 0$ (a.s.), or in other words if $X(t)$ converges almost surely to the union of the faces of $\mathcal{X}$ that do not contain p.

We first present an example, which showcase that even in a very simple game the noise can be taken in an adversarial way so that non-dominated strategies become extinct with positive probability.

dominated strategies can survive if noise big enough

Example 4.1.1 (A bad noise scenario)**.** Let us consider a game with only two pure actions, i.e., $\mathcal{A} = \{\alpha, \beta\}$, and with payoffs $1 = v_\beta(x) > v_\alpha(x) = 0$ for all $x \in \mathcal{X}$, so that β dominates α. In a deterministic setting, we would then observe that strategy α becomes extinct along the trajectories of (ID), but what about the stochastic regime ?

For simplicity, assume that the strategic profile of the population evolves according to the stochastic imitation dynamics (SID) with an underlying imitation protocol such that $r_{\alpha\beta} - r_{\beta\alpha} = 1$[1], and with independent noise of constant variance across the space. Accordingly, let $\sigma^2 = \sigma^2_{\alpha\beta} + \sigma^2_{\alpha\beta}$ be a measure of the intensity of this noise.

Even in this simple setting, if the noise amplitude is taken big enough and if the dynamics starts near α, then we will observe trajectories of (SID) with respect to which β becomes extinct even if it strictly dominates α. We illustrate this behavior in Figure 4.1.1, where he have sampled trajectories of (SID) starting from $X_\beta(0) = 0.1$, five of which are in a *low noise* regime (cf. fig. 4.1.1a), and the other five in a *high noise* regime (cf. fig. 4.1.1b).

♦

1 For instance, this is the case when taking the imitation of success protocol $r_{\alpha\beta} = K + v_\beta$.

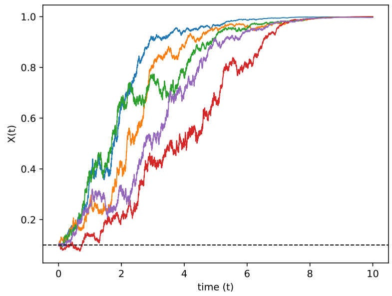

(a) Low noise regime ($\sigma = 0.5$)

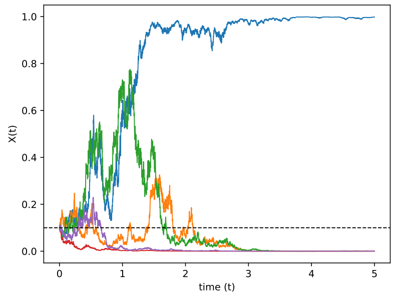

(b) High noise regime ($\sigma = 2$)

Figure 4.1.1: Trajectories of $X(t) \equiv X_\beta(t)$ evolving with respect to the dynamics (SID) with initial condition $X(0) = 0.1$ for two different levels of noise. Even if the pure strategy β strictly dominates α, we observe in Figure 4.1.1b that it becomes extinct along some trajectories of (SID) when the noise is taken big enough.

Motivated by Example 4.1.1, we expect that we need some conditions on the noise amplitude in order to observe rational behavior of the dynamics, i.e., dominated strategies that indeed become extinct along the trajectories of (SID). Such conditions are what we provide in Theorem 4.1.2 down below.

perturbed payoff vector $\tilde{v}$

For reasons that would become clear right afterward, an important role will be played by the perturbed vector field $\tilde{v}\colon \mathcal{X} \to \mathbb{R}^{\mathcal{A}}$ whose α-th component is given by

$$\tilde{v}_\alpha(x) = \sum\nolimits_\beta x_\beta[r_{\beta\alpha} - r_{\alpha\beta}] - \frac{1}{2}\sigma_\alpha^T(x)\widehat{C}(\alpha,\alpha)\sigma_\alpha(x), \qquad (\tilde{V})$$

with $\sigma_\alpha(x)$ and $\widehat{C}(\alpha,\alpha)$ defined as in Remark 3.3.3.

extinction of strategies in (SID)

Theorem 4.1.2. *Let $X(t)$ be an interior solution orbit of* (SID) *and let $p \in \mathcal{X}$ be a mixed strategy. Then p becomes extinct along $X(t)$ whenever there exists $p' \in \mathcal{X}$ such that*

$$\langle \tilde{v}(x), p' - p\rangle > 0 \quad \textit{for all } x \in \mathcal{X}, \qquad (4.1.1)$$

where $\tilde{v}$ is given by ($\tilde{V}$).

Remark 4.1.3. Theorem 4.1.2 yields an explicit threshold for the noise and how badly things can happen when random perturbations are too high compared to the payoffs. Indeed, the extinction result do *not* require domination of p in the original game $\Gamma(\mathcal{N},\mathcal{A},v)$, meaning that even a non dominated strategy p can become extinct with positive probability if the noise is chosen bad enough so that condition (4.1.1) holds. This also shows that such behaviors could *always* happen whatever the imitation protocols we choose, as the vector field $\tilde{v}$ will never be equal to the full payoff vector v due to the additive noise term.

pertubed game $\tilde{\Gamma}(\mathcal{N},\mathcal{A},\tilde{v})$

Remark 4.1.4. Noticing the similarity with mixed strategies domination (3.1.1), condition (4.1.1) can also be interpreted as saying that p is dominated by p' in the *perturbed game* $\tilde{\Gamma}(\mathcal{N},\mathcal{A},\tilde{v})$. Imhof (2005) and Mertikopoulos and Viossat (2016) introduced similar modified games in order to study the long-run behavior of stochastic replicator dynamics with aggregate shocks and payoff shocks respectively, putting our result in the continuation of such works. In particular, we obtain the same noise conditions as Mertikopoulos and Viossat (2016) when considering the dynamics (RDPS) generated by an imitation of success protocol and independent target noise.

Proof of Theorem 4.1.2. Let $X(t)$ be an interior solution orbit of (SID) (which exists whenever $X(0) \in \mathcal{X}^{\circ}$ thanks to Proposition 3.3.4), and let $p, p' \in \mathcal{X}$ verifying condition (4.1.1). We define the Kullback-Leibler divergence between p and $x \in \mathcal{X}$ as $\mathrm{KL}(p, x) = \sum_{\alpha} p_{\alpha} \log \frac{p_{\alpha}}{x_{\alpha}}$. By classical arguments (see Weibull, 1995), we notice that $\mathrm{KL}(p, x) < \infty$ if and only if $x_{\alpha} > 0$ for all $\alpha \in \mathrm{supp}(p)$, and so p become extinct along $X(t)$ whenever $\lim_{t\to\infty} \mathrm{KL}(p, X(t)) \to \infty$ (a.s.). Accordingly, and following the analysis of Cabrales (2000), we introduce the cross-entropy function

Kullback-Leibler divergence KL

cross-entropy function V

$$V(x) \equiv \mathrm{KL}(p, x) - \mathrm{KL}(p', x) = \sum_{\alpha} (p_{\alpha} \log p_{\alpha} - p'_{\alpha} \log p'_{\alpha}) + \sum_{\alpha} (p'_{\alpha} - p_{\alpha}) \log x_{\alpha} \tag{4.1.2}$$

so that $\mathrm{KL}(p, X(t)) \to \infty$ (a.s.) if and only if $V(X(t)) \to \infty$ (a.s.) (notice that $\mathrm{KL}(p', x) \geq 0$ by usual properties of the Kullback-Leibler divergence). To that end, let $Y_{\alpha} = \log X_{\alpha}$. Following Remark 3.3.5 we know that X_{α} is a semimartingale, for which Itô's formula (Theorem 2.2.5) with $f(x) = \log x$ then readily yields

dual process $Y_{\alpha} = \log X_{\alpha}$

$$dY_{\alpha} = \frac{dX_{\alpha}}{X_{\alpha}} - \frac{1}{2X_{\alpha}^2} dX_{\alpha} \cdot dX_{\alpha}. \tag{4.1.3}$$

From the compact form (SID′) of the stochastic imitation dynamics and Proposition 2.2.2, we can compute

$$dX_{\alpha} \cdot dX_{\alpha} = X_{\alpha}^2 \sigma_{\alpha}^T dW_{\alpha} \cdot dW_{\alpha} \sigma_{\alpha} = X_{\alpha}^2 \sigma_{\alpha}^T dW_{\alpha} \cdot dW_{\alpha}^T \sigma_{\alpha} = X_{\alpha}^2 \sigma_{\alpha}^T \widehat{C}(\alpha, \alpha) \sigma_{\alpha} dt, \tag{4.1.4}$$

where we have dropped the dependence of the coefficients on X for notational convenience. Putting the expression back into (4.1.3), it then holds that

$$dY_{\alpha} = \left(\sum_{\beta} \left[r_{\beta\alpha} - r_{\alpha\beta} \right] - \frac{1}{2} \sigma_{\alpha}^T \widehat{C}(\alpha, \alpha) \sigma_{\alpha} \right) dt + \sigma_{\alpha}^T dW_{\alpha} = \tilde{v}_{\alpha}(X) dt + \sigma_{\alpha}^T(X) dW_{\alpha}, \tag{4.1.5}$$

with $\tilde{v}$ given by ($\tilde{V}$). As the first sum in the cross-entropy function's expression do not depend on time, we therefore obtain

$$dV = \sum_{\alpha} (p' - p) dY_{\alpha} = \sum_{\alpha} (p'_{\alpha} - p_{\alpha}) \tilde{v}_{\alpha}(X) dt + \sum_{\alpha} (p'_{\alpha} - p_{\alpha}) \sigma_{\alpha}^T(X) dW_{\alpha} \tag{4.1.6}$$

$$= \langle \tilde{v}(X), p' - p \rangle dt + \sigma_{p'p}^T(X) dW, \tag{4.1.7}$$

where $\sigma_{p'p}^T = \begin{bmatrix} (p'_1 - p_1)\sigma_1^T & \ldots & (p'_A - p_A)\sigma_A^T \end{bmatrix}$ and $W^T = \begin{bmatrix} W_1^T & \ldots & W_A^T \end{bmatrix}$. Notice that W is still a vector of Brownian motions defined on the same underlying filtration. By assumption, let $m > 0$ such that $\langle \tilde{v}(x), p' - p \rangle \geq m$ for all $x \in \mathcal{X}$ (such a constant exists because $\tilde{v}$ is continuous on the compact $\mathcal{X}$). The stochastic differential equation (4.1.7) then yields

$$V(X(t)) \geq V(X(0)) + mt + \int_0^t \sigma_{p'p}^T(X(s))\, dW(s). \tag{4.1.8}$$

Since the coefficients $\sigma_{p'p}$ are continuous on the compact set $\mathcal{X}$, they are in particular bounded. Furthermore, X is adapted to the filtration generated by W by Theorem 2.2.7, and so is $\sigma_{p'p}^T(X)$ by continuity. Therefore, Corollary 2.3.4 states that $mt + \int_0^t \sigma_{p'p}^T(X(s))\, dW(s) \sim mt$ (a.s.), from which we deduce immediately the limit $V(X(t)) \to \infty$, hence proving the theorem. □

4.2 RATES OF EXTINCTION

Theorem 4.1.2 only provides a quantitative answer to the question "what conditions are needed for the extinction of mixed strategies ?". However, in many practical applications we could be interested into more qualitative considerations, such as the rates at which extinctions occur.

In the following propositions, we investigate this matter by first proposing an asymptotic rate of almost sure extinction (Proposition 4.2.1), then by giving two non-asymptotic results on the probability to observe deviation from the mean and on the average time needed for a strategy to fall under a fixed threshold (Proposition 4.2.6).

asymptotic rate of extinction

Proposition 4.2.1. *Under the same notations as Theorem 4.1.2, if (4.1.1) holds then for all $\varepsilon > 0$ and for all t big enough, the extinction occurs with rate*

$$\min\{X_\alpha(t) : \alpha \in \operatorname{supp}(p)\} \leq K \exp\Big\{-mt + 2(1+\varepsilon)\hat{\sigma}_{p'p}\sqrt{t \log\log t}\Big\} \quad (a.s.) \tag{4.2.1}$$

where

i) K is a constant depending only on the initial conditions of (SID);

ii) $m = \inf_{x\in\mathcal{X}} \langle \tilde{v}(x), p' - p\rangle > 0$;

iii) $\hat{\sigma}^2_{p'p} = \sup_{x\in\mathcal{X}} \sigma^T_{p'p}(x)\widetilde{C}\sigma_{p'p}(x) < \infty$ with $\sigma^T_{p'p} = \begin{bmatrix}(p'_1 - p_1)\sigma^T_1 & \dots & (p'_A - p_A)\sigma^T_A\end{bmatrix}$ and $\widetilde{C} = [\,\widehat{C}(\alpha,\beta)\,]_{\alpha,\beta\in\mathcal{A}}$.

Corollary 4.2.2. *If p and p' are both pure strategies, i.e., if $p = \alpha$ and $p' = \beta$ where $\alpha, \beta \in \mathcal{A}$, then the extinction rate of Proposition 4.2.1 becomes*

$$X_\alpha(t) \leq K \exp\Big\{-mt + 2(1+\varepsilon)\hat{\sigma}_{\alpha\beta}\sqrt{t\log\log t}\Big\} \quad (a.s.) \tag{4.2.2}$$

with

i) $m = \inf_{x\in\mathcal{X}} \tilde{v}_\beta(x) - \tilde{v}_\alpha(x) > 0$;

ii) $\hat{\sigma}^2_{\alpha\beta} = \sup_{x\in\mathcal{X}} \sigma^T_{\alpha\beta}(x)\widetilde{C}(\alpha,\beta)\sigma_{\alpha\beta}(x) < \infty$ with $\widetilde{C}(\alpha,\beta) = \begin{bmatrix}\widehat{C}(\alpha,\alpha) & \widehat{C}(\alpha,\beta)\\ \widehat{C}(\beta,\alpha) & \widehat{C}(\beta,\beta)\end{bmatrix}$ and $\sigma^T_{\alpha\beta} = \begin{bmatrix}\sigma_\alpha & -\sigma_\beta\end{bmatrix}$.

Remark 4.2.3. The asymptotic rate of elimination for pure strategies in Corollary 4.2.2 is similar to those obtained by Imhof (2005) and Bravo and Mertikopoulos (2017) for replicator dynamics with aggregate shocks and stochastic exponential learning respectively, even if different thresholds of the noise were needed to observe extinction of dominated strategies. This may indicate that the random perturbations in game dynamics have only an impact on whether a dominated strategy becomes extinct, but not on the speed at which such strategies tend to zero.

Heuristically, this phenomenon is due to the fact that the drift of the underlying Brownian motion always pushes the trajectories toward the boundary of the simplex in a rate of order t as soon as the noise condition is verified, much faster than the diffusion part which can only perturb the system up to an order $\sqrt{t\log\log t}$.

Proof of Proposition 4.2.1. Following in the steps of the proof of Theorem 4.1.2, inequality (4.1.8) yields

$$\mathrm{KL}(p, X(t)) - \mathrm{KL}(p', X(t)) = V(X(t)) \geq V(X(0)) + mt + \xi_{p'p}(t) \tag{4.2.3}$$

where $\xi_{p'p}(t) = \int_0^t \sigma_{p'p}^T(X(s))\, dW(s)$. Since $\mathrm{KL}(p', X(t)) \geq 0$ for all t, the previous expression can be rearranged to show that

$$\mathrm{KL}(p, X(t)) \geq V(X(0)) + mt + \xi_{p'p}(t). \tag{4.2.4}$$

Now, let $X^*(t) = \min\{X_\alpha(t) : \alpha \in \mathrm{supp}(p)\}$, and notice that

$$\mathrm{KL}(p, X(t)) = \sum_{\alpha \in \mathrm{supp}(p)} p_\alpha \log \frac{p_\alpha}{X_\alpha(t)} \leq \sum_{\alpha \in \mathrm{supp}(p)} p_\alpha \log \frac{p_\alpha}{X^*(t)} \tag{4.2.5}$$

$$= \sum_{\alpha \in \mathrm{supp}(p)} p_\alpha \log p_\alpha - \log(X^*(t)), \tag{4.2.6}$$

where we have used that $\sum_{\alpha \in \mathrm{supp}(p)} p_\alpha = 1$. Combining this inequality with (4.2.4), we then readily get

$$X^*(t) \leq \exp\left\{\sum\nolimits_\alpha p_\alpha \log p_\alpha - \mathrm{KL}(p, X(t))\right\} \leq K \exp\{-mt - \xi_{p'p}(t)\} \tag{4.2.7}$$

where $K = \exp\left\{\sum_\alpha p_\alpha \log p_\alpha - V(X(0))\right\} > 0$ only depends on the initial condition and on the dominated strategy. As in the proof of Theorem 4.1.2, $\sigma_{p'p}(X)$ is bounded and adapted to the filtration generated by W, so in particular it is in $L_{ad}(W)$ where W is a martingale. Therefore, in virtue of Theorem 2.2.3, $\xi_{p'p}$ is itself a martingale, and so by the time-change theorem (Theorem 2.2.9) there exists a standard Brownian motion $\widetilde{W}$ (defined on a possibly enlarged probability space) such that $\xi_{p'p}(t) = \widetilde{W}\left(\left[\xi_{p'p}\right]_t\right)$. Using Proposition 2.2.2, the quadratic variation $\left[\xi_{p'p}\right]_t$ of ξ_t can then be computed explicitly as

$$d\left[\xi_{p'p}\right]_t = d\xi_{p'p} \cdot d\xi_{p'p} = \sigma_{p'p}^T dW \sigma_{p'p}^T dW = \sigma_{p'p}^T \widetilde{C} \sigma_{p'p} dt \tag{4.2.8}$$

where $\widetilde{C} = Cor(W, W) = [\,\widehat{C}(\alpha, \beta)\,]_{\alpha,\beta \in \mathcal{A}}$. Letting $\rho(t) = \left[\xi_{p'p}\right]_t$, we then have that $\rho(t) \leq \hat{\sigma}_{p'p}^2 t$ where $\hat{\sigma}_{p'p}^2 = \sup_{x \in \mathcal{X}} \sigma_{p'p}^T(x) \widetilde{C} \sigma_{p'p}(x) < \infty$ by compactness, and $\hat{\sigma}_{p'p}^2 > 0$ because $\widetilde{C}$ is symmetric definite (it is a correlation matrix). If $\rho(t) \to \infty$, then the law of iterated logarithm yields that for all $\varepsilon > 0$,

$$\left|\xi_{p'p}(t)\right| \leq (1+\varepsilon)\sqrt{2\rho(t) \log\log \rho(t)} \leq (1+\varepsilon)\sqrt{2\hat{\sigma}_{p'p}^2 t \log\log 2\hat{\sigma}_{p'p}^2 t} \tag{4.2.9}$$

$$\leq 2(1+\varepsilon)\hat{\sigma}_{p'p}\sqrt{t \log\log t} \quad (a.s.) \tag{4.2.10}$$

for t large enough. On the other hand, when $\rho(t) \to \rho(\infty) < \infty$, $\xi_{p'p}$ is almost surely bounded, and so

$$\left|\xi_{p'p}(t)\right| \leq 2(1+\varepsilon)\hat{\sigma}_{p'p}\sqrt{t \log\log t} \quad (a.s.) \tag{4.2.11}$$

for t big enough and for all $\varepsilon > 0$. The rate of extinction is then obtained by combining (4.2.10) and (4.2.11), and putting the bound back into (4.2.7). □

For the non-asymptotic rates of extinction we need a more precise control on the noise coefficients, accordingly that they do not depend "too much" on the trajectory of the dynamics. This heuristic condition is made precise in the following assumption, that we will also use in Chapter 5 to derive similar rates for the convergence toward Nash equilibria.

constant individual noise

Assumption 4.2.4 (Constant individual noise)**.** The noise coefficients of (SID′) can be decomposed as $\sigma_\alpha^T(X)dW_\alpha = \eta_\alpha^T dW_\alpha + dF$ where η_α do not depend on X and F is a continuous martingale that may depend on X.

Remark 4.2.5. The decomposition of Assumption 4.2.4 can be interpreted as a linear separability between an *individual noise* η_α and a *common noise* F, where the individual noises have constant variance across the space. For instance, stochastic imitation dynamics under target noise or source noise (see Example 3.3.1) can be decomposed in this form whenever their diffusion coefficient have constant variance. In particular, this is also the case for the replicator dynamics with payoff shocks studied by Mertikopoulos and Viossat (2016).

Proposition 4.2.6. *Under the same notations as Theorem 4.1.2 and Proposition 4.2.1, assume that (4.1.1) holds and that Assumption 4.2.4 is verified. Then, we can derive the following rates of extinction for* $X^*(t) = \min\{X_\alpha(t) : \alpha \in \operatorname{supp}(p)\}$*:*

concentration inequality

a) *(Concentration inequality) for all* $\delta > 0$ *and* $t \geq 0$*,*

$$\mathbb{P}(X^*(t) > \delta) \leq \frac{1}{2}\operatorname{erfc}\left[\frac{1}{\sqrt{2}\hat{\eta}_{p'p}}\left(m\sqrt{t} - \frac{K_* - \log\delta}{\sqrt{t}}\right)\right], \tag{4.2.12}$$

where $K_* = \log(K)$, $\hat{\eta}_{p'p}^2 = \eta_{p'p}^T\widetilde{C}\eta_{p'p}$ *with* $\eta_{p'p}^T = \left[(p_1' - p_1)\eta_1^T \quad \dots \quad (p_A' - p_A)\eta_A^T\right]$, *and* $\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}}\int_z^\infty e^{-s^2}ds$ *is the complementary error function.*

hitting time

b) *(Hitting time) for all* $\delta > 0$*,*

$$\mathbb{E}[\tau_\delta] \leq \frac{\left[K_* - \log\delta\right]_+}{m} \tag{4.2.13}$$

where $\tau_\delta = \inf\{t \geq 0 : X^*(t) \leq \delta\}$.

Remark 4.2.7. Even if the noise terms σ_α do not seem to appear in (4.2.13), it is important to remember that m is defined as $\inf_{x\in\mathcal{X}}\langle\tilde{v}(x), p' - p\rangle$, where $\tilde{v}_\alpha$ itself depends on σ_α from its expression ($\tilde{V}$). Accordingly, to compute the rates (4.2.12) and (4.2.13) in practice, we would need to estimate at first the amplitude of the noise across the different pure strategies.

Proof of Proposition 4.2.6. Coming back to the stochastic differential equation (4.1.7) verified by the cross-entropy function V, we notice that if the noise coefficients decompose as $\sigma_\alpha^T(X)dW_\alpha = \eta_\alpha^T dW_\alpha + dF$, then

$$\begin{aligned} dV &= \langle\tilde{v}(X), p' - p\rangle dt + \sum\nolimits_\alpha(p_\alpha' - p_\alpha)\eta_\alpha^T dW_\alpha + \left[\sum\nolimits_\alpha(p_\alpha' - p_\alpha)\right]dF && (4.2.14)\\ &= \langle\tilde{v}(X), p' - p\rangle dt + \eta_{p'p}^T dW, && (4.2.15)\end{aligned}$$

where the last equality comes from the fact that $\sum_\alpha(p_\alpha' - p_\alpha) = 0$ and that F do not depend on the action α. Continuing as in the proofs of Theorem 4.1.2 and Proposition 4.2.1, we readily obtain that

$$X^*(t) \leq K\exp\{-mt - \xi_{p'p}(t)\} \tag{4.2.16}$$

with $\xi_{p'p}$ now given by $\xi_{p'p}(t) = \int_0^t \eta_{p'p}^T dW(s)$. From the usual argument based on the time-change theorem, we can then write $\xi_{p'p}(t) = \widetilde{W}(\rho(t))$ where

$$\rho(t) = \left[\xi_{p'p}\right]_t = \int_0^t \eta_{p'p}^T \widetilde{C} \eta_{p'p}\, ds = \hat{\eta}_{p'p}^2 t. \tag{4.2.17}$$

Equation (4.2.16) and the representation of $\xi_{p'p}$ as a time-changed Brownian motion constitute the starting point for the proofs of both rates of extinction.

a) From (4.2.16), we deduce that for all $\delta > 0$

$$X^*(t) > \delta \Longrightarrow K \exp\{-mt - \xi_{p'p}(t)\} > \delta, \tag{4.2.18}$$

which can be put into probabilities to obtain

$$\mathbb{P}(X^*(t) > \delta) \leq \mathbb{P}\big(K \exp\{-mt - \xi_{p'p}(t)\} > \delta\big) \tag{4.2.19}$$

$$= \mathbb{P}\Big(\widetilde{W}(\rho(t)) < -mt + K_* - \log\delta\Big) \tag{4.2.20}$$

where $K_* = \log(K)$. As $\rho(t)$ is a deterministic quantity (reason why we need the independence of $\eta_{p'p}$ from X), we have $\widetilde{W}(\rho(t)) \sim \mathcal{N}(0, \rho(t))$ by definition of the Brownian motion. Letting $l = -mt + K_* - \log\delta$ and writing ρ instead of $\rho(t)$ for readability, we can then compute explicitly the right-hand side of (4.2.20) :

$$\mathbb{P}\Big(\widetilde{W}(\rho) < l\Big) = \mathbb{P}\Big(\widetilde{W}(\rho) > -l\Big) \tag{4.2.21}$$

$$= \int_{-l}^{\infty} \frac{1}{\sqrt{2\pi\rho}} \exp\{-\frac{y^2}{2\rho}\}\, dy \tag{4.2.22}$$

$$= \frac{1}{2}\operatorname{erfc}\left[-\frac{l}{\sqrt{2\rho}}\right] \tag{4.2.23}$$

$$= \frac{1}{2}\operatorname{erfc}\left[\frac{1}{\sqrt{2}\hat{\eta}_{p'p}}\left(m\sqrt{t} - \frac{K_* - \log\delta}{\sqrt{t}}\right)\right] \tag{4.2.24}$$

thus proving (4.2.12).

b) Taking the complementary of (4.2.18), we notice that

$$\tau_\delta = \inf\{t \geq 0 : X^*(t) \leq \delta\} \leq \inf\Big\{t \geq 0 : K \exp\{-mt - \xi_{p'p}(t)\} \leq \delta\Big\} \tag{4.2.25}$$

$$= \inf\Big\{t \geq 0 : mt + \widetilde{W}(\rho(t)) \geq a\Big\} \tag{4.2.26}$$

$$=: \tau_a \tag{4.2.27}$$

where $a = \left[K_* - \log\delta\right]_+$ (if $\delta \geq e^{K_*} = K$ then both τ_δ and τ_a are equal to 0). Now, as $\rho(t) = \hat{\eta}_{p'p}^2 t$ is a deterministic time change, we can write

$$\tau_a = \frac{1}{\hat{\eta}_{p'p}^2} \inf\left\{s \geq 0 : \frac{m}{\hat{\eta}_{p'p}^2} s + \widetilde{W}(s) \geq a\right\} = \frac{\tau_a^h}{\hat{\eta}_{p'p}^2} \tag{4.2.28}$$

where τ_a^h is the hitting time at level a of a Brownian motion with positive drift $\mu = m/\hat{\eta}_{p'p}^2$. A classical application of Dynkin's formula then yields $\mathbb{E}\left[\tau_a^h\right] = a/\mu$, from which we deduce that

$$\mathbb{E}[\tau_\delta] \leq \frac{1}{\hat{\eta}_{p'p}^2} \mathbb{E}\left[\tau_a^h\right] \leq \frac{a}{m}, \tag{4.2.29}$$

hence proving the second bound (4.2.13).

□

missing assumption in Bravo and Mertikopoulos (2017)

Remark 4.2.8. Bravo and Mertikopoulos (2017) have proposed a similar result than (4.2.12) for their stochastic regularized learning dynamics but without stating a condition such as Assumption 4.2.4 for the noise coefficients. Their proof of this result is however flawed whenever the noise depends on $X(t)$, as the main idea behind both of our results is that $\widetilde{W}(\rho(t)) \sim \mathcal{N}(0, \rho(t))$, which is obviously false when $\rho(t)$ is random and dependent of $X(t)$. Even a conditioning argument, i.e., looking at the distribution of $\widetilde{W}(\rho(t))$ given the filtration $\mathcal{F}_t^X$ generated by $\{X(s) : s \leq t\}$, could not solve this issue as the Brownian motion $\widetilde{W}$ is constructed by the time-change theorem on a martingale which is dependent on $X(t)$, and so $\widetilde{W}$ is itself dependent on $X(t)$.

5

STABILITY OF EQUILIBRIA

In this chapter, we focus ourselves on the study of the equilibrium stability under stochastic imitations dynamics. Particularly, we investigate if the noise conditions established in Chapter 4 are also sufficient to observe stochastic stability around Nash equilibria of the game; and if so, at what rates do such convergences occur.

5.1 STOCHASTIC STABILITY

As mentioned in Chapter 3, the asymptotic stability of Nash equilibria for deterministic imitation dynamics is well understood in the game theoretic literature. The "folk theorem" of evolutionary game theory indeed states that *a*) Lyapunov stable states are Nash; *b*) limits of interior trajectories are Nash; and *c*) strict Nash equilibria are asymptotically stable under (ID).

folk theorem for (ID)

However, in a stochastic regime we would expect that such results stay true only under some noise conditions as it was the case for the extinction of dominated strategies in Chapter 4. This is indeed what was also shown by Hofbauer and Imhof (2009) and Mertikopoulos and Viossat (2016) for different kinds of stochastic replicator dynamics, where the equilibrium states need to be Nash in a modified game. On the other hand, the stochastic replicator dynamics generated by perturbed exponential learning do not need such conditions, and stability of its Nash equilibria always hold true (see Mertikopoulos and Moustakas, 2010 and Bravo and Mertikopoulos, 2017).

need noise conditions

Furthermore, as we are now dealing with stochastically perturbed dynamics systems where oscillations around rest points could still be observed, the notions of Lyapunov and asymptotically stabilities need to be adapted in this respect. Following Khasminskii (2012), we then define :

Definition 5.1.1. Let $x^* \in \mathcal{X}$. We say that

1. x^* is *stochastically (Lyapunov) stable* under (SID) if, for every $\varepsilon > 0$ and for every neighborhood $\mathcal{U}_0$ of x^* in $\mathcal{X}$, there exists an even smaller neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of x^* such that

   stochastically stable

$$\mathbb{P}(X(t) \in \mathcal{U}_0 \text{ for all } t \geq 0) \geq 1 - \varepsilon \tag{5.1.1}$$

   whenever $X(0) \in \mathcal{U}$.

2. x^* is *stochastically asymptotically stable* under (SID) if it is stochastically stable and attracting : for every $\varepsilon > 0$ and for every neighborhood $\mathcal{U}_0$ of x^* in $\mathcal{X}$, there exists an even smaller neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of x^* such that

   stochastically asymptotically stable

$$\mathbb{P}\Big(X(t) \in \mathcal{U}_0 \text{ for all } t \geq 0 \text{ and } \lim_{t\to\infty} X(t) = x^*\Big) \geq 1 - \varepsilon \tag{5.1.2}$$

whenever $X(0) \in \mathcal{U}$.

Unlike the deterministic case, stable points of the dynamics will not be described as Nash equilibria of the game, but instead by a modified condition based on the random perturbations of the payoff vector. Accordingly, we say that $x^* \in \mathcal{X}$ verifies the equilibrium condition (EQ^σ) with respect to vector field $\tilde{v}$ if

perturbed equilibrium (EQ^σ)

$$\tilde{v}_\alpha(x^*) \geq \tilde{v}_\beta(x^*) \quad \text{for all } \alpha \in \operatorname{supp}(x^*) \text{ and for all } \beta \in \mathcal{A}, \qquad (\mathrm{EQ}^\sigma)$$

strict perturbed equilibrium (EQ^σ)

where $\tilde{v}$ is given by ($\tilde{V}$). In particular, it is said to verify *strictly* the equilibrium condition (EQ^σ) if the inequality is strict for all $\beta \notin \operatorname{supp}(x^*)$.

In the following theorem, we first characterize the points of the space which are either limits of interior trajectories or stochastically stable states with respect to (SID). Its proof will be delayed to Section 5.3 in order to streamline the presentation.

limits of interior trajectories and stochastically stable states are perturbed equilibria

Theorem 5.1.2. *Let $x^* \in \mathcal{X}$ and let $X(t)$ be an interior solution orbit of* (SID)*. Assume that we have either*

1. $\mathbb{P}(\lim_{t\to\infty} X(t) = x^*) > 0$,
2. *or x^* be a stochastically Lyapunov stable state.*

Then x^* verifies the equilibrium condition (EQ^σ).

Nash equilibrium in perturbed game $\tilde{\Gamma}$

Remark 5.1.3. Similarly to Remark 4.1.4 on the extinction of dominated strategy, the equilibrium condition (EQ^σ) can be interpreted as saying that the limits of interior trajectories and the stochastically stable states are Nash equilibria in the *modified game* $\tilde{\Gamma}(\mathcal{N}, \mathcal{A}, \tilde{v})$. Accordingly, this result can again be put into comparison with those of Hofbauer and Imhof (2009) and Mertikopoulos and Viossat (2016) on stochastic replicator dynamics with random shocks.

Remark 5.1.4. According to Theorem 5.1.2, limits of interior trajectories and stochastically stable states are not necessarily Nash equilibria in the original game $\Gamma(\mathcal{N}, \mathcal{A}, v)$ unless the noise is mild enough. This is to be contrasted with the deterministic folk theorem which stated that such points were always Nash equilibria, but also to the stochastic study of perturbed exponential learning from Mertikopoulos and Moustakas (2010) and Bravo and Mertikopoulos (2017) who obtain the characterization as Nash equilibria regardless of the level of noise.

Now that the stable points are characterized, we may wonder if the last result of the folk theorem also hold in a stochastic regime, i.e., if the strict Nash equilibria are asymptotically stochastically stable under (SID). Surprisingly, this is indeed the case as we show in the following theorem and its associated corollary.

perturbed equilibria are stochastically asymptotically stable

Theorem 5.1.5. *Let $x^* \in \mathcal{X}$. If x^* verifies strictly the equilibrium condition* (EQ^σ)*, then x^* is stochastically asymptotically stable under* (SID).

Remark 5.1.6. In the continuation of Remark 5.1.3, we can interpret the result of Theorem 5.1.5 as saying that strict Nash equilibria of the modified game $\tilde{\Gamma}(\mathcal{N}, \mathcal{A}, \tilde{v})$ are stochastically asymptotically stable under (SID).

strict Nash equilibria are stochastically asymptotically stable

Corollary 5.1.7. *If $x^* \in \mathcal{X}$ is a strict Nash equilibrium of the game $\Gamma(\mathcal{N}, \mathcal{A}, v)$, then it is stochastically asymptotically stable under* (SID).

Proof. As x^* is a strict Nash equilibrium of Γ, it is in particular a pure strategy. Accordingly, let $\alpha^* \in \mathcal{A}$ such that $\operatorname{supp}(x^*) = \{\alpha^*\}$. From the expressions in Remark 3.3.3, we can then compute

$$\sigma_{\alpha^*}^T(x^*)\widehat{C}(\alpha^*, \alpha^*)\sigma_{\alpha^*}(x^*) = \sigma_{\alpha^*\alpha^*}^2(x^*) + \sigma_{\alpha^*\alpha^*}^2(x^*) - 2\sigma_{\alpha^*\alpha^*}^2(x^*) = 0. \tag{5.1.3}$$

Combined with the fact that x^* is a Nash equilibrium for the payoff vector v and the monotone condition (3.2.4), this implies that for all $\alpha \neq \alpha^*$,

$$\tilde{v}_{\alpha^*}(x^*) - \tilde{v}_\alpha(x^*) > \frac{1}{2}\sigma_\alpha^T(x^*)\widehat{C}(\alpha, \alpha)\sigma_\alpha(x^*) \geq 0; \tag{5.1.4}$$

i.e., x^* verifies strictly the equilibrium (EQ^σ). Thus proving the stability through Theorem 5.1.5. □

Remark 5.1.8. In fact, we can show that the strict equilibrium condition (EQ^σ) of Theorem 5.1.5 is true if and only if x^* is pure (due to its interpretation as a strict Nash equilibrium of the modified game $\tilde{\Gamma}$), and

$$r_{\alpha^*\alpha}(v, x^*) - r_{\alpha\alpha^*}(v, x^*) < \frac{1}{2}\left[\sigma_{\alpha\alpha*}^2(x^*) + \sigma_{\alpha^*\alpha}^2(x^*) - 2\sigma_{\alpha\alpha^*}\sigma_{\alpha^*\alpha}C(\alpha^*, \alpha; \alpha, \alpha^*)\right] \tag{5.1.5}$$

for all $\alpha \neq \alpha^*$, where α^* is the pure action associated to x^*. For instance, if we consider the imitation of success protocol $r_{\alpha\beta} = K + v_\alpha$ with independent target noise, this condition simplifies to

$$v_\alpha(x^*) - v_{\alpha^*}(x^*) < \frac{1}{2}\left[\sigma_{\alpha^*}^2 + \sigma_\alpha^2\right] \tag{5.1.6}$$

which is exactly the same as the one provided in Remark 3.6 of Mertikopoulos and Viossat (2016).

Proceeding as in this remark, we will provide some intuition to the condition (5.1.5). Consider that $\mathcal{A} = \{\alpha, \alpha^*\}$ and that the diffusion coefficient have constant variance. Then, we can show that

$$dX_\alpha = X_\alpha(1 - X_\alpha)[(r_{\alpha^*\alpha} - r_{\alpha\alpha^*})dt + \sigma dW] \tag{5.1.7}$$

where $\sigma^2 = \sigma_{\alpha\alpha^*}^2 + \sigma_{\alpha^*\alpha}^2 - 2\sigma_{\alpha\alpha^*}\sigma_{\alpha^*\alpha}C(\alpha^*, \alpha; \alpha, \alpha^*)$ and dW is a Brownian motion obtained by the time-change theorem and a rescaling. A discrete counterpart of the process is therefore given by the random walk $X_\alpha^d(n)$ verifying

$$X_\alpha^d(n+1) - X_\alpha^d(n) = X_\alpha^d(n)\Big(1 - X_\alpha^d(n)\Big)\Big[(r_{\alpha^*\alpha} - r_{\alpha\alpha^*})\delta + \sigma\xi_n\sqrt{\delta}\Big] \tag{5.1.8}$$

where ξ_n is a Rademacher random variable and $\delta > 0$ approximates the infinitesimal time step dt. Whenever δ and $X_\alpha^d(n)$ are small enough, it can be shown after some computations that conditioned on the event $\{\xi_{n+1} = -\xi_n\}$, we have

$$X_\alpha^d(n+2) - X_\alpha^d(n) = 2\delta X_\alpha^d(n)\left[r_{\alpha^*\alpha} - r_{\alpha\alpha^*} - \frac{1}{2}\sigma^2\right] + o(\delta) + o\Big(X_\alpha^d(n)\Big). \tag{5.1.9}$$

Therefore, the discrete process $X_\alpha^d(n)$ is decreasing if and only if condition (5.1.5) holds. In other words, if a trajectory starts near the vertex associated to the pure strategy α^*, then after two small time steps the process will be closer to this vertex with probability $3/4$ (probability that $\xi_{n+1} = -\xi_n$), and so strategy α^* will be strictly attracting and stable for the stochastic dynamics.

Remark 5.1.9. Corollary 5.1.7 is a very surprising result on stochastic imitation dynamics. Indeed, it states that whatever the noise we choose, correlated or not, strict Nash equilibria would *always* be stochastically asymptotically stable. Such an observation was already made in Mertikopoulos and Viossat (2016) for the replicator dynamics with payoff shocks, but here we generalize this fact to any imitation protocol as long as it verifies the monotone condition (3.2.4). Furthermore, such a behavior is entirely due to how we have injected the noise into our deterministic dynamics. For instance, the stochastic variant of replicator dynamics with aggregate shock from biology do not verify this condition, i.e., strict Nash equilibria of the original game are not necessarily strict Nash equilibria in the modified game (cf. Theorem 4.11 of Hofbauer and Imhof, 2009).

5.2 RATES OF CONVERGENCE

As for the extinction of dominated strategies in Chapter 4, some applications may need more qualitative results than just stating the asymptotic stability of strict Nash equilibria. In particular, we would like to know at what rates do the convergence occurs, either asymptotically or non-asymptotically, to control more efficiently the game dynamics and its behavior.

Similarly to Proposition 4.2.1 and Proposition 4.2.6, we propose first an asymptotic rate of almost sure convergence, then results on the probability to observe a deviation from the mean and on the mean time needed for the trajectory to be arbitrarily close to the equilibrium. However, contrarily to dominated strategies, the convergence and stability of strict Nash equilibria only hold with probability greater than $1-\eta$, and as such we need to condition on such events to obtain finite rates.

As in the previous section, we prefer to only state and comment the results here in order to clarify the discussion, postponing the proofs to Sections 5.4 and 5.5.

asymptotic rate of convergence

Proposition 5.2.1. *Assume $\alpha^* \in \mathcal{A}$ is a pure Nash equilibrium of the game $\Gamma(\mathcal{N}, \mathcal{A}, v)$, and let $\mathcal{U}$ be a neighborhood of x^* so that $\inf_{\alpha \neq \alpha^*, x \in \mathcal{U}} \tilde{v}_{\alpha^*}(x) - \tilde{v}_\alpha(x) > 0$[1]. Then, for all $\eta > 0$, there exists a neighborhood $\mathcal{U}_\eta \subseteq \mathcal{U}$ of x^* such that for all $\varepsilon > 0$ and t big enough,*

$$X_{\alpha^*}(t) \geq 1 - K(A-1)\exp\left\{-mt + 2(1+\varepsilon)\sigma_*\sqrt{t \log\log t}\right\} \tag{5.2.1}$$

with probability at least $1-\eta$ whenever $X(t)$ is a solution orbit of (SID) with $X(0) \in \mathcal{U}_\eta$, where

i) K is a constant depending only on the initial conditions of (SID);

ii) $m = \inf_{\alpha \neq \alpha^*, x \in \mathcal{U}} \tilde{v}_{\alpha^*}(x) - \tilde{v}_\alpha(x) > 0$;

iii) $\sigma_*^2 = \sup_{\alpha \neq \alpha^*, x \in \mathcal{U}} \sigma_{*\alpha}^T(x)\overline{C}(\alpha^*, \alpha)\sigma_{*\alpha}(x) < \infty$ with $\overline{C}(\alpha^*, \alpha) = \begin{bmatrix} \widehat{C}(\alpha^*, \alpha^*) & \widehat{C}(\alpha^*, \alpha) \\ \widehat{C}(\alpha, \alpha^*) & \widehat{C}(\alpha, \alpha) \end{bmatrix}$ and $\sigma_{*\alpha}^T = \begin{bmatrix} \sigma_{\alpha^*} & -\sigma_\alpha \end{bmatrix}$.

1 Such a neighborhood always exists by the pure Nash definition and the continuity of $\tilde{v}$.

Remark 5.2.2. Another way to formulate bound (5.2.1) is by noticing that $\|X(t)-x^*\|_1 = 2(1-X_{\alpha^*}(t))$ where $x^* \in \mathcal{X}$ is such that $\operatorname{supp}(x^*) = \{\alpha^*\}$. It follows that we can rewrite (5.2.1) as

$$\|X(t)-x^*\|_1 \leq 2K(A-1)\exp\Big\{-mt + 2(1+\varepsilon)\sigma_* \sqrt{t\log\log t}\Big\}. \tag{5.2.2}$$

under the same conditions and notations as Proposition 5.2.1. In this form, the rate of convergence to Nash equilibrium can be put into comparison with the one derived by Giannou et al. (2021) for stochastically perturbed regularized learning in discrete time, for which the Gibbs entropic regularization and constant stepsizes also yield a decreasing leading term of order e^{-mt}.

concentration inequality

Proposition 5.2.3. *Under the same conditions as Proposition 5.2.1 and if Assumption 4.2.4 holds true,*

$$\mathbb{P}(X_{\alpha^*}(t) < 1-\delta) \leq \eta + \frac{1}{2}\sum_{\alpha\in\mathcal{A}^*} \operatorname{erfc}\left[\frac{1}{\sqrt{2}\hat{\eta}_{*\alpha}}\left(m\sqrt{t} - \frac{\tilde{K} - \log(\frac{\delta}{A-1})}{\sqrt{t}}\right)\right] \tag{5.2.3}$$

for all $\delta > 0$ *and* $t \geq 0$, *where* $\tilde{K} = \log K$ *and* $\hat{\eta}^2_{*\alpha} = \eta^T_{*\alpha}\overline{C}(\alpha^*,\alpha)\eta_{*\alpha}$ *with* $\eta^T_{*\alpha} = \begin{bmatrix}\eta_{\alpha^*} & -\eta_\alpha\end{bmatrix}$.

Remark 5.2.4. In particular, we can remove the sum appearing in the bound (5.2.3) by writing

$$\mathbb{P}(X_{\alpha^*}(t) < 1-\delta) \leq \eta + \frac{1}{2}(A-1)\operatorname{erfc}\left[\frac{1}{\sqrt{2}\hat{\eta}}\left(m\sqrt{t} - \frac{\tilde{K} - \log(\frac{\delta}{A-1})}{\sqrt{t}}\right)\right] \tag{5.2.4}$$

where the value of $\hat{\eta}$ is given by

$$\hat{\eta} = \begin{cases} \max_{\alpha\in\mathcal{A}^*} \hat{\eta}^2_{*\alpha} & \text{if } \frac{\delta}{A-1} > Ke^{-mt} \\ \min_{\alpha\in\mathcal{A}^*} \hat{\eta}^2_{*\alpha} & \text{otherwise} \end{cases}. \tag{5.2.5}$$

For large t, the rate is therefore driven by the amplitude of the maximum noise across the system.

In the next proposition, our goal is to derive a bound on the mean time needed for a "good" trajectory of (SID) to become arbitrarily close to a strict Nash equilibrium. In fact, we provide not only one but three different bounds for this expectation, in order to take into accounts situations where the size of the action space is very big, e.g., in games with billion of different strategies that we can encounter in modern applications.

hitting time

Proposition 5.2.5. *Under the same conditions as Proposition 5.2.1 and Proposition 5.2.3, let* $\tau_\delta = \inf\{t \geq 0 : X_{\alpha^*}(t) \geq 1-\delta\}$ *and denote* E *the event* $\{X(t) \in \mathcal{U}$ *for all* $t \geq 0\}$. *Then, for all* $\delta > 0$, τ_δ *verifies the following bounds :*

$$\mathbb{E}[\tau_\delta \mathbb{1}_E] \leq \begin{cases} \dfrac{(A-1)\big[\tilde{K} - \log(\frac{\delta}{A-1})\big]_+}{m} \\[2ex] \dfrac{\sqrt{A-1}\big[\tilde{K} - \log(\frac{\delta}{A-1})\big]_+}{m}\sqrt{1 + \dfrac{\hat{\eta}^2_{max}}{\big[\tilde{K} - \log\delta/(A-1)\big]_+}} \\[2ex] \dfrac{\log(A-1)}{\lambda} + \dfrac{\big[\tilde{K} - \log(\frac{\delta}{A-1})\big]_+}{\lambda\hat{\eta}^2_{min}}\left(m - \sqrt{m^2 - 2\lambda\hat{\eta}^2_{max}}\right) \end{cases} \tag{5.2.6}$$

where $\hat{\eta}^2_{min} = \min_{\alpha\in\mathcal{A}^*} \hat{\eta}_{*\alpha}$, $\hat{\eta}^2_{max} = \max_{\alpha\in\mathcal{A}^*} \hat{\eta}_{*\alpha}$, *and* $0 < \lambda < \frac{m^2}{2\hat{\eta}^2_{max}}$.

Remark 5.2.6. Focusing for now only the first two bounds of Proposition 5.2.5, we can see that the second one is smaller if and only if the condition

$$\delta < K(A-1)\exp\left\{\frac{\hat{\eta}^2_{max}}{A-2}\right\} \tag{5.2.7}$$

holds. In particular, we see that it is the case if A is big, e.g., the second bound is indeed better in a high-dimensional context. Interestingly, we see that it is also the case if some pure strategy is highly perturbed by noise in the dynamics.

minimize the dependence on A

Remark 5.2.7. It is possible to find analytically the value of λ which optimizes the third bound of Proposition 5.2.5, however its expression is quite convoluted so we do not provide its expression here. However, if our goal is to minimize the impact of the action space size A on the bound, it seems reasonable to pick λ as big as possible, i.e., to let $\lambda \to m^2/(2\hat{\eta}^2_{max})$. This yields the alternate bound

$$\mathbb{E}[\tau_\delta \mathbb{1}_E] \le \frac{2\hat{\eta}^2_{max}\log(A-1)}{m^2} + \frac{2\left[\tilde{K} - \log\left(\frac{\delta}{A-1}\right)\right]_+}{m} \cdot \frac{\hat{\eta}^2_{max}}{\hat{\eta}^2_{min}}. \tag{5.2.8}$$

5.3 PROOFS OF THEOREM 5.1.2 AND THEOREM 5.1.5

Unlike Hofbauer and Imhof (2009), our proofs of the equilibrium stability will not use arguments based on the existence of well-chosen stochastic Lyapunov functions, but will instead follow the idea brought by Bravo and Mertikopoulos (2017) and Mertikopoulos and Viossat (2016) of studying convergence of the dual processes $Y_\alpha = \log X_\alpha$ We recall from the proof of Theorem 4.1.2 that these processes verify the stochastic differential equation

dual process $Y_\alpha = \log X_\alpha$

$$dY_\alpha = \tilde{v}(X)dt + \sigma_\alpha^T(X)dW_\alpha. \tag{5.3.1}$$

We first prove an auxiliary stability result from which Theorem 5.1.2 could be derived, giving sense to the intuition that if the process remains around x^* with positive probability, then x^* should verify some equilibrium condition with respect to a vector field perturbed by the noise.

staying around point implies equilibrium

Proposition 5.3.1. *Let* $x^* \in \mathcal{X}$. *If every neighborhood* $\mathcal{U}$ *of* x^* *in* $\mathcal{X}$ *admits with positive probability an interior solution orbit* $X(t)$ *of* (SID) *and a time* t_0 *such that* $X(t) \in \mathcal{U}$ *for all* $t \ge t_0$, *then* x^* *verifies the equilibrium condition* (EQ$^\sigma$).

Proof. Assume that x^* do not verify condition (EQ$^\sigma$). Equivalently, this means that there exists $\alpha \in \mathrm{supp}(x^*)$ and $\beta \in \mathcal{A}$ such that $\tilde{v}_\alpha(x^*) < \tilde{v}_\beta(x^*)$. Let $\mathcal{U}$ be a neighborhood of x^* such that $\tilde{v}_\beta(x) - \tilde{v}_\alpha(x) \ge m > 0$ and $x_\alpha \ge m' > 0$ for all $x \in \mathcal{X}$, which always exists by the compactness of $\mathcal{X}$ and the continuity of $\tilde{v}$. Accordingly, let E be the event $\{\exists t_0$ such that $X(t) \in \mathcal{U}$ for all $t \ge t_0\}$, which has positive probability by assumption. Conditioning on E and using (5.3.1), we then have for t big enough

$$Y_\alpha(t) - Y_\beta(t) = Y_\alpha(0) - Y_\beta(0) + \int_0^{t_0} \left[\tilde{v}_\alpha - \tilde{v}_\beta\right] ds + \int_{t_0}^{t} \left[\tilde{v}_\alpha - \tilde{v}_\beta\right] ds - \xi(t) \tag{5.3.2}$$

$$\le K - mt - \xi(t) \tag{5.3.3}$$

where $\xi(t) = \int_0^t \sigma_\beta^T \, dW_\beta - \int_0^t \sigma_\alpha^T \, dW_\alpha$ and

$$K = \log\left(\frac{X_\alpha(0)}{X_\beta(0)}\right) + \left(m + \sup_{\mathcal{X}}\left|\tilde{v}_\alpha(x) - \tilde{v}_\beta(x)\right|\right) t_0. \tag{5.3.4}$$

Noticing that[2] $|K| < \infty$ (a.s.) and that ξ can be written into the form $\xi(t) = \int_0^t \sigma^T(X) \, dW$, where W is a vector of Brownian motions and σ is a continuous bounded process adapted to the filtration generated by W, an application of Corollary 2.3.4 readily yields $Y_\alpha(t) - Y_\beta(t) \to -\infty$ (a.s.) conditionally on E. This in turn implies that $X_\alpha(t) \to 0$ on E, contradicting our assumption that $X(t)$ stays into $\mathcal{U}$ for all $t \geq t_0$. Hence x^* must verify the equilibrium condition (EQ$^\sigma$), proving the proposition. □

With Proposition 5.3.1 in hand, we are now ready to prove Theorem 5.1.2, which boils down to simply show that limits of interior solutions and stochastically stable states verify the assumption of the proposition.

Proof of Theorem 5.1.2. If $\mathbb{P}(\lim_{t\to\infty} X(t) = x^*) > 0$, then for every neighborhood $\mathcal{U}$ of x^* there exists with positive probability a t_0 such that $X(t) \in \mathcal{U}$ for all $t \geq t_0$, and $X(t)$ is an interior solution orbit by definition. As such, it verifies the assumption of Proposition 5.3.1 and so the equilibrium result (EQ$^\sigma$) follows.

Similarly, if x^* is stochastically Lyapunov stable, then there exists a neighborhood $\mathcal{U}_0$ of x^* such that $X(t) \in \mathcal{U}$ for all $t \geq 0$ with positive probability whenever $X(0) \in \mathcal{U}_0$. In particular, as $\mathcal{U}_0$ is a neighborhood with respect to the relative topology induced by $\mathcal{X}$, the set $\mathcal{U}_0 \cap \mathcal{X}^{\circ}$ is nonempty, and so the condition of Proposition 5.3.1 holds with $t_0 \equiv 0$ when taking an initial condition in this intersection. □

On its end, the proof of Theorem 5.1.5 is more tricky and convoluted. Here, we follow the same steps as Bravo and Mertikopoulos (2017) and Mertikopoulos and Viossat (2016), where we state that showing the asymptotically stochastic stability in (SID) reduces to proving that the hitting time of an associated Brownian motion with drift is finite with probability lesser than an arbitrary $\varepsilon > 0$, then use a result based on Girsanov theorem to conclude.

We also mention here that the proof of the stability in both of these articles contains a small mistake in their reasoning, where they forget that some bound is only true on an event and not on the entire space. Our following theorem thus correct this issue, but boils down essentially to the same underlying idea.

Proof of Theorem 5.1.5. Let $x^* \in \mathcal{X}$ verifying strictly the equilibrium condition (EQ$^\sigma$). As mentioned in Remark 5.1.6, this can be interpreted by saying that x^* is a strict Nash equilibrium of the modified game $\tilde{\Gamma}$. Hence, x^* is in particular a pure strategy. Accordingly, let $\alpha^* \in \mathcal{A}$ be its associated pure action, and write $\mathcal{A}^* \equiv \mathcal{A} \setminus \{x^*\}$. Introducing the stochastic process $Z_\alpha = Y_\alpha - Y_{\alpha^*}$ and noticing from an easy computation that $X_{\alpha^*}(t) = \left[1 + \sum_{\alpha\in\mathcal{A}^*} \exp(Z_\alpha)\right]^{-1}$, we have

$$\lim_{t\to\infty} X(t) = x^* \text{ if and only if } Z_\alpha(t) \to -\infty \text{ for all } \alpha \in \mathcal{A}^*. \tag{5.3.5}$$

To proceed, fix a threshold $\varepsilon > 0$ and let $\mathcal{U}_0$ be a neighborhood of x^* in $\mathcal{X}$. As x^* is a strict Nash equilibrium in $\tilde{\Gamma}$, we know in particular that there exists a neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of x^* and a constant $m > 0$ such that $\tilde{v}_{\alpha^*}(x) - \tilde{v}_\alpha(x) \geq m > 0$ for all $x \in \mathcal{U}$ and $\alpha \in \mathcal{A}^*$. As the processes X and Z are

2 By assumption $X(t)$ is an interior orbit of (SID), so its support is all of $\mathcal{A}$.

(almost surely) continuous, (5.3.5) implies that we can pick some $M > 0$ big enough so that $X(t) \in \mathcal{U}$ whenever $Z_\alpha(t) \leq -M$ for all $\alpha \in \mathcal{A}^*$.

The main idea behind this proof is to show that if M is well-chosen and if the initial conditions verify $Z_\alpha(0) \leq -2M$, then $X(t) \in \mathcal{U} \subseteq \mathcal{U}_0$ for all $t \geq 0$ and $X(t) \to x^*$ with probability at least $1 - \varepsilon$.

Assume that $Z_\alpha(t) \leq -2M$ for all $\alpha \in \mathcal{A}^*$, and let $\tau_{\mathcal{U}} = \inf\{t \geq 0 : X(t) \notin \mathcal{U}\}$ be the first exist time of $X(t)$ from $\mathcal{U}$. On the event $\{\tau_{\mathcal{U}} = \infty\}$, we have

$$Z_\alpha(t) = Z_\alpha(0) + \int_0^t [\tilde{v}_\alpha - \tilde{v}_{\alpha^*}]\, ds - \xi(t) \leq -2M - mt - \xi_\alpha(t) \tag{5.3.6}$$

where $\xi_\alpha(t) = \int_0^t \sigma_{*\alpha}^T(X) dW_{*\alpha}$ with $\sigma_{*\alpha}^T = \begin{bmatrix}\sigma_{\alpha^*} & -\sigma_\alpha\end{bmatrix}$ and $W_{*\alpha} = \begin{bmatrix}W_{\alpha^*} & W_\alpha\end{bmatrix}$. From the usual argument based on the time-change theorem for martingales, there exists a Brownian motion $\widetilde{W}_\alpha$ such that $\xi_\alpha(t) = \widetilde{W}_\alpha(\rho_\alpha(t))$ where $\rho_\alpha(t)$ is the quadratic variation of ξ_α verifying

$$\rho_\alpha(t) \leq \left[\sup_{x \in \mathcal{X}, \alpha \in \mathcal{A}^*} \sigma_{*\alpha}^T(x) Cor(W_{*\alpha}, W_{*\alpha}) \sigma_{*\alpha}(x)\right] \times t. \tag{5.3.7}$$

For convenience, let us denote $\hat{\sigma}^2$ the supremum appearing in equation (5.3.7). The law of iterated logarithms combined with (5.3.6) can then be used to deduce that on $\{\tau_{\mathcal{U}} = \infty\}$, we have $Z_\alpha(t) \to -\infty$ (a.s.) for all $\alpha \in \mathcal{A}^*$. As such, it is sufficient to prove that $\mathbb{P}(\tau_{\mathcal{U}} = \infty) \geq 1 - \varepsilon$ to obtain the stochastically asymptotic stability.

To do so, let us define the random time $\tau_0 = \inf\left\{t \geq 0 : \inf_{\alpha \in \mathcal{A}^*} \widetilde{W}_z(t) \leq -M - \frac{mt}{\hat{\sigma}^2}\right\}$. We first show that if $\tau_0 = \infty$, then we also have $\tau_{\mathcal{U}} = \infty$. For the sole purpose of proving a contradiction, assume on the other hand that we are on the event $\{\tau_0 = \infty\} \cap \{\tau_{\mathcal{U}} < \infty\}$. From the expression of Z_α, we then get

$$Z_\alpha(\tau_{\mathcal{U}}) \leq -2M - m\tau_{\mathcal{U}} - \widetilde{W}_\alpha(\rho_\alpha(\tau_{\mathcal{U}})), \tag{5.3.8}$$

which is indeed defined because $\tau_{\mathcal{U}} < \infty$ and $X(t) \in \mathcal{U}$ for all $t < \tau_{\mathcal{U}}$. However, on $\{\tau_0 = \infty\}$ we have $\widetilde{W}_\alpha(t) > -M - \frac{mt}{\hat{\sigma}^2}$ for all $t \geq 0$ and all $\alpha \in \mathcal{A}^*$, so in particular

$$\widetilde{W}_\alpha(\rho_\alpha(\tau_{\mathcal{U}})) > -M - \frac{m\rho(\tau_{\mathcal{U}})}{\hat{\sigma}^2} \geq -M - m\tau_{\mathcal{U}} \tag{5.3.9}$$

where the last inequality comes from $\rho_\alpha(t) \leq \hat{\sigma}^2 t$. Combined with (5.3.8), it yields $Z_\alpha(\tau_{\mathcal{U}}) \leq -M$ for all $\alpha \in \mathcal{A}^*$, and so $X(\tau_{\mathcal{U}}) \in \mathcal{U}$. This rises a contradiction by definition of the exit time $\tau_{\mathcal{U}}$, hence proving that $\tau_{\mathcal{U}} = \infty$ whenever $\tau_0 = \infty$. Taking the probability of both events, it then holds that $\mathbb{P}(\tau_{\mathcal{U}} = \infty) \geq \mathbb{P}(\tau_0 = \infty)$. But the event $\{\tau_0 < \infty\}$ can be expressed as $\{\tau_0 < \infty\} = \bigcup_{\alpha \in \mathcal{A}^*} \{\tau_\alpha < \infty\}$ where $\tau_\alpha = \inf\left\{t \geq 0 : \widetilde{W}_\alpha(t) + \frac{mt}{\hat{\sigma}^2} = -M\right\}$. We notice that τ_α is an hitting time of a Brownian motion with positive drift, for which the analysis in Subsection 3.5.C of Karatzas and Shreve (1998) yields $\mathbb{P}(\tau_\alpha < \infty) = \exp\left\{\frac{-2mM}{\hat{\sigma}^2}\right\}$. Therefore, if we take M big enough so that $M > -\frac{\hat{\sigma}^2 \log \varepsilon / A}{2m}$, then $\mathbb{P}(\tau_0 < \infty) \leq \sum_{\alpha \in \mathcal{A}^*} \mathbb{P}(\tau_\alpha < \infty) < \varepsilon$, from which we finally deduce $\mathbb{P}(\tau_{\mathcal{U}} = \infty) \geq \mathbb{P}(\tau_0 = \infty) \geq 1 - \varepsilon$. By the previous discussion, we have then proven that x^* is stochastically asymptotically stable. □

5.4 PROOFS OF PROPOSITION 5.2.1 AND PROPOSITION 5.2.3

The proofs of the rates of convergence to strict Nash equilibria follow the same kind of ideas as those used for the speed of extinction of dominated strategies in Chapter 4. Indeed, we will show that the pure strategy associated to the strict Nash equilibrium dominates all other pure actions conditioned on a event with positive probability, and so controlling the conditional speed of extinction for each of these pure actions yields a speed of convergence towards the strict Nash equilibrium. Unlike dominated strategies, as all of our results will only hold on an event, we also need to be particularly careful with conditioning arguments.

First, we provide some common groundwork and important bounds that will be used throughout all proofs concerning rates of convergence.

Let α^* be a strict Nash equilibrium of the game Γ. Following the same arguments as in the proof of Corollary 5.1.7, we know that α^* also strictly verifies the equilibrium condition (EQ^σ), i.e., $\tilde{v}_{\alpha^*}(x^*) > \tilde{v}_\alpha(x^*)$ for all $\alpha \in \mathcal{A}^*$, where $\operatorname{supp}(x^*) = \{\alpha^*\}$. Accordingly, let $\mathcal{U}$ be a neighborhood of x^* such that $\inf_{\alpha \neq \alpha^*, x \in \mathcal{U}} \tilde{v}_{\alpha^*}(x) - \tilde{v}_\alpha(x) > 0$, which exists by continuity of $\tilde{v}$.

Now, Corollary 5.1.7 yields that x^* is asymptotically stochastically stable with respect to (SID), and so in particular there exists for all $\eta > 0$ a smaller neighborhood $\mathcal{U}_\eta \subseteq \mathcal{U}$ of x^* such that $X(t) \in \mathcal{U}$ for every $t \geq 0$ with probability at least $1-\eta$ whenever $X(0) \in \mathcal{U}_\eta$. For simplicity, let us denote $E = \{X(t) \in \mathcal{U} \text{ for all } t \geq 0\}$, and assume that we choose $X(0) \in \mathcal{U}_\eta$. Fixing $\alpha \in \mathcal{A}^*$ and using the stochastic differential equation formulation (5.3.1) of the dual processes $Y_\alpha = \log X_\alpha$, we then have on E :

$$dY_\alpha - dY_{\alpha^*} \leq -m\,dt - d\xi_\alpha \leq -m\,dt - \inf_{\alpha \in \mathcal{A}^*} d\xi_\alpha \tag{5.4.1}$$

where $\xi_\alpha(t) = \int_0^t \sigma_{*\alpha}(X(s))\, dW_{*\alpha}(s)$ with $\sigma_{*\alpha}^T = \begin{bmatrix} \sigma_{\alpha^*} & -\sigma_\alpha \end{bmatrix}$ and $W_{*\alpha}^T = \begin{bmatrix} W_{\alpha^*} & W_\alpha \end{bmatrix}$. Turning back to the stochastic integral interpretation, we therefore get on E

$$Y_\alpha(t) \leq Y_{\alpha^*}(t) + Y_\alpha(0) - Y_{\alpha^*}(0) - mt - \inf_{\alpha \in \mathcal{A}^*} \xi_\alpha(t) \leq \log(K) - mt - \inf_{\alpha \in \mathcal{A}^*} \xi_\alpha(t), \tag{5.4.2}$$

where $\log(K) = -Y_{\alpha^*}(0)$ only depends on the initial condition. Writing $X_\alpha = \exp(Y_\alpha)$ then immediately yields

$$X_\alpha(t) \leq K \exp\left\{-mt - \inf_{\alpha \in \mathcal{A}^*} \xi_\alpha(t)\right\} \tag{5.4.3}$$

conditioned on the event E.

On the other hand, as $X(t)$ is in the simplex $\mathcal{X}$ for all $t \geq 0$, we have $X_{\alpha^*}(t) = 1 - \sum_{\alpha \in \mathcal{A}^*} X_\alpha(t)$, which combined with (5.4.3) readily gives the bound

$$X_{\alpha^*}(t) \geq 1 - K(A-1) \exp\left\{-mt - \inf_{\alpha \in \mathcal{A}^*} \xi_\alpha(t)\right\} \tag{5.4.4}$$

whenever the event E is true. In particular, it means that (5.4.4) holds with probability at least $1-\eta$.

The idea behind the following proofs is then to bound the random quantity $\inf_{\alpha \in \mathcal{A}^*} \xi_\alpha(t)$ appearing on the right-hand side of (5.4.4) by invoking as usual the time-change theorem for martingales.

Proof of Proposition 5.2.1. Let $\alpha \in \mathcal{A}^*$ be fixed. Using the time-change theorem for martingales, there exists a standard Brownian motion $\widetilde{W}_\alpha$ such that $\xi_\alpha(t) = \widetilde{W}_\alpha(\rho_\alpha(t))$ where $\rho_\alpha(t) = [\xi_\alpha]_t$ is the quadratic variation of ξ_α.

If $\rho_\alpha(t) \to \infty$, then for all $\varepsilon > 0$ the law of iterated logarithms yields $|\xi_\alpha(t)| \leq (1+\varepsilon)\sqrt{2\rho_\alpha(t)\log\log\rho_\alpha(t)}$ (a.s.) for t big enough. In particular, if we are on the event E, then a quick computation gives $\rho_\alpha(t) \leq \sigma_*^2 t$ where $\sigma_*^2 = \sup_{\alpha\in\mathcal{A}^*,\, x\in\mathcal{U}} \sigma_{*\alpha}^T(x)\overline{C}(\alpha^*,\alpha)\sigma_{*\alpha}(x) < \infty$, and so

$$|\xi_\alpha(t)| \leq 2(1+\varepsilon)\sigma_*\sqrt{t\log\log t} \quad (a.s.) \tag{5.4.5}$$

conditionally on E [3], for t large enough. On the other hand, if $\rho_\alpha(t) \to \rho_\alpha(\infty) < \infty$, then ξ_α is bounded on E, which again yields that $|\xi_\alpha(t)| \leq 2(1+\varepsilon)\sigma_*\sqrt{t\log\log t}$ for t big enough. It allows us to conclude that (5.4.5) is true on E whatever the limit of ρ_α, and as its right-hand side does not depend on α, we therefore get

$$\inf_{\alpha\in\mathcal{A}^*}|\xi_\alpha(t)| \leq 2(1+\varepsilon)\sigma_*\sqrt{t\log\log t} \quad (a.s.) \tag{5.4.6}$$

conditionally on E when t is big enough. The rate (5.2.1) then follows by injecting bound (5.4.6) inside (5.4.4). □

Proof of Proposition 5.2.3. Unlike Proposition 5.2.1, we now suppose that Assumption 4.2.4 holds true, i.e., that the noise term can be decomposed as $\sigma_\alpha^T(X)dW_\alpha = \eta_\alpha^T dW_\alpha + dF$. When taking the dual difference $dY_\alpha - dY_{\alpha^*}$ as in (5.4.1), the term dF will then cancel out, so that (5.4.4) still holds true on E but with ξ_α now defined by $\xi_\alpha(t) = \int_0^t \eta_{*\alpha}(X(s))\, dW_{*\alpha}(s)$ where $\eta_{*\alpha}^T = \begin{bmatrix}\eta_{\alpha^*} & -\eta_\alpha\end{bmatrix}$. In particular, the time-change theorem yields $\xi_\alpha(t) = \widetilde{W}_\alpha(\hat\eta_{*\alpha}^2 t)$ where $\hat\eta_{*\alpha}^2 = \eta_{*\alpha}^T\overline{C}(\alpha^*,\alpha)\eta_{*\alpha}$ [4].

We can therefore write the following flow of inclusions :

$$\begin{aligned}
\{X_{\alpha^*}(t) < 1-\delta\}\cap E &\subseteq \left\{K(A-1)\exp\left[-mt - \inf_{\alpha\in\mathcal{A}^*}\widetilde{W}_\alpha\left(\hat\eta_{*\alpha}^2 t\right)\right] > \delta\right\}\cap E && (5.4.7)\\
&= \left\{\inf_{\alpha\in\mathcal{A}^*}\widetilde{W}_\alpha\left(\hat\eta_{*\alpha}^2 t\right) < L(\delta)\right\}\cap E && (5.4.8)\\
&\subseteq \left\{\inf_{\alpha\in\mathcal{A}^*}\widetilde{W}_\alpha\left(\hat\eta_{*\alpha}^2 t\right) < L(\delta)\right\} && (5.4.9)\\
&= \bigcup_{\alpha\in\mathcal{A}^*}\left\{\widetilde{W}_\alpha\left(\hat\eta_{*\alpha}^2 t\right) < L(\delta)\right\} && (5.4.10)
\end{aligned}$$

where $L(\delta) = \tilde{K} - mt - \log\left(\frac{\delta}{A-1}\right)$. Taking the probability of the events, it then means that

$$\mathbb{P}(\{X_{\alpha^*}(t) < 1-\delta\}\cap E) \leq \sum_{\alpha\in\alpha^*}\mathbb{P}\left(\widetilde{W}_\alpha\left(\hat\eta_{*\alpha}^2 t\right) < L(\delta)\right). \tag{5.4.11}$$

---

3 Recall that $\mathbb{P}(E) > 0$, so $\mathbb{P}(A \mid E) = 1$ for all event A with $\mathbb{P}(A) = 1$.

4 Notice that the integrand of ξ_α do not depend on time by Assumption 4.2.4, so its quadratic variation is readily given by $[\xi_\alpha]_t = \hat\eta_{*\alpha}^2 t$ following usual computations.

Using the same computations as in the proof of Proposition 4.2.6 for each term of the sum, this yields the explicit bound

$$\mathbb{P}(\{X_{\alpha^*}(t) < 1-\delta\} \cap E) \leq \frac{1}{2} \sum_{\alpha \in \mathcal{A}^*} \operatorname{erfc}\left[\frac{1}{\sqrt{2}\hat{\eta}_{*\alpha}}\left(m\sqrt{t} - \frac{\tilde{K} - \log(\frac{\delta}{A-1})}{\sqrt{t}}\right)\right]; \quad (5.4.12)$$

from which we deduce (5.2.3) by noticing that

$$\mathbb{P}(X_{\alpha^*}(t) < 1-\delta) \leq \mathbb{P}(\{X_{\alpha^*}(t) < 1-\delta\} \cap E) + \mathbb{P}(E^c) \quad (5.4.13)$$

$$\leq \mathbb{P}(\{X_{\alpha^*}(t) < 1-\delta\} \cap E) + \eta. \quad (5.4.14)$$

□

5.5 PROOF OF PROPOSITION 5.2.5

The general idea behind this proof belongs to the same realm as the one of Proposition 4.2.6 b), i.e., bounding the random time by the hitting time of a Brownian motion with drift, for which we know many properties thanks to Girsanov theorem. But the main difference here is that the bound will in fact be the maximum of such hitting times, yielding more freedom about *how* we can deal with this quantity.

To make this remark more precise, let $\tau_\delta = \inf\{t \geq 0 : X_{\alpha^*}(t) \geq 1-\delta\}$ and denote E the event $\{X(t) \in \mathcal{U}$ for all $t \geq 0\}$ as in the statement of the proposition. Based on the bound (5.4.4) on E for a noise verifying Assumption 4.2.4 (cf. proof of Proposition 5.2.3), we then have

$$\tau_\delta \, \mathbb{1}_E \leq \inf\left\{t \geq 0 : K(A-1)\exp\left[-mt - \inf_{\alpha \in \mathcal{A}^*} \xi_\alpha(t)\right] \leq \delta\right\} \quad (5.5.1)$$

$$= \inf\left\{t \geq 0 : \inf_{\alpha \in \mathcal{A}^*} mt + \xi_\alpha(t) \geq a\right\} \quad (5.5.2)$$

$$= \max_{\alpha \in \mathcal{A}^*} \tau_\alpha \quad (5.5.3)$$

where $a = [\tilde{K} - \log(\frac{\delta}{A-1})]_+$[5], and $\tau_\alpha = \inf\{t \geq 0 : mt + \xi_\alpha(t) = a\}$. From the classical application of time-change theorem, $\xi_\alpha(t) = \widetilde{W}_\alpha(\hat{\eta}^2_{*\alpha} t)$ for a Brownian motion $\widetilde{W}_\alpha$, so we can rewrite each random time τ_α as

$$\tau_\alpha = \inf\left\{t \geq 0 : mt + \widetilde{W}_\alpha\left(\hat{\eta}^2_{*\alpha} t\right) = a\right\} = \frac{1}{\hat{\eta}^2_{*\alpha}} \tau^h_\alpha, \quad (5.5.4)$$

where τ^h_α is the first hitting time at level a of a Brownian motion with positive drift $\mu_\alpha = m/\hat{\eta}^2_{*\alpha}$. All in all, we therefore obtain

$$\mathbb{E}[\tau_\delta \, \mathbb{1}_E] \leq \mathbb{E}\left[\max_{\alpha \in \mathcal{A}^*} \frac{1}{\hat{\eta}^2_{*\alpha}} \tau^h_\alpha\right]. \quad (5.5.5)$$

As mentioned at the start of this section, the question is now how to bound this expectation of a maximum of hitting times. Indeed, an obvious first idea would be to bound it by the expectation of the sum, but this will lead to a multiplicative factor $(A-1)$ in front of our bound, which may not be ideal when dealing with large action spaces. Accordingly, we want to use our knowledge on the distribution of the hitting times in order to construct tighter bounds,

5 If $\log[\delta/(A-1)] \geq \tilde{K}$, then the random times can be shown to all be equal to 0.

especially with respect to the size of the action space. Such constructions will be based on the following lemma, which yields bounds on the maximum for each increasing convex function from $\mathbb{R}^+$ to $\mathbb{R}$. This approach is inspired by classical results for maximal bounds in concentration inequalities (see e.g., Subsection 2.5 of Boucheron et al., 2013).

convex bound on the max

Lemma 5.5.1. *Let $X_1, \dots, X_n$ be random variables with finite mean, and let $\psi\colon \mathbb{R}^+ \to \mathbb{R}$ be an increasing convex function such that $\mathbb{E}[\psi(|X_i|)] < \infty$ for all $i = 1, \dots, n$. Then,*

$$\mathbb{E}\left[\max_{1\leq i\leq n}|X_i|\right] \leq \psi^{-1}\left(\sum_{i=1}^{n} \mathbb{E}[\psi(|X_i|)]\right) \tag{5.5.6}$$

where ψ^{-1} is the generalized inverse of ψ.

Proof. In virtue of Jensen's inequality, we have $\psi(\mathbb{E}[\max_i|X_i|]) \leq \mathbb{E}\left[\phi(\max_i|X_i|)\right]$. Then, using the facts that ψ is increasing and the classical bound $\max_i |X_i| \leq \sum_i|X_i|$, we therefore get

$$\psi\left(\mathbb{E}\left[\max_i|X_i|\right]\right) \leq \mathbb{E}\left[\max_i \psi(|X_i|)\right] \leq \mathbb{E}\left[\sum\nolimits_i \psi(|X_i|)\right], \tag{5.5.7}$$

from which (5.5.6) is readily obtained by applying the (increasing) function ψ^{-1} on both sides. □

In virtue of Lemma 5.5.1 applied to the right-hand side of (5.5.5), we then have

$$\mathbb{E}[\tau_\delta\, \mathbb{1}_E] \leq \psi^{-1}\left(\sum_{\alpha\in\mathcal{A}^*} \mathbb{E}\left[\psi\left(\frac{1}{\hat{\eta}^2_{*\alpha}}\tau^h_\alpha\right)\right]\right) \tag{5.5.8}$$

for all increasing convex function $\psi\colon \mathbb{R}^+ \to \mathbb{R}$. For instance, this inequality holds for power functions $x \mapsto x^k$ where $k > 0$, which leads to the computation of the moments of τ^h_α; and for the exponential function $x \mapsto e^{\lambda x}$ with $\lambda > 0$. In fact, those functions are exactly those that are used to obtain the bounds of Proposition 5.2.5.

An interesting quantity to derive in all of these cases is therefore the moment generating function $\varphi_\alpha(\lambda)$ of τ^h_α, especially when $\lambda > 0$. However, up to our knowledge, most of the literature on this subject is only concerned on the computation of the *Laplace transform* of τ^h_α, i.e., when λ is negative, using a direct application of Girsanov theorem (see Karatzas and Shreve, 1998). Such an approach cannot be taken for $\lambda > 0$, as the moment generating function of the hitting time of a standard Brownian motion without drift is infinite. But this does not mean that it is also infinite for a Brownian motion with drift. Indeed, the following lemma provides an explicit expression for the moment generating function of τ^h_α for λ small enough but still positive whenever the drift is not zero.

moment generating function of hitting time

Lemma 5.5.2. *Let $\tau = \inf\{t \geq 0 : \mu t + W_t = a\}$ be the hitting time at level a of a Brownian motion with positive drift μ. Then, the moment generating function $\varphi_\tau(\lambda) \equiv \mathbb{E}\left[e^{\lambda\tau}\right]$ of τ is finite for all $\lambda < \mu^2/2$, and is given by $\varphi_\tau(\lambda) = \exp\left\{a\left(\mu - \sqrt{\mu^2 - 2\lambda}\right)\right\}$.*

Proof. From a classical application of Girsanov theorem (see e.g., Subsection 3.5.C of Karatzas and Shreve, 1998 for a complete discussion), it can be shown that the distribution of τ is absolutely continuous with respect to Lebesgue measure and that its associated density f_τ is given by

$$f_\tau(t) = \frac{a}{\sqrt{2\pi t^3}} \exp\left\{-\frac{(a-\mu t)^2}{2t}\right\} \mathbb{1}_{(0,\infty]}(t). \tag{5.5.9}$$

From the equivalent $-\frac{(a-\mu t)^2}{2t} + \lambda t \sim \left(-\frac{\mu^2}{2} + \lambda\right)t$, we see that $e^{\lambda t} f_\tau(t)$ is integrable on $\mathbb{R}$ as long as $-\mu^2/2 + \lambda < 0$, i.e., when $\lambda < \mu^2/2$, which immediately yields that ϕ_τ is finite for such λ's. On the other hand, an easy computation gives

$$-\frac{(a-\mu t)^2}{2t} + \lambda t = -\frac{1}{2t}\left(a - \sqrt{\mu^2 - 2\lambda}t\right)^2 + a\left(\mu - \sqrt{\mu^2 - 2\lambda}\right), \tag{5.5.10}$$

so that $e^{\lambda t} = \exp\left\{a\left(\mu - \sqrt{\mu^2 - 2\lambda}\right)\right\} f_{\tilde{\tau}}(t)$ where $f_{\tilde{\tau}}$ is the density of the first hitting time of a Brownian motion with drift $\sqrt{\mu^2 - 2\lambda}$. Therefore,

$$\phi_\tau(\lambda) = \exp\left\{a\left(\mu - \sqrt{\mu^2 - 2\lambda}\right)\right\} \int_{\mathbb{R}} f_{\tilde{\tau}}(t)\, dt = \exp\left\{a\left(\mu - \sqrt{\mu^2 - 2\lambda}\right)\right\} \tag{5.5.11}$$

which complete the proof. □

Armed with this lemma and the previously established inequalities, we are know ready to prove the various bounds of Proposition 5.2.5.

Proof of (5.2.6). From Lemma 5.5.2, we know that the moment generating function of τ^h_α is finite on an open interval containing 0, and so τ^h_α admits finite moments of all orders. In particular, its first two moments can be recovered as $\mathbb{E}\left[\tau^h_\alpha\right] = \varphi'(0)$ and $\mathbb{E}\left[\left(\tau^h_\alpha\right)^2\right] = \varphi''(0)$, leading after some computations to

$$\mathbb{E}\left[\tau^h_\alpha\right] = \frac{a}{\mu_\alpha}{}^{6} \quad \text{and} \quad \mathbb{E}\left[\left(\tau^h_\alpha\right)^2\right] = \frac{a^2}{\mu_\alpha^2}\left(1 + \frac{1}{a\mu_\alpha}\right). \tag{5.5.12}$$

Applying inequality (5.5.8) with function $x \mapsto x$ therefore gives

$$\mathbb{E}[\tau_\delta \mathbb{1}_E] \leq \sum_{\alpha \in \mathcal{A}^*} \frac{1}{\hat{\eta}^2_{*\alpha}} \mathbb{E}\left[\tau^h_\alpha\right] = \sum_{\alpha \in \mathcal{A}^*} \frac{a}{m} = \frac{(A-1)a}{m} \tag{5.5.13}$$

which yields the first bound of Proposition 5.2.5.

Similarly, when we choose instead the function $x \mapsto x^2$, we obtain

$$\mathbb{E}[\tau_\delta \mathbb{1}_E] \leq \left[\sum_{\alpha \in \mathcal{A}^*} \frac{a^2}{m^2}\left(1 + \frac{\hat{\eta}^2_{*\alpha}}{ma}\right)\right]^{1/2} \leq \frac{\sqrt{A-1}a}{m}\sqrt{1 + \frac{\hat{\eta}^2_{max}}{ma}}, \tag{5.5.14}$$

leading to the second bound.

To derive the last bound, let us define the function $\phi_\lambda \colon x \mapsto e^{\lambda x}$ with $0 < \lambda < \frac{m^2}{2\hat{\eta}^2_{max}}$, which is obviously convex and increasing on $\mathbb{R}^+$. We have $\phi_\lambda\left(\frac{\tau^h_\alpha}{\hat{\eta}^2_{*\alpha}}\right) = \exp\left\{\frac{\lambda}{\hat{\eta}^2_{*\alpha}}\tau^h_\alpha\right\}$, and thanks to the assumption on λ we can see that $\frac{\lambda}{\hat{\eta}^2_{*\alpha}} <$

6 In fact this expectation can be recovered more easily by a direct application of Dynkin's formula, but we still prove it here for completeness.

$\mu_\alpha^2/2$ for all $\alpha \in \mathcal{A}^*$. Hence, Lemma 5.5.2 yields that $\mathbb{E}\left[\phi_\lambda\left(\frac{\tau_\alpha^h}{\hat{\eta}_{*\alpha}^2}\right)\right] = \varphi_{\tau_\alpha^h}(\lambda/\hat{\eta}_{*\alpha}^2)$ is finite and given by

$$\mathbb{E}\left[\phi_\lambda\left(\frac{\tau_\alpha^h}{\hat{\eta}_{*\alpha}^2}\right)\right] = \exp\left\{\frac{a}{\hat{\eta}_{*\alpha}^2}\left(m - \sqrt{m^2 - 2\hat{\eta}_{*\alpha}^2\lambda}\right)\right\} \tag{5.5.15}$$

$$\leq \exp\left\{\frac{a}{\hat{\eta}_{min}^2}\left(m - \sqrt{m^2 - 2\hat{\eta}_{max}^2\lambda}\right)\right\}. \tag{5.5.16}$$

Accordingly, by applying (5.5.8) with function ϕ_λ we readily get

$$\mathbb{E}[\tau_\delta \mathbb{1}_E] \leq \phi_\lambda^{-1}\left(\sum_{\alpha \in \mathcal{A}^*} \mathbb{E}\left[\phi_\lambda\left(\frac{\tau_\alpha^h}{\hat{\eta}_{*\alpha}^2}\right)\right]\right) \tag{5.5.17}$$

$$\leq \frac{1}{\lambda}\log\left[(A-1)\exp\left\{\frac{a}{\hat{\eta}_{min}^2}\left(m - \sqrt{m^2 - 2\hat{\eta}_{max}^2\lambda}\right)\right\}\right] \tag{5.5.18}$$

$$\leq \frac{\log(A-1)}{\lambda} + \frac{a}{\lambda\hat{\eta}_{min}^2}\left(m - \sqrt{m^2 - 2\hat{\eta}_{max}^2\lambda}\right); \tag{5.5.19}$$

thus proving the third and last bound of the proposition.

□

Part II

FURTHER TOPICS IN STOCHASTIC DYNAMICS

6

STOCHASTIC EXPONENTIAL DYNAMICS

We discuss an extension of stochastic imitation dynamics (SID) to more general noise frameworks, that we name *stochastic exponential dynamics*. These dynamics include many stochastic variants of the replicator dynamics introduced in the literature, and we prove that their stability can then be studied in an unified manner.

6.1 A GENERAL FRAMEWORK

Up to now we have only studied the stochastic behavior of imitation dynamics under a particular source of noise coming directly from the revision protocol. However such deterministic dynamics could also emerge from totally different frameworks, and so could lead to various kind of noise terms depending on how the perturbations were injected into the system. In this section, we therefore introduce a more general dynamics that encapsulate many of these different noise frameworks.

First, notice that all deterministic imitations dynamics (monotone or not) can be written into the form

$$\dot{x}_\alpha = x_\alpha L_\alpha(x) \tag{ED}$$

where L is a Lipschitz continuous vector field.

Taking a step away from the imitation framework, we then say that a deterministic (game) dynamics is an *exponential dynamics* [1] if it can be written into the form (ED) where L is a Lipschitz continuous vector field verifying the consistency condition $\sum_\alpha x_\alpha L_\alpha(x) = 0$. Under this condition, it can be shown that there exists a unique solution of (ED) staying in $\mathcal{X}^\circ$ for all $t \geq 0$ whenever it starts from $\mathcal{X}^\circ$. In this chapter we are mostly interested into the stochastic aspects of (ED) so we will delay to Chapter 7 a study of its deterministic stability properties.

exponential dynamics (ED)

To take into account arbitrary stochastic perturbations in (ED), we introduce the associated *stochastic exponential dynamics* as

stochastic exponential dynamics (SED)

$$dX_\alpha = X_\alpha \hat{L}_\alpha(X)dt + X_\alpha \sigma_\alpha^T(X)dW_\alpha \tag{SED}$$

where $\hat{L}_\alpha = L_\alpha + I_\alpha$ with I_α an Itô correction term that comes from the way that the noise was injected into the framework, i.e., an intrinsic bias due to the environment and that may depend on the noise coefficients. The reason of why we consider this bias term will becomes clear in the examples of the

1 This denomination comes from the fact that if $x \in \mathcal{X}$ is a solution to (ED), then it verifies $x_\alpha(t) = x_\alpha(0)\exp\left\{\int_0^t L_\alpha(x(s))\,ds\right\}$. Such dynamics are also sometimes called *general* imitation dynamics (cf. Mertikopoulos and Viossat, 2022) or *regular selection dynamics* (cf. Ritzberger and Weibull, 1995), but we prefer here to employ the term "exponential" to separate ourselves from the precise framework of imitation protocols.

next section. As for stochastic imitation dynamics, W_α corresponds to a vector of possibly correlated Brownian motions and σ_α to the amplitude of the noise along this vector.

To ensure existence and uniqueness of solutions to (SED) in $\mathcal{X}^\circ$ (cf. proof of Proposition 3.3.4), we also assume that both I_α and σ_α are Lipschitz continuous and verify the consistency conditions

$$\sum\nolimits_\alpha x_\alpha I_\alpha(x) = 0 \quad \text{and} \quad \sum\nolimits_\alpha x_\alpha \sigma_\alpha^T(x) dW_\alpha = 0 \tag{6.1.1}$$

for all $x \in \mathbb{R}^A$.

6.2 EXAMPLES OF STOCHASTIC GAME DYNAMICS

The framework of stochastic exponential dynamics has the advantage to include many important game dynamics from the literature, such as different forms of stochastic replicator dynamics. We provide below examples of such dynamics that can be expressed as stochastic exponential dynamics. These examples would also illustrate why we consider an Ito correction term in the formulation of (SED).

Example 6.2.1 (Stochastic Imitation Dynamics)**.** Of course, we can recover the stochastic imitation dynamics (SID) introduced in Chapter 3, and so even if the underlying conditional imitation rates are not monotone. Indeed, we just need to take

$$L_\alpha(x) = \sum\nolimits_\beta x_\beta \left[r_{\beta\alpha} - r_{\alpha\beta}\right] \tag{6.2.1}$$

$$I_\alpha(x) = 0 \tag{6.2.2}$$

$$\sigma_\alpha^T(x) dW = \sum\nolimits_\beta x_\beta \left[\sigma_{\beta\alpha}(x) dW_{\beta\alpha} - \sigma_{\alpha\beta}(x) dW_{\alpha\beta}\right] \tag{6.2.3}$$

in the expression of the previous section. ♦

Replicator dynamics with aggregate shocks

Example 6.2.2 (Replicator Dynamics with Aggregate Shocks)**.** Let us consider a finite population of a species, for which each individual is of some genotype taken in the set $\mathcal{A} = \{1, \dots, A\}$. The idea is to model how the genotype distribution of the population evolves along time given their reproductive fitness among others. To do that, let $z_\alpha(t)$ be the *absolute number* of individuals having genotype $\alpha \in \mathcal{A}$ at time t, and define $x_\alpha(t) = z_\alpha(t)/\sum_\beta z_\beta(t)$ the *proportion* of the population having this genotype. We also assume that the *fitness* of a genotype α among others, i.e., the number of offspring per individual, is given by a Lipschitz function $v_\alpha(x)$.

In a natural manner, the evolution of the population can then be characterized by the ordinary differential equation

$$\dot{z}_\alpha = z_\alpha v_\alpha(x), \tag{6.2.4}$$

from which a direct application of the chain rule yields the replicator dynamics of Taylor and Jonker (1978) :

$$\dot{x}_\alpha = x_\alpha \left[v_\alpha(x) - \sum\nolimits_\beta v_\beta(x) x_\beta\right]. \tag{6.2.5}$$

However, when we look at biological reproduction, the fitness is not only determined by how some genotype performs against others, but also by numerous interference of the nature such as weather-like effects. It was then proposed by Fudenberg and Harris (1992) to approximate these interference of nature as random noise perturbing the system (6.2.4), and so considered instead the stochastic biological model

$$dZ_\alpha = Z_\alpha[v_\alpha dt + \sigma_\alpha dW_\alpha]. \tag{6.2.6}$$

Itô's formula can then be used to find the associated stochastic equation verified by $X_\alpha = Z_\alpha / \sum_\beta Z_\beta$, leading to the *replicator dynamics with aggregate shocks* :

$$dX_\alpha = X_\alpha\left[v_\alpha - \sum\nolimits_\beta v_\alpha X_\alpha\right]dt + X_\alpha\left[\sigma_\alpha dW_\alpha - \sum\nolimits_\beta \sigma_\beta X_\beta dW_\beta\right] \tag{6.2.7}$$
$$- X_\alpha\left[\sigma_\alpha^2 X_\alpha - \sum\nolimits_\beta \sigma_\beta^2 X_\beta^2\right]dt. \tag{6.2.8}$$

This stochastic dynamics cannot be recovered from the stochastic imitation dynamics model, however it can be from stochastic exponential dynamics. Indeed, we can take

$$L_\alpha(x) = v_\alpha(x) - \sum\nolimits_\beta x_\beta v_\beta(x) \tag{6.2.9}$$
$$I_\alpha(x) = -\left(\sigma_\alpha^2 x_\alpha - \sum\nolimits_\beta \sigma_\beta^2 x_\beta^2\right) \tag{6.2.10}$$
$$\sigma_\alpha^T(x)dW = \sigma_\alpha dW_\alpha - \sum\nolimits_\beta \sigma_\beta x_\beta dW_\beta. \tag{6.2.11}$$

In particular, we notice that the main difference between this dynamics and (SID) is the presence of a non-null correction term I_α, which is due to the application of Itô's formula to derive the stochastic differential equation verified by X. Or in other words, comes directly from how we have decided to inject the noise into the system, here by considering the influence of nature shocks in genotype fitness. ♦

Stochastic exponential learning

Example 6.2.3 (Stochastic Exponential Learning)**.** Looking from the point of view of learning in games, consider that we are now interested in a population of players that can each play any action in $\mathcal{A}$. Letting $x = (x_\alpha)_{\alpha\in\mathcal{A}}$ be the probability (or frequency) that a player chooses action α, we denote $v_\alpha(x)$ the payoff of players having chosen action α and $y_\alpha(t)$ their cumulative payoffs at time t. If players follow the exponential learning scheme

$$\dot{y_\alpha} = v_\alpha(x)dt \tag{6.2.12}$$
$$x_\alpha = \frac{\exp(y_\alpha)}{\sum_\beta \exp(y_\beta)}, \tag{6.2.13}$$

i.e., they play the *soft best response* to the cumulative payoffs, then an application of the chain rule yields that x also verifies the replicator dynamics (RD).

Accordingly, and to take into account the incertitude in one and others' payoffs, Mertikopoulos and Moustakas (2010) considered the stochastic variant

$$dY_\alpha = v_\alpha(X)dt + \sigma_\alpha(X)dW_\alpha \tag{6.2.14}$$
$$X_\alpha = \frac{\exp(Y_\alpha)}{\sum_\beta \exp(Y_\beta)}. \tag{6.2.15}$$

As in the previous example, we can then apply Itô's formula to obtain an equation on X, yielding the so-called *stochastic exponential learning dynamics* :

$$dX_\alpha = X_\alpha\left[v_\alpha - \sum_\beta v_\alpha X_\alpha\right]dt + X_\alpha\left[\sigma_\alpha dW_\alpha - \sum_\beta \sigma_\beta X_\beta dW_\beta\right] \tag{6.2.16}$$

$$+ \frac{X_\alpha}{2}\left[\sigma_\alpha^2(1-2X_\alpha) - \sum_\beta \sigma_\beta^2 X_\beta(1-2X_\beta)\right]dt. \tag{6.2.17}$$

Here again, such a dynamics can be recovered from the general framework of stochastic exponential dynamics by choosing

$$L_\alpha(x) = v_\alpha(x) - \sum_\beta x_\beta v_\beta(x) \tag{6.2.18}$$

$$I_\alpha(x) = \frac{1}{2}\left[\sigma_\alpha^2(1-2x_\alpha) - \sum_\beta \sigma_\beta^2 x_\beta(1-2x_\beta)\right] \tag{6.2.19}$$

$$\sigma_\alpha^T(x)dW = \sum_\beta x_\beta\left[\sigma_{\beta\alpha}(x)dW_{\beta\alpha} - \sigma_{\alpha\beta}(x)dW_{\alpha\beta}\right]. \tag{6.2.20}$$

♦

Stochastic exponential dynamics therefore constitute a good framework to encapsulate stochastic variants of replicator dynamics and even more complex stochastic dynamics such as the one obtained from imitation procedures. In the next section, we will show that this framework is also great to unify the study of stability in these stochastic game dynamics.

6.3 LONG-RUN BEHAVIOR

Interestingly enough, most results from Chapters 4 and 5 still hold for stochastic dynamics (SED) under an appropriate vector field counterpart to ($\tilde{V}$).

perturbed vector field $\tilde{L}$

Inspired by those previous approaches, let us define $\tilde{L}$ as the vector field given components-wise by

$$\tilde{L}_\alpha(x) = \hat{L}_\alpha(x) - \frac{1}{2}\sigma_\alpha^T(x)C(\alpha,\alpha)\sigma_\alpha(x) \tag{$\tilde{L}$}$$

where $C(\alpha,\beta)$ is the pairwise correlation matrix between vectors W_α and W_β. Accordingly, we say that a strategy $x^* \in \mathcal{X}$ verifies the equilibrium condition (EQ_L^σ) with respect to $\tilde{L}$ if

perturbed equilibrium (EQ_L^σ)

$$\tilde{L}_\alpha(x^*) \geq \tilde{L}_\beta(x^*) \quad \text{for all } \alpha \in \operatorname{supp}(x^*) \text{ and for all } \beta \in \mathcal{A}. \tag{EQ_L^σ}$$

strict perturbed equilibrium (EQ_L^σ)

In particular, x^* is said to verify *strictly* the condition if the inequality in (EQ_L^σ) holds strictly for all $\beta \notin \operatorname{supp}(x^*)$.

Repeating the same analysis as before, we then obtain the following global proposition on the asymptotic behavior of (SED) :

Proposition 6.3.1. *Let $X(t)$ be an interior solution orbit of the stochastic exponential dynamics (SED). Then we have :*

1. *Let $p \in \mathcal{X}$. If there exists $p' \in \mathcal{X}$ such that $\langle \tilde{L}(x), p' - p\rangle > 0$ for all $x \in \mathcal{X}$, then p becomes extinct along $X(t)$.*
2. *If $\mathbb{P}(\lim_{t\to\infty} X(t) = x^*) > 0$, then x^* verifies (EQ_L^σ).*
3. *If x^* is stochastically Lyapunov stable, then it verifies (EQ_L^σ).*
4. *If x^* verifies strictly the equilibrium condition (EQ_L^σ), then it is stochastically asymptotically stable under (SED).*

Proof. The proof is exactly the same as those of Theorems 4.1.2, 5.1.2 and 5.1.5, with the exception that we now use the vector field $\tilde{L}$ instead of $\tilde{v}$. Properties 1. and 4. will also be proved as special cases of Theorem 7.3.2 in the next chapter. □

generalization of results for stochastic replicator dynamics

Remark 6.3.2. Robustness conditions and equilibrium results of Proposition 6.3.1 derived for the special cases mentioned in Examples 6.2.1–6.2.3 above are exactly the same as the ones obtained in the literature, see e.g., Cabrales (2000), Imhof (2005) and Hofbauer and Imhof (2009) for the replicator dynamics with aggregate shocks, Mertikopoulos and Moustakas (2010) and Bravo and Mertikopoulos (2017) for stochastic exponential learning.

Remark 6.3.3. Under the adjusted notations and conditions of Proposition 6.3.1, the rates of extinction and convergence provided in Chapters 4 and 5 also hold for (SED). Their proof are again very similar so we omit them to avoid repetitiveness.

Remark 6.3.4. It is however important to mention that Corollary 5.1.7 do *not* generally holds for (SED). Indeed, this result was strongly linked to the underlying dependence of $\tilde{v}$ and σ on the underlying payoff vectors v, for which we do not make any assumption in the construction of stochastic exponential dynamics. Especially, we do not even need to consider a structure of evolutionary game to study (SED), and as such there is no general reason to talk about "Nash equilibria of the original game".

7

VARIATIONALLY ALIGNED SETS

In this chapter, we introduce and study a new geometrical property of sets that we name *variational alignment*. In particular, we show that this property is strongly linked with dominated strategies and strict Nash equilibria, and that it provides an interesting framework to characterize stability in both deterministic and stochastic exponential dynamics.

7.1 DEFINITION & EXAMPLES

Intuitively, we define a variationally aligned set as a "strongly attracting" set, in the sense that if the dynamics takes a step following a deviation outward of the set, then the mean payoff will always be reduced.

Definition 7.1.1 (Variationally aligned set)**.** Let $v\colon \mathcal{X} \to \mathbb{R}^A$ be a vector field. The set $S \subseteq \mathcal{X}$ is called *variationally aligned* with respect to v if it admits a finite set of deviations $\mathcal{Z} = \{z_1, \ldots, z_r\} \subset \mathbb{R}^A$ such that :

variationally aligned set

i) $\max_{z\in\mathcal{Z}} \langle v(x), z\rangle < 0$ for all x in a neighborhood $\mathcal{U}_a$ of S;

ii) Every interior sequence $\{x^n\}$ of $\mathcal{X}$ verifies $x^n \to S$ whenever $\max_{z\in\mathcal{Z}} \langle y^n, z\rangle \to -\infty$, with $\{y^n\}$ the sequence of $\mathbb{R}^A$ given componentwise by $y^n_\alpha = \log x^n_\alpha$, $\alpha \in \mathcal{A}$.

If we can take $\mathcal{U}_a$ to be all of $\mathcal{X}$, then we will say that S is *globally variationally aligned*; and otherwise we will say that it is only *locally variationally aligned*.

Let $L\colon \mathcal{X} \to \mathbb{R}^A$ be an arbitrary Lipschitz continuous vector field. To keep the terminology consistent throughout this chapter, we will say (by a slight abuse of language) that :

1. a strategy $p \in \mathcal{X}$ is *dominated under L* if there exists $p' \in \mathcal{X}$ such that $\langle L(x), p - p'\rangle < 0$ for all $x \in \mathcal{X}$;

dominated under L

2. a state $x^* \in \mathcal{X}$ is a *strict Nash equilibrium under L* if $L_\alpha(x^*) > L_\beta(x^*)$ for every $\alpha \in \operatorname{supp}(x^*)$ and all $\beta \notin \operatorname{supp}(x^*)$.

strict Nash equilibrium under L

With these notations in hand, we can now provide some important examples of variationally aligned sets, which illustrate why we are that interested such a property.

Example 7.1.2 (Pure dominated strategies)**.** Let $\mathcal{A}^- \subset \mathcal{A}$ be a set of dominated pure strategies under the vector field v, meaning that for all $\alpha \in \mathcal{A}^-$, there exists $\beta_\alpha \in \mathcal{A}$ such that $v_\alpha(x) < v_{\beta_\alpha}(x)$ on $\mathcal{X}$. Let S be the face of $\mathcal{X}$ which does not contain any pure strategy of $\mathcal{A}^-$. Then S is globally variationally aligned.

pure dominated strategies are variationally stable

Indeed, let us construct the finite set $\mathcal{Z} = \{z^\alpha\}_{\alpha\in\mathcal{A}^-}$ whose elements are given by $z^\alpha = e_\alpha - e_{\beta_\alpha}$[1]. For such a set, we have $\langle v(x), z^\alpha\rangle = v_\alpha(x) - v_{\beta_\alpha}(x) < 0$ for all $z^\alpha \in \mathcal{Z}$ and all $x \in \mathcal{X}$ by domination assumption, which verifies the first point of Definition 7.1.1 with $\mathcal{U}_a = \mathcal{X}$. On the other hand, $\langle z^\alpha, y^n\rangle = y^n_\alpha - y^n_{\beta_\alpha} > \log(x^n_\alpha)$, meaning that $x^n_\alpha \to 0$ for every $\alpha \in \mathcal{A}^-$ whenever $\langle z^\alpha, y^n\rangle \to -\infty$ for all $z^\alpha \in \mathcal{Z}$. Therefore the second point of the definition also holds true, hence proving that S is globally variationally aligned with respect to v. ♦

mixed dominated strategies are variationally stable

Example 7.1.3 (Mixed dominated strategies)**.** Let $\mathcal{X}^- \subset \mathcal{X}$ be a finite set of dominated mixed strategies under the vector field v, meaning that for all $p \in \mathcal{X}^-$, there exists $q^p \in \mathcal{X}$ such that $\langle v(x), p - q^p\rangle < 0$ on $\mathcal{X}$. Let S be the union of faces of $\mathcal{X}$ which do not contain any mixed strategy of $\mathcal{X}^-$. Then S is globally variationally aligned.

Indeed, let us construct the finite set $\mathcal{Z} = \{z^p\}_{p\in\mathcal{X}^-}$ whose elements are given by $z^p = p - q^p$. For each $p \in \mathcal{X}$ and $x \in \mathcal{X}$, the Kullback-Lieber divergence between p and x (seen as discrete probability measures) is given by $\mathrm{KL}(p, x) = \sum_\alpha p_\alpha \log \frac{p_\alpha}{x_\alpha}$. For $p \in \mathcal{X}^-$, let S_p be the union of faces that do not contain p. By classical arguments (see Weibull, 1995), we notice that $\mathrm{KL}(p, x) < \infty$ if and only if $x_\alpha > 0$ for all $\alpha \in \mathrm{supp}(p)$, and so $x^n \to S_p$ whenever $\mathrm{KL}(p, x^n) \to \infty$. Moreover, we can decompose this divergence as

$$\mathrm{KL}(p, x) = \mathrm{KL}(q^p, x) + \sum\nolimits_\alpha \left(p_\alpha \log p_\alpha - q^p_\alpha \log q^p_\alpha\right) - \langle z^p, y\rangle, \tag{7.1.1}$$

so that $\mathrm{KL}(p, x^n) \to \infty$ if and only if $\langle z^p, y^n\rangle \to -\infty$ (notice that $\mathrm{KL}(q^p, x) \geq 0$ by usual properties of the Kullback-Leibler divergence). Accordingly, $x^n \to S = \bigcap_{p\in\mathcal{X}^-} S_p$ whenever $\langle z^p, y^n\rangle \to -\infty$ for every $z^p \in \mathcal{Z}$, i.e., $\mathcal{Z}$ verifies the second condition of Definition 7.1.1. As the first condition also holds immediately for $\mathcal{U}_a = \mathcal{X}$ by the domination condition, we have hence proven that S is globally variationally aligned with respect to v. ♦

strict Nash equilibria are variationally stable

Example 7.1.4 (Strict Nash equilibrium)**.** Let x^* be a strict Nash equilibrium under the vector field v. Then $S = \{x^*\}$ is locally variationally aligned.

As x^* is strict, it is in particular a pure strategy, and so we introduce $\alpha^* \in \mathcal{A}$ the pure action such that $x^* = e_{\alpha^*}$. We then construct the finite set $\mathcal{Z} = \{z^\alpha\}_{\alpha\neq\alpha^*}$ whose elements are given by $z^\alpha = e_\alpha - e_{\alpha^*}$. For each $z^\alpha \in \mathcal{Z}$, we have $\langle z^\alpha, y^n\rangle = y^n_\alpha - y^n_{\alpha^*} \geq \log x^n_\alpha$, so that $x^n_\alpha \to 0$ whenever $\langle z^\alpha, y^n\rangle \to -\infty$. Accordingly, $\mathcal{Z}$ verifies the second condition of Definition 7.1.1. Furthermore, by definition of the strict Nash equilibrium and continuity of the payoff vector v, there exists a neighborhood $\mathcal{U}_a$ of x^* such that $\langle z^\alpha, v(x)\rangle = v_\alpha(x) - v_{\alpha^*}(x) < 0$ for all $\alpha \neq \alpha^*$ and all $x \in \mathcal{U}_a$, hence proving that S is locally variationally aligned. ♦

closed under better-replies sets are variationally stable

Example 7.1.5 (Closed under better-replies)**.** For this example, let us go back to the multipopulation setup of evolutionary games, and assume that individuals are interacting through asymmetric random matching. This means that the payoff functions are given by

$$v_{k\alpha}(x) = \sum\nolimits^1_{\alpha_1} \ldots \sum\nolimits^N_{\alpha_N} u_k(\alpha_1, \ldots, \alpha, \ldots, \alpha_N) x_{1\alpha_1} \ldots \delta_{\alpha_k\alpha} \ldots x_{N\alpha_N}, \tag{7.1.2}$$

normal-form game

better-reply correspondence γ_k

where $u_k \colon \mathcal{A} \to \mathbb{R}$ are the payoffs of pure strategy combinations. In this case, the game $\Gamma(\mathcal{N}, \mathcal{A}, v)$ is called a *normal-form game* . Inspired by the study of stability on normal-form games performed by Ritzberger and Weibull (1995), we then introduce for each population $k \in \mathcal{N}$ the *better-reply* correspondence

1 Here, e_α should be understand as the unit vector in $\mathbb{R}^{\mathcal{A}}$ with value 1 at coordinate α and 0 for all other components.

$\gamma_k : \mathcal{X} \to \mathcal{A}_k$ defined by

$$\gamma_k(x) = \left\{\beta \in \mathcal{A}_k : \left\langle v_k(e_\beta; x_{-k}), e_\beta \right\rangle \geq \langle v_k(x), x_k \rangle \right\} \tag{7.1.3}$$
$$= \left\{\beta \in \mathcal{A}_k : \left\langle v_k(x), e_\beta - x_k \right\rangle \geq 0 \right\} \tag{7.1.4}$$

for all $x \in \mathcal{X}$. In other words, this correspondence yields the pure actions of population k that achieve a payoff at least as good as the one of the mixed strategy profile x. If we let S be a face[2] of the polyhedron $\mathcal{X}$ that is closed by $\gamma = (\gamma_k)_{k \in \mathcal{N}}$ (we say that S is *closed under better-replies*), then it is locally variationally aligned.

closed under better-replies

To pursue, notice that the fact that S is closed under better-replies implies $\left\langle v_k(x), e_\beta - x_k \right\rangle < 0$ for every $\beta \notin S_k$[3], $x \in S$ and $k \in \mathcal{N}$. In particular, let us fix some pure strategies $\alpha_k^* \in S_k$ and $\beta \notin S_k$, and take any $\alpha_{-k} \in \prod_{i \neq k} S_j$, so that $x^* = \left(\alpha_k^*; \alpha_{-k}\right)$ is a vertex of the face S. For this specific element of S, we then have $v_{k\beta}\left(\alpha_k^*; \alpha_{-k}\right) < v_{k\alpha_k^*}\left(\alpha_k^*; \alpha_{-k}\right)$, which boils down to $u_k(\beta; \alpha_{-k}) < u_k\left(\alpha_k^*; \alpha_{-k}\right)$ for any $\alpha_{-k} \in \prod_{i \neq k} S_j$. Now, pick some $x \in S$. As S is a face, we immediately have $\mathrm{supp}(x) \subset S$, so that $u_k(\beta; \alpha_{-k}) < u_k\left(\alpha_k^*; \alpha_{-k}\right)$ for all $\alpha_k = (\alpha_k; \alpha_{-k}) \in \mathrm{supp}(x)$. This inequality is true for any pure actions in the support of x, which combined with the specific form (7.1.2) of the payoff function v readily yields $v_{k\beta}(x) < v_{k\alpha_k^*}(x)$ for every $\beta \notin S_k$. Letting $\mathcal{Z}_k = \left\{e_\beta - e_{\alpha_k^*} : \beta \notin S_k\right\}$, we therefore have $\langle v_k(x), z_k \rangle < 0$ for every $x \in S$, $z_k \in \mathcal{Z}_k$ and $k \in \mathcal{N}$. By continuity of v and compactness of $\mathcal{X}$, it means that S verifies the first condition of Definition 7.1.1 (adapted to a multipopulation setting) in a local neighborhood of S. The second condition is also easily shown using the same idea as in Example 7.1.2, hence proving that S is locally variationally aligned. ♦

As was shown in the previous examples, variationally aligned sets seem to constitute a good generalization of what we were already studying in the previous chapters of the thesis, i.e., dominated strategies and strict Nash equilibria, and even contains more exotic sets such as those closed under better-replies.

In the next two sections, we will now study the stability of such sets, both in deterministic and in stochastic regimes.

7.2 DETERMINISTIC STABILITY

Before delving into the study of variationally stable sets, we first define what we mean by *stability* of sets in deterministic dynamical systems.

Definition 7.2.1 (Deterministic stability)**.** Let $S \subseteq \mathcal{X}$ and let $x(t)$ be a solution orbit of (ED). We say that

1. S is *Lyapunov stable* under (ED) if, for every neighborhood $\mathcal{U}_0$ of S in $\mathcal{X}$, there exists an even smaller neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of S such that $x(t) \in \mathcal{U}_0$ for all $t \geq 0$ whenever $x(0) \in \mathcal{U}$.

Lyapunov stable set

2. S is *asymptotically stable* under (ED) if it is Lyapunov stable and attracting : for every neighborhood $\mathcal{U}_0$ of S in $\mathcal{X}$, there exists an even smaller neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of S such that $x(t) \in \mathcal{U}_0$ for all $t \geq 0$ and $x(t) \to S$ whenever $x(0) \in \mathcal{U}$.

asymptotically stable set

2 Here it should be understand as $S = \prod_k S_k$ where, for each $k \in \mathcal{N}$, S_k is a face of the simplex $\mathcal{X}_k$.

3 In the sense that $\beta \in \mathcal{A}_k$ but $e_\beta \notin S_k$.

globally attracting set

3. S is *globally attracting* under (ED) if $x(t) \to S$ for any interior initial condition.

In the following theorem, we state the main result in deterministic setting, accordingly that variationally aligned sets are always stable under (ED), and that the local (resp. global) aspects are preserved.

variationally stable sets are stable

Theorem 7.2.2. *Suppose $S \subseteq \mathcal{X}$ is variationally aligned with respect to the vector field L. Then we have :*

1. *If S is globally variationally aligned, then it is globally attracting under* (ED)*.*
2. *If S is locally variationally aligned, then it is asymptotically stable under* (ED)*.*

In particular, we observe that the conditions and results of Theorem 7.2.2 do not assume any underlying structure of a population game, meaning that it provides a general and geometrical aspect to the characterization of stability in any dynamics on the simplex of the form (ED).

Before proving Theorem 7.2.2, we first propose important special cases based on the examples of the previous section.

Corollary 7.2.3. *Dominated strategies under L become extinct along any interior solution orbit of* (ED)*.*

Proof. Let $\mathcal{X}^-$ be a finite set of dominated strategies under L, and let S be the union of faces of $\mathcal{X}$ that do not contain any strategy in $\mathcal{X}^-$. From Example 7.1.3, we know that S is globally variationally aligned with respect to L, and so Theorem 7.2.2 states that S is globally attracting under (ED). In particular it means that any interior solution orbit $x(t)$ of (ED) verifies $x(t) \to S$, hence proving the extinction of dominated strategies. □

Corollary 7.2.4. *Strict Nash equilibria under L are asymptotically stable with respect to* (ED)*.*

Proof. From Example 7.1.4 we know that a strict Nash equilibrium is locally variationally aligned with respect to L, and so is asymptotically stable due to Theorem 7.2.2. □

monotone vector field

Let $\Gamma(\mathcal{N}, \mathcal{A}, \tilde{v})$ be a population game and let $L\colon \mathcal{X} \to \mathbb{R}^A$ be a vector field. We say that L is *monotone* with respect to v if for every $x \in \mathcal{X}$,

$$v_\beta(x) > v_\alpha(x) \Longrightarrow L_\beta(x) > L_\alpha(x) \quad \text{for all } \alpha, \beta \in \mathcal{A}. \tag{7.2.1}$$

An important example of a monotone vector field is the one given componentwise by $L_\alpha(x) = \sum_\beta x_\beta \left[r_{\beta\alpha} - r_{\alpha\beta} \right]$ where $r_{\alpha\beta}$ is a monotone conditional imitation rate (cf. Chapter 3).

Using this monotone condition, we can then propose a stability result for closed under better-replies sets in normal-form games, as was already shown for similar dynamics by Ritzberger and Weibull (1995).

Corollary 7.2.5. *Let $\Gamma(\mathcal{N}, \mathcal{A}, v)$ be a normal-form game and let S be a face of $\mathcal{X}$ closed under better-replies. Then S is asymptotically stable under* (ED) *whenever L is monotone with respect to v.*

Proof. Following Example 7.1.5, S is locally variationally stable under v with deviation set $\mathcal{Z}_k = \left\{ e_\beta - e_{\alpha_k^*} : \beta \notin S_k \right\}$ and some alignment neighborhood $\mathcal{U}_a$.

We will show that under the conditions of the corollary, S is in fact also locally variationally aligned with respect to L.

Indeed, the second point of Definition 7.1.1 is immediately verified as it does not depend on the underlying vector field. Furthermore, by alignment and monotony,

$$\left\langle L_k(x), e_\beta - e_{\alpha_k^*} \right\rangle = L_{k\beta}(x) - L_{k\alpha_k^*}(x) < 0 \tag{7.2.2}$$

for any $x \in \mathcal{U}_a$ and $e_\beta - e_{\alpha_k^*} \in \mathcal{Z}_k$. This proves that S is locally variationally aligned with respect to L, and so by Theorem 7.2.2 we deduce that it is asymptotically stable under (ED). □

Remark 7.2.6. Using the same reasoning as in the proof of Corollary 7.2.5, we can show that the conclusion of Corollary 7.2.3 (resp. of Corollary 7.2.4) also holds for dominated strategies (resp. for strict Nash equilibria) of a population game $\Gamma(\mathcal{N}, \mathcal{A}, v)$ whenever L is monotone with respect to the payoff vector v. In particular, it provides a alternate proof of the extinction of dominated strategies and stability of strict Nash equilibria in monotone imitation dynamics.

Proof of Theorem 7.2.2. Let $x(t)$ be an interior solution orbit of (ED) and let us define the dual variable $y(t)$ on $\mathbb{R}^A$ given by $y_\alpha(t) = \log x_\alpha(t)$ for all $\alpha \in \mathcal{A}$.

If globally variationally aligned : Let $\mathcal{Z}$ be a finite set of deviations verifying the conditions of Definition 7.1.1 for S with $\mathcal{U}_a = \mathcal{X}$. According to the variationnal alignment of S, it is therefore sufficient to show that $v_z(t) \to -\infty$ for all $z \in \mathcal{Z}$, where $v_z(t) = \langle z, y(t) \rangle$. However, we have

$$v_z'(t) = \langle z, \dot{y} \rangle = \langle z, L(x) \rangle \leq -m \tag{7.2.3}$$

where $m = -\sup_{x \in \mathcal{X}} \langle L(x), z \rangle > 0$ by assumption and compactness of $\mathcal{X}$. Integrating on both sides it yields $v_z(t) \leq v_z(0) - mt$ whose right-hand side tends to $-\infty$ for every $z \in \mathcal{Z}$, hence proving the global attractiveness of S.

If locally variationally aligned : Let $\mathcal{Z}$ be a finite set of deviations verifying the conditions of Definition 7.1.1 for S and let $\mathcal{U}_a$ be an associated alignment neighborhood. As before, let $v_z(t) = \langle z, y(t) \rangle$.

Let $\mathcal{U}_0$ be a neighborhood of S in $\mathcal{X}$. From the variational alignment of S, there exists a neighborhood[4] $\mathcal{U} \subseteq \mathcal{U}_0$ of S and a constant $m > 0$ such that $\langle z, L(x) \rangle \leq -m$ for all $x \in \mathcal{U}$ and all $z \in \mathcal{Z}$. Furthermore, the continuity of v_z and the second condition of Definition 7.1.1 implies that we can pick some $M > 0$ big enough so that $x(t) \in \mathcal{U}$ whenever $v_z(t) \leq -m$ for all $z \in \mathcal{Z}$.

The main idea behind this proof is to show that if M is well-chosen and if the initial conditions verify $v_z(0) \leq -M$, then $x(t) \in \mathcal{U} \subseteq \mathcal{U}_0$ for all $t \geq 0$ and $x(t) \to S$.

Accordingly, assume that $v_z(0) \leq -M$ and let $T = \inf\{t \geq 0 : x(t) \notin \mathcal{U}\}$ be the first exist time of x from $\mathcal{U}$. If $T < \infty$, we can write $v_z(T) \leq v_z(0) - mT \leq -M$, and so $x(T) \in \mathcal{U}$. This leads to a contradiction as $x(T) \notin \mathcal{U}$ by construction, hence showing that $T = \infty$. In particular, it means that $x(t) \in \mathcal{U}$ for all $t \geq 0$ whenever $v_z(0) \leq -M$, hence proving the Lyapunov stability of S.

Proceeding as in the first part of the proof, we can then show that $v_z(t) \leq v_z(0) - mt$ for all $t \geq 0$ whenever $v_z(0) \leq -M$, from which we readily deduce the attracting property of S. □

4 Take for instance $\mathcal{U} \subset \mathcal{U}_a \cap \mathcal{U}_0$ small enough.

7.3 STOCHASTIC STABILITY

In the previous section, we have seen that variational alignment provides a great tool to study stability in (ED), but is also the case when we consider general random perturbations such as the ones in (SED) ?

In Chapter 5, we have already talked about stochastic stability but only for a unique state. In the following definition we extend this notion to the stochastic stability of a whole set, mirroring the definition that we have stated for the deterministic setup.

Definition 7.3.1 (Stochastic stability)**.** Let $S \subseteq \mathcal{X}$ and let $X(t)$ be a solution orbit of (SED). We say that

stochastically stable set

1. S is *stochastically (Lyapunov) stable* under (SED) if, for every $\varepsilon > 0$ and for every neighborhood $\mathcal{U}_0$ of S in $\mathcal{X}$, there exists an even smaller neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of S such that

$$\mathbb{P}(X(t) \in \mathcal{U}_0 \text{ for all } t \geq 0) \geq 1 - \varepsilon \tag{7.3.1}$$

whenever $X(0) \in \mathcal{U}$.

stochastically asymptotically stable set

2. S is *stochastically asymptotically stable* under (SED) if it is stochastically stable and attracting : for every $\varepsilon > 0$ and for every neighborhood $\mathcal{U}_0$ of S in $\mathcal{X}$, there exists an even smaller neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of S such that

$$\mathbb{P}\left(X(t) \in \mathcal{U}_0 \text{ for all } t \geq 0 \text{ and } \lim_{t\to\infty} X(t) \in S\right) \geq 1 - \varepsilon \tag{7.3.2}$$

whenever $X(0) \in \mathcal{U}$.

globally stochastically attracting set

3. S is *globally stochastically attracting* under (SED) if $X(t) \to S$ (a.s.) for any interior initial condition.

Interestingly, the stability results of Theorem 7.2.2 can be carried over to the stochastic setup almost identically, under the important exception that we are now asking for variational stability under the perturbed vector field $\tilde{L}$ instead of only L.

variationally stable sets are stochastically stable

Theorem 7.3.2. *Suppose $S \subseteq \mathcal{X}$ is variationally aligned with respect to the vector field $\tilde{L}$ defined by* ($\tilde{L}$)*. Then we have :*

1. *If S is globally variationally aligned, then it is globally stochastically attracting under* (SED)*.*

2. *If S is locally variationally aligned, then it is stochastically asymptotically stable under* (SED)*.*

As for the deterministic case, we first provide special cases of Theorem 7.3.2 before proving it. In particular, we show that all of Corollaries 7.2.3–7.2.5 can be extended in the stochastic setting, changing only L by its perturbed version $\tilde{L}$.

Corollary 7.3.3.
1. *Dominated strategies under $\tilde{L}$ become extinct along any interior solution orbit of* (SED)*.*

2. *Strict Nash equilibria under $\tilde{L}$ are asymptotically stochastically stable with respect to* (SED)*.*

3. *Let $\Gamma(\mathcal{N}, \mathcal{A}, v)$ be a normal-form game and let S be a face of $\mathcal{X}$ closed under better-replies. Then S is stochastically asymptotically stable under* (SED) *whenever $\tilde{L}$ is monotone with respect to v.*

Proof. Follows the same reasoning as the proofs of Corollaries 7.2.3–7.2.5, with the exception of using Theorem 7.3.2 instead of Theorem 7.2.2 as we are now in a stochastic setup. □

Remark 7.3.4. Noticing the equivalence of the conditions, results 1. and 2. of Corollary 7.3.3 provide a general proof for the stability conclusions 1. and 4. stated in Proposition 6.3.1.

Remark 7.3.5. Corollary 7.3.3 also gives an (implicit) noise condition so that closed under better-replies sets remain stable in a stochastic regime, accordingly that $\tilde{L}$ needs to stay monotone with respect to v. In particular, this require the Itô's correction term and the noise to be either small enough or structured enough (e.g., also monotone with respect to v) to preserve monotony.

Proof of Theorem 7.3.2. Let $X(t)$ be an interior solution orbit of (SED) and let us define the dual stochastic process $Y(t)$ on $\mathbb{R}^A$ given by $Y_\alpha(t) = \log X_\alpha(t)$ for all $\alpha \in \mathcal{A}$.

If globally variationally aligned : Let $\mathcal{Z}$ be a finite set of deviations verifying the conditions of Definition 7.1.1 for S with $\mathcal{U}_a = \mathcal{X}$. According to the variationnal alignment of S, it is therefore sufficient to show that $V_z(t) \to -\infty$ (a.s.) for all $z \in \mathcal{Z}$, where $V_z(t) = \langle z, Y(t)\rangle$. A similar use of Itô's formula as in the proof of Theorem 4.1.2 leads to the stochastic differential form

$$dY_\alpha = \left(\hat{L}_\alpha - \frac{1}{2}\sigma_\alpha^T C(\alpha,\alpha)\sigma_\alpha\right)dt + \sigma_\alpha^T dW_\alpha = \tilde{L}_\alpha(X)dt + \sigma_\alpha^T(X)dW_\alpha, \tag{7.3.3}$$

with $\tilde{L}$ given by ($\tilde{L}$). As such, we can write

$$dV_z(t) = \langle z, dY(t)\rangle = \left\langle z, \tilde{L}(X)\right\rangle dt - d\xi_z \leq -m dt - d\xi_z \tag{7.3.4}$$

where $\xi_z(t) = -\sum_\alpha \int_0^t z_\alpha \sigma_\alpha^T(X(s))\, dW_\alpha(s)$ and $m = -\sup\limits_{x\in\mathcal{X}}\left\langle \tilde{L}(x), z\right\rangle > 0$ by assumption. In stochastic integral form, this yields

$$V_z(t) \leq V_z(0) - mt - \xi_z(t). \tag{7.3.5}$$

The right-hand side of (7.3.5) can be expressed as in Corollary 2.3.4, which gives that $V_z(t) \sim -mt \to -\infty$ (a.s.) for all $z \in \mathcal{Z}$, hence proving the first statement.

If locally variationally aligned : As the first point, let $\mathcal{Z}$ be a finite set verifying the conditions of Definition 7.1.1 for S and let $\mathcal{U}_a$ be an associated alignment neighborhood. Furthermore, let us denote V_z the process defined by $V_z(t) = \langle z, Y(t)\rangle$.

To proceed, fix a threshold $\varepsilon > 0$ and let $\mathcal{U}_0$ be a neighborhood of S in $\mathcal{X}$. As in the proof of Theorem 7.2.2, there exists a neighborhood $\mathcal{U} \subseteq \mathcal{U}_0$ of S and a constant $m > 0$ such that $\left\langle z, \tilde{L}(x)\right\rangle \leq -m < 0$ for all $x \in \mathcal{U}$ and $z \in \mathcal{Z}$. Furthermore, we can pick some $M > 0$ big enough so that $X(t) \in \mathcal{U}$ whenever $V_z(t) \leq -M$ for all $z \in \mathcal{Z}$.

Analogously to the deterministic case, we will show that if M is well-chosen and if $V_z(0) \leq -2M$, then $X(t) \in \mathcal{U} \subseteq \mathcal{U}_0$ for all $t \geq 0$ and $X(t) \to S$ with probability at least $1 - \varepsilon$.

Assume that $V_z(t) \leq -2M$ for all $z \in \mathcal{Z}$, and let $\tau_{\mathcal{U}} = \inf\{t \geq 0 : X(t) \notin \mathcal{U}\}$ be the first exist time of $X(t)$ from $\mathcal{U}$. On the event $\{\tau_{\mathcal{U}} = \infty\}$ the process $X(t)$ will always remain in $\mathcal{U}$, and so we notice that the same reasoning as for the first part of the proof can be carried on to prove that $X(t) \to S$ (a.s.) conditionally on $\{\tau_{\mathcal{U}} = \infty\}$[5]. As such, to show the asymptotic stochastic stability of S, it is sufficient to prove that the event $\{\tau_{\mathcal{U}} = \infty\}$ occurs with probability greater than $1 - \varepsilon$.

To do so, we first notice that the time-change theorem for martingales states that the process $\xi_z(t)$ can be written $\widetilde{W}_z(\rho_z(t))$ where $\widetilde{W}_z$ is a standard Brownian motion (on a possible enlarged probability space) and ρ_z is the quadratic variation of ξ_z. By compactness of $\mathcal{X}$ and Lipschitz continuity of the integrand, there also exists a deterministic constant $0 < K < \infty$ which does not depend on z and such that $\rho_z(t) \leq Kt$ for all $t \geq 0$ and all $z \in \mathcal{Z}$. Let us define the random time $\tau_0 = \inf\left\{t \geq 0 : \inf_{z \in \mathcal{Z}} \widetilde{W}_z(t) \leq -M - \frac{mt}{K}\right\}$. We first show that if $\tau_0 = \infty$, then we also have $\tau_{\mathcal{U}} = \infty$.

For the sole purpose of proving a contradiction, assume on the other hand that we are on the event $\{\tau_0 = \infty\} \cap \{\tau_{\mathcal{U}} < \infty\}$. From the expression (7.3.5) of V_z, we then get

$$V_z(\tau_{\mathcal{U}}) \leq -2M - m\tau_{\mathcal{U}} - \widetilde{W}_z(\rho_z(\tau_{\mathcal{U}})), \tag{7.3.6}$$

which is indeed well-defined because $\tau_{\mathcal{U}} < \infty$ and $X(t) \in \mathcal{U}$ for all $t < \tau_{\mathcal{U}}$. However, on $\{\tau_0 = \infty\}$ we have $\widetilde{W}_z(t) > -M - \frac{mt}{K}$ for all $z \in \mathcal{Z}$ and all $t \geq 0$, so in particular

$$\widetilde{W}_z(\rho_z(\tau_{\mathcal{U}})) > -M - \frac{m\rho_z(\tau_{\mathcal{U}})}{K} \geq -M - m\tau_{\mathcal{U}} \tag{7.3.7}$$

where the last inequality comes from $\rho_z(t) \leq Kt$. Combined with (7.3.6), it yields $V_z(\tau_{\mathcal{U}}) \leq -M$ for all $z \in \mathcal{Z}$ and so $X(\tau_{\mathcal{U}}) \in \mathcal{U}$. This rises a contradiction by definition of the exit time $\tau_{\mathcal{U}}$, hence proving that $\tau_{\mathcal{U}} = \infty$ whenever $\tau_0 = \infty$.

Taking the probability of both events, it then holds that $\mathbb{P}(\tau_{\mathcal{U}} = \infty) \geq \mathbb{P}(\tau_0 = \infty)$. But the event $\{\tau_0 < \infty\}$ can be expressed as $\{\tau_0 < \infty\} = \bigcup_{z \in \mathcal{Z}} \{\tau_z < \infty\}$ where $\tau_z = \inf\left\{t \geq 0 : \widetilde{W}_z(t) + \frac{mt}{K} = -M\right\}$. As τ_z is an hitting time of a Brownian motion with positive drift, it enjoys the bound $\mathbb{P}(\tau_z < \infty) = \exp\left\{\frac{-2mM}{K}\right\}$ (cf. Karatzas and Shreve, 1998). Therefore, if we take M big enough so that $M > -\frac{K \log \varepsilon / |\mathcal{Z}|}{2m}$, then $\mathbb{P}(\tau_0 < \infty) \leq \sum_{z \in \mathcal{Z}} \mathbb{P}(\tau_z < \infty) < \varepsilon$, from which we finally deduce $\mathbb{P}(\tau_{\mathcal{U}} = \infty) \geq \mathbb{P}(\tau_0 = \infty) \geq 1 - \varepsilon$. By the previous discussion, we have then proven that the set S is stochastically asymptotically stable as required. □

---

5 In fact, S can be seen as globally variationally aligned when we restrain the whole space $\mathcal{X}$ to only $\mathcal{U}$, and so a slight generalization of the first result of Theorem 7.3.2 yields that S is globally stochastically attracting conditioned on the event that the trajectory stays in $\mathcal{U}$ for all t.

8

PERSPECTIVES

To conclude this manuscript, we describe some research perspectives around robustness of game dynamics envisaged for future works.

8.1 VARIATIONALLY ALIGNED & EVOLUTIONARILY STABLE

Continuing in the analysis of variationally aligned sets in deterministic and stochastic exponential dynamics, Ritzberger and Weibull (1995) have shown that in the particular case of deterministic *sign-preserving dynamics* in a normal-form game, a face is asymptotically stable if and only if it is closed under better-replies.

sign-preserving dynamics

In Chapter 7, we have proved that closed under better-replies sets (in normal-form games) are also stochastically asymptotically stable under some monotony condition on the perturbed vector field $\tilde{L}$, but is the other way around also true ? In other words, do all stochastically asymtotically stable faces are in fact closed under better-replies when the noise is mild enough ?

Another interrogation about variationally aligned sets is their links with what are called *evolutionarily stable sets* (see Thomas, 1985 and Balkenborg and Schlag, 2001). This class of sets is included in the class of Nash equilibria, and are generally a good way to characterize asymptotically stable states of a game, more so than strict Nash equilibria. However, in general games, evolutionarily stable sets are neither included, nor include, variationally aligned sets. But we can conjecture that any variationally aligned face contains an evolutionarily set, which would therefore be the subset that attracts strongly the trajectory inside the face.

evolutionarily stable sets

An important direction would then be to study how the convergence and stability occur inside a variationally aligned face, both in a deterministic and a stochastic setting, to infer more local properties than just the stability of the whole face.

8.2 ABRUPT SHOCKS

In what we have done up to now, we always assumed that the noise is characterized as some continuous stochastic process, for instance as a stochastic integral with respect to correlated Brownian motions.

Such a model is well suited when the interactions between individuals are subject to overall shocks of their underlying environment, such as the weather or telecommunication noise. However, we may also want to study the impact of *abrupt* events that cause an immediate shift to the population states. Such rare events are common in nature, for instance we can think of earthquakes, thunder, or any other natural disasters.

In these case, we may want to consider more general types of noise that are not necessarily continuous everywhere, i.e., that may contain some jumps at random times. Theoretically, this can be done by generalizing the stochastic analysis setup of Chapter 2 to arbitrary semimartingales. For instance, stochastic integrals may be taken with respect to a *Lévy process* (see Applebaum, 2009), which can be decomposed as a sum between a Brownian motion, a compound Poisson process and a compensated Poisson process.

Lévy process

This kind of study of abrupt shocks in game dynamics was already explored by Vlasic (2018) for the biological replicator dynamics in a multipopulation symmetric game, who showed the extinction of dominated strategies and the stability of strict Nash equilibria under some conditions on the jumps amplitude.

We conjecture that if the jumps are small enough and occur rarely often, then our long-run results for stochastic exponential dynamics remain true under a discontinuous noise, as the trajectories would stay in the same neighborhood with high probability. However, most of our proofs would need to be carefully adapted, as the usual use of the time-change theorem can only be done on a *continuous* local martingale.

8.3 STOCHASTIC MIRROR DYNAMICS

Going back to the stochastic exponential learning dynamics of Mertikopoulos and Moustakas (2010) defined in Example 6.2.3, it can in fact be seen as a particular case of another general stochastic dynamics introduced by Bravo and Mertikopoulos (2017) and called the *stochastic regularized learning dynamics*. These dynamics are of the form :

stochastic regularized learning

$$dY_\alpha = v_\alpha(X)dt + \sigma_\alpha(X)dW_\alpha \tag{8.3.1}$$

$$X_\alpha(t) = Q(\eta(t)Y(t)), \tag{8.3.2}$$

where Q is a *mirror map* or *regularized best response* from $\mathbb{R}^{\mathcal{A}}$ to $\mathcal{X}$ (see e.g., Giannou et al., 2021 for a more exhaustive geometric treatment of such maps in discrete time).

mirror map Q

In particular, our rates for both extinction of mixed strategies and convergence to strict Nash equilibria can be extended almost seamlessly to cover general regularizers, but we do not delve into this topic to avoid having to define many technical tools of regularized optimization.

These dynamics have a strong link to *mirror descent dynamics* that are well-know in discrete time optimization (see Nesterov, 2009), and that can be defined for any primal space $\mathcal{X}$, not only for the simplex on $\mathcal{A}$.

mirror descent dynamics

A stochastic variant of continuous-time mirror dynamics has also been studied by Mertikopoulos and Staudigl (2017) in the context of monotone variational inequalities with $\mathcal{X}$ being a compact convex set. In particular, they established global convergence in the ergodic sense and large deviations principles for the concentration around the average.

However, we can wonder if our local study of equilibria stability and extinction of dominated strategies can also be carried out to these stochastic dynamics in general primal spaces, such as convex polytopes, compact convex sets or even Riemannian manifolds.

The case of convex polytopes should behave similarly to the simplex as the faces have the same geometric properties (e.g., are flat with sharp angles at the vertices, so should become strong traps for the trajectories). However, the extension for smooth spaces will surely need some new tools to study stochastic stability of equilibria, such as a more thorough exploration of the local geometry.

BIBLIOGRAPHY

[1] Ethan Akin. Domination or equilibrium. *Mathematical Biosciences*, 50(3-4):239–250, 1980.

[2] David Applebaum. *Lévy Processes and Stochastic Calculus*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2 edition, 2009.

[3] Patrick Bajari, Han Hong, and Denis Nekipelov. *Game Theory and Econometrics: A Survey of Some Recent Research*, volume 3 of *Econometric Society Monographs*, page 3–52. Cambridge University Press, 2013.

[4] Paolo Baldi. *Stochastic Calculus: An Introduction Through Theory and Exercises*. Universitext. Springer International Publishing, 2017.

[5] Dieter Balkenborg and Karl H. Schlag. Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595, 2001.

[6] Jonas Björnerstedt and Jörgen W. Weibull. Nash equilibrium and evolution by imitation. In Kenneth Joseph Arrow, E. Colombatto, M. Perlman, and C. Schmidt, editors, *The Rational Foundations of Economic Behavior*, pages 155–181. St. Martin's Press, New York, NY, 1996.

[7] Daan Bloembergen, Karl Tuyls, Daniel Hennes, and Michael Kaisers. Evolutionary dynamics of multi-agent learning: A survey. *J. Artif. Int. Res.*, 53(1):659–697, may 2015.

[8] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[9] Mario Bravo and Panayotis Mertikopoulos. On the robustness of learning in games with stochastically perturbed payoff observations. *Games and Economic Behavior*, 103(John Nash Memorial issue):41–66, May 2017.

[10] Antonio Cabrales. Stochastic replicator dynamics. *International Economic Review*, 41 (2):451–81, May 2000.

[11] Erhan Çınlar. *Probability and Stochastics*. Graduate Texts in Mathematics. Springer New York, 2011.

[12] Drew Fudenberg and Christopher Harris. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57(2):420–441, August 1992.

[13] Angeliki Giannou, Emmanouil Vasileios Vlatakis-Gkaragkounis, and Panayotis Mertikopoulos. The convergence rate of regularized learning in games: From bandits and uncertainty to optimism and beyond. In *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.

[14] Tanmoy Hazra and Kushal Anjaria. Applications of game theory in deep learning: a survey. *Multimedia Tools and Applications*, 81, 02 2022.

[15] Josef Hofbauer and Lorens A. Imhof. Time averages, recurrence and transience in the stochastic replicator dynamics. *The Annals of Applied Probability*, 19(4):1347–1368, 2009.

[16] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, July 2003.

[17] Josef Hofbauer, Sylvain Sorin, and Yannick Viossat. Time average replicator and best reply dynamics. *Mathematics of Operations Research*, 34(2):263–269, May 2009.

[18] Lorens A. Imhof. The long-run behavior of the stochastic replicator dynamics. *The Annals of Applied Probability*, 15(1B):1019–1045, 2005.

[19] Ioannis Karatzas and Steven E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, Berlin, 1998.

[20] Rafail Z. Khasminskii. *Stochastic Stability of Differential Equations*. Number 66 in Stochastic Modelling and Applied Probability. Springer-Verlag, Berlin, 2 edition, 2012.

[21] Rafail Z. Khasminskii and N. Potsepun. On the replicator dynamics behavior under Stratonovich type random perturbations. *Stochastic Dynamics*, 6:197–211, 2006.

[22] Hui-Hsiung Kuo. *Introduction to Stochastic Integration*. Springer, Berlin, 2006.

[23] Jean-Francois Le Gall. *Brownian Motion, Martingales, and Stochastic Calculus*. Graduate Texts in Mathematics. Springer International Publishing, 2016.

[24] Panayotis Mertikopoulos and Aris L. Moustakas. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, 20(4): 1359–1388, July 2010.

[25] Panayotis Mertikopoulos and Mathias Staudigl. Convergence to Nash equilibrium in continuous games with noisy first-order feedback. In *CDC '17: Proceedings of the 56th IEEE Annual Conference on Decision and Control*, 2017.

[26] Panayotis Mertikopoulos and Yannick Viossat. Imitation dynamics with payoff shocks. *International Journal of Game Theory*, 45(1-2):291–320, March 2016.

[27] Panayotis Mertikopoulos and Yannick Viossat. Survival of dominated strategies under imitation dynamics. *Journal of Dynamics and Games*, 9(4, William H. Sandholm memorial issue):499–528, October 2022.

[28] John H. Nachbar. Evolutionary selection dynamics in games. *International Journal of Game Theory*, 19:59–89, 1990.

[29] John F. Nash. *Non-cooperative games*. PhD thesis, Princeton University, 1950.

[30] Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.

[31] Bernt Øksendal. *Stochastic Differential Equations*. Springer-Verlag, Berlin, 6 edition, 2007.

[32] Daniel Revuz and Marc Yor. *Continuous Martingales and Brownian Motion*. Springer, Berlin, 3rd edition 1999 edition, 1999.

[33] Klaus Ritzberger and Jörgen W. Weibull. Evolutionary selection in normal-form games. *Econometrica*, 63(6):1371–99, November 1995.

[34] Larry Samuelson and Jianbo Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391, 1992.

[35] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2010.

[36] Karl H. Schlag. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78(1):130–156, 1998.

[37] J. Maynard Smith. The Logic of Animal Conflict. 246(5427):15–18, November 1973.

[38] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.

[39] Bernhard Thomas. On evolutionarily stable sets. *Journal of Mathematical Biology*, 22: 105–115, Jun 1985.

[40] Andrew Vlasic. Long-run analysis of the stochastic replicator dynamics in the presence of random jumps, 2018.

[41] John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.

[42] John von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1947.

[43] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.

[44] H Peyton Young and Dean Foster. Cooperation in the long-run. *Games and Economic Behavior*, 3(1):145–156, 1991.

COLOPHON

This manuscript was typeset with $\LaTeX\,2_\varepsilon$ using Hermann Zapf's Palatino type face (the actual Type 1 PostScript fonts used were URW Palladio L and FPL). The monospaced text (hyperlinks, etc.) was typeset in *Bera Mono*, originally developed by Bitstream, Inc. as "Bitstream Vera" (with Type 1 PostScript fonts by Malte Rosenau and Ulrich Dirr).

The typographic style of this dissertation was inspired by the authoritative genius of Bringhurst's *Elements of Typographic Style*, ported to LaTeX by André Miede, the original designer of the `classicthesis` template. Any unsightly deviations from these works should be attributed solely to the author's (not always successful) efforts to conform to the awkward A4 paper size.

The robustness of game dynamics under random perturbations

Saint-Martin-d'Hères, June 23, 2023

Pierre-Louis Cauvin