

ON THE ROBUSTNESS OF LEARNING IN GAMES WITH STOCHASTICALLY PERTURBED PAYOFF OBSERVATIONS

MARIO BRAVO AND PANAYOTIS MERTIKOPOULOS

ABSTRACT. Motivated by the scarcity of accurate payoff feedback in practical applications of game theory, we examine a class of learning dynamics where players adjust their choices based on past payoff observations that are subject to noise and random disturbances. First, in the single-player case (corresponding to an agent trying to adapt to an arbitrarily changing environment), we show that the stochastic dynamics under study lead to no regret almost surely, irrespective of the noise level in the player's observations. In the multi-player case, we find that dominated strategies become extinct and we show that strict Nash equilibria are stochastically stable and attracting; conversely, if a state is stable or attracting with positive probability, then it is a Nash equilibrium. Finally, we provide an averaging principle for 2-player games, and we show that in zero-sum games with an interior equilibrium, time averages converge to Nash equilibrium for any noise level.

CONTENTS

1. Introduction	1
2. The model	4
3. Regret minimization	10
4. Extinction of dominated strategies	13
5. Long-term stability and convergence analysis	16
6. Time averages in 2-player games	18
7. Discussion	21
Appendix	23
Appendix A. Technical proofs	23
Appendix B. Auxiliary results	29
References	32

2010 *Mathematics Subject Classification.* Primary 60H10, 37N40, 91A26; secondary 60H30, 60J70, 91A22.

Key words and phrases. Dominated strategies; learning; Nash equilibrium; regret minimization; regularization; robustness; stochastic game dynamics; stochastic stability.

The authors are greatly indebted to Roberto Cominetti for arranging the visit of the second author to the University of Chile and for his many constructive comments. The authors would also like to express their gratitude to Mathias Staudigl for his many insightful comments and suggestions, to Bill Sandholm, Josef Hofbauer, and Yannick Viossat for helpful discussions, and to two anonymous referees for their detailed remarks and recommendations.

Part of this work was carried out during the authors' visit to the Hausdorff Research Institute for Mathematics at the University of Bonn in the framework of the Trimester Program "Stochastic Dynamics in Economics and Finance" and during the second author's visit to the University of Chile. MB was partially supported by Fondecyt grant No. 11151003, and the Núcleo Milenio Información y Coordinación en Redes ICM/FIC RC130003. PM was partially supported by the French National Research Agency (grant nos. NETLEARN-13-INFR-004 and GAGA-13-JS01-0004-01) and the French National Center for Scientific Research (grant no. REAL.NET-PEPS-JCJC-INS2I-2014.)

1. INTRODUCTION

A central question in game-theoretic learning is whether the outcome of a learning process (viewed here as a plausible model for the behavior of optimizing agents) is also justifiable from the point of view of rationality – e.g. whether the process leads to a Nash equilibrium or a state where no dominated strategies are present. In that regard, one of the most widely studied learning procedures is the exponential weight (EW) algorithm that was originally introduced by Vovk (1990) and Littlestone and Warmuth (1994) in the context of multi-armed bandit problems. In a game-theoretic setting, the algorithm simply prescribes that players score their actions based on their cumulative payoffs and then assign choice probabilities proportionally to the exponential of each action’s score. As such, the EW algorithm in continuous time (Sorin, 2009) is described by the dynamics

$$\begin{aligned} \dot{y}_{k\alpha} &= v_{k\alpha}, \\ x_{k\alpha} &= \frac{\exp(y_{k\alpha})}{\sum_{\beta} \exp(y_{k\beta})}, \end{aligned} \tag{EW}$$

where, somewhat informally, $v_{k\alpha}$ denotes the payoff to the α -th action of player k , $y_{k\alpha}$ is the action’s performance score (cumulative payoff) over time and $x_{k\alpha}$ is the corresponding mixed strategy weight.

Given its long history and its link with the replicator dynamics of evolutionary game theory (discussed below), the rationality properties of (EW) are also relatively well understood. To name the most important ones: *a*) dominated strategies become extinct in the long run; *b*) limits of interior trajectories and stable rest points are Nash equilibria; *c*) strict equilibria are stable and attracting; and *d*) empirical distributions of play converge to equilibrium in 2-player zero-sum games with an interior equilibrium (Hofbauer and Sigmund, 1998; Sandholm, 2010). More recently, Sorin (2009) also showed that (EW) is *universally consistent*, i.e. players have no regret for following (EW) instead of any other fixed strategy in a dynamically changing environment.

On the other hand, a crucial limitation in the above considerations is that players are assumed to have perfect observations of their actions’ rewards. In practical applications of game theory (e.g. in economics, finance and traffic networks), “perfect feedback” requirements are often too stringent, especially in games with massively many actions and/or players. Accordingly, an important question that arises is whether (EW) retains its rationality properties in the presence of noise and stochastically perturbed payoff observations. Somewhat surprisingly, this is indeed the case (at least for some of them): even if the payoffs in (EW) are perturbed by a Brownian noise term of arbitrarily high variance, dominated strategies become extinct and strict equilibria remain stable and attracting with high probability (Mertikopoulos and Moustakas, 2009, 2010). Thus, even when the players’ true payoffs are masked by noise and uncertainty, the reinforcement learning principle behind (EW) allows players to weed out the noise and leads to similar outcomes as in the noiseless, deterministic regime.

Motivated by the robustness of exponential learning in noisy environments, we examine here a broad class of game-theoretic learning procedures where players adjust their strategies by playing an approximate best response to the vector of their actions’ cumulative payoffs – possibly subject to random disturbances and noise. In the case of perfect payoff observations, this scheme boils down to the reinforcement

learning dynamics considered by Mertikopoulos and Sandholm (2016) who showed that the properties discussed above still hold in the absence of noise. With this in mind, our main contribution is to show that feedback noise does not really matter: even if the players’ payoff observations are subject to *arbitrarily high* (and possibly state-dependent or correlated) noise, dominated strategies become extinct (roughly at the same rate as in deterministic environments); strict Nash equilibria remain (stochastically) stable and attracting; and, in the converse direction, if a state is (stochastically) stable or attracting with positive probability, then it is also a Nash equilibrium. Finally, if players use a decreasing learning parameter to adjust the weight of their scoring process over time, the stochastic learning dynamics under study lead to no regret (a.s.) and their time average converges to equilibrium in 2-player zero-sum games with an interior equilibrium.

Our analysis also highlights an important difference between “static” solution concepts (such as Nash equilibria and strategic dominance) and more “dynamic” notions (such as regret minimization). Whereas the speed of convergence to static target states is accelerated by the use of a large, constant learning parameter (noise and disturbances notwithstanding), the rate of regret minimization is optimized by using a learning parameter that decays proportionally to $t^{-1/2}$ (and which guarantees an $\mathcal{O}(\sqrt{t \log \log t})$ bound for the players’ cumulative regret under uncertainty). This disparity only appears in the noisy regime and is due to the fact that players need to be more conservative when facing a fluid environment that varies with time in an (a priori) unpredictable fashion. Otherwise, if players have access to noiseless payoff observations, they can be significantly more greedy and achieve lower regret faster by using a constant learning parameter.

1.1. Related work. The long-term rationality properties of exponential learning in a game-theoretic setting were first studied in conjunction with those of the *replicator dynamics*, one of the most widely studied dynamical systems for population evolution under natural selection. Indeed, a simple differentiation of (EW) reveals that the evolution of the players’ mixed strategies under (EW) follows the differential equation:

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

which is simply the (multi-population) replicator dynamics of Taylor and Jonker (1978). In this context, Akin (1980), Nachbar (1990) and Samuelson and Zhang (1992) showed that dominated strategies become extinct, while it is well known that *a*) stable states of (RD) are Nash equilibria; *b*) strict Nash equilibria are asymptotically stable; and *c*) time averages of replicator orbits converge to equilibrium in 2-player games provided that no strategy share becomes arbitrarily small (Hofbauer and Sigmund, 1998, 2003).

This two-way relationship between exponential learning and the replicator dynamics was noted early on by Rustichini (1999) in his study of reinforcement learning models for games – i.e. learning how to react to a given situation so as to maximize a numerical reward (Sutton and Barto, 1998). From a game-theoretic point of view, the learning models of Börgers and Sarin (1997) and Erev and Roth (1998) are also closely related to the replicator dynamics – and, hence, to (EW) – while Fudenberg and Levine (1998), Hofbauer and Sandholm (2002), and Hopkins (2002) studied a smooth variant of the well-known fictitious play algorithm where the players play a perturbed best response to their opponents’ empirical frequency

of play.¹ Up to mild technical differences, the correlated version of these models can be seen as a noiseless version of our reinforcement learning model, viz. the dynamics of Mertikopoulos and Sandholm (2016) with the parameter choice $\eta(t) = 1/t$; we explore this relation in more detail in Sections 3 and 6.

Of course, a crucial aspect of these considerations is whether players have accurate observations of their actions' rewards or only a noisy estimate thereof: if the former is not the case, noise and fluctuations could potentially lead to suboptimal outcomes with high probability. In biology and evolutionary game theory (where payoffs measure the reproductive fitness of a biological species or the average payoff of populations of nonatomic players respectively), Fudenberg and Harris (1992) accounted for such fluctuations by introducing a stochastic variant of the replicator dynamics where evolution is perturbed by “aggregate shocks” that reflect the impact of weather-like effects and other perturbations.² In this framework, Cabrales (2000), Imhof (2005) and Hofbauer and Imhof (2009) showed that dominated strategies are still eliminated if the variability of the shocks across different genotypes (strategies) is not too high, while Imhof (2005) and Hofbauer and Imhof (2009) showed that strict Nash equilibria of a modified game are stochastically asymptotically stable under the replicator dynamics with aggregate shocks.³

On the other hand, Mertikopoulos and Moustakas (2009, 2010) showed that the dynamics obtained by Fudenberg and Harris (1992) do not coincide with the stochastic replicator dynamics induced by (EW) in the presence of random disturbances and measurement noise – in contrast to the noiseless case where (EW) and (RD) *do* coincide. As we mentioned above, this learning variant of the stochastic replicator dynamics actually retains the rationality properties of the deterministic system (EW)/(RD) without any caveats on the noise: dominated strategies become extinct and strict Nash equilibria remain stochastically asymptotically stable without any conditions on the perturbations' magnitude. This shows that the origin of the noise is crucial in the determination of the dynamics' long-term properties: whereas the “aggregate shocks” of evolutionary environments lead to rational behavior in a modified game, no such modifications are required when learning under uncertainty.

1.2. Paper outline. In Section 2, we present our model for learning in the presence of noise and we derive the system of coupled stochastic differential equations (SDEs) that governs the evolution of the players' mixed strategies. In Section 3, we show that if the players use a smoothly decreasing learning parameter, then the learning dynamics under study lead to no regret (a.s.), whatever the noise level. In Section 4, we investigate dominated strategies: we show that dominated strategies become extinct (a.s.) and we derive an explicit bound for their extinction rate. Section 5 focuses on the dynamics' long-term stability and convergence properties: we show that *a*) stochastically (Lyapunov) stable states and states that attract trajectories of play with positive probability are Nash equilibria; and *b*) strict Nash equilibria are stochastically asymptotically stable, irrespective of the fluctuations' magnitude.

¹For a closely related model, see also Cominetti, Melo, and Sorin (2010).

²Khasminskii and Potsepun (2006) also consider a Stratonovich-based model while Vlasic (2012) examines the case of random jumps incurred by catastrophic, earthquake-like events

³Whether the equilibria of the original game are themselves asymptotically stable, depends on the intensity of the noise on different strategies and also on the exact way that the noise enters the process (Mertikopoulos and Viossat, 2016).

In Section 6, we provide an averaging principle for 2-player games in the spirit of Hofbauer, Sorin, and Viossat (2009); thanks to this principle, we then show that empirical distributions of play converge to Nash equilibrium in zero-sum games (again, no matter the noise level). Finally, in Section 7, we discuss the relaxation of some of the assumptions on the noise process – and, in particular, the independence of the observation noise across players and strategies.

To streamline our presentation, we included several numerical illustrations in the main text (Figs. 1 and 2) and we relegated the more convoluted proofs to a series of appendices at the end.

2. THE MODEL

After a few preliminaries to set notation and terminology, this section focuses on the basic properties of a broad class of reinforcement learning dynamics under noise and uncertainty. In the noiseless, deterministic regime (Section 2.2), our model essentially boils down to the class of dynamics recently studied by Mertikopoulos and Sandholm (2016). The full stochastic framework (Section 2.3) is then obtained by positing that the agents’ observations are perturbed at each moment in time by a zero-mean stochastic process (an Itô diffusion).

2.1. Preliminaries.

Notation. If $\mathcal{S} = \{s_\alpha\}_{\alpha=1}^n$ is a finite set, the real vector space generated by \mathcal{S} will be denoted by $\mathbb{R}^{\mathcal{S}}$ and we will write $\{e_s\}_{s \in \mathcal{S}}$ for its canonical basis; for concision, we will also use α to refer interchangeably to α or s_α , writing e.g. x_α instead of x_{s_α} . The set $\Delta(\mathcal{S})$ of probability measures on \mathcal{S} will be identified with the n -dimensional simplex $\Delta = \{x \in \mathbb{R}^{\mathcal{S}} : \sum_\alpha x_\alpha = 1 \text{ and } x_\alpha \geq 0\}$ and the relative interior of Δ will be written Δ° .

If $\{\mathcal{S}_k\}_{k \in \mathcal{N}}$ is a finite family of sets, we will use the shorthand $(\alpha_k; \alpha_{-k})$ for the tuple $(\dots, \alpha_{k-1}, \alpha_k, \alpha_{k+1}, \dots)$ and we will write \sum_α^k instead of $\sum_{\alpha \in \mathcal{S}_k}$. Finally, any statement of the form $f(t) = o(g(t))$ for a stochastic process f should be interpreted in the a.s. sense, i.e. “for every $C > 0$, there exists a.s. some (random) t_0 such that $|f(t)| \leq C |g(t)|$ for all $t \geq t_0$ ” – and likewise for the statement “ $f(t) = \mathcal{O}(g(t))$ ”. Definitions from game theory. A *finite game in normal form* is a tuple $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \mathcal{A}, u)$ consisting of a) a finite set of *players* $\mathcal{N} = \{1, \dots, N\}$; b) a finite set \mathcal{A}_k of *actions* (or *pure strategies*) per player $k \in \mathcal{N}$; and c) the players’ *payoff functions* $u_k: \mathcal{A} \rightarrow \mathbb{R}$, where $\mathcal{A} \equiv \prod_k \mathcal{A}_k$ denotes the set of all joint action profiles $(\alpha_1, \dots, \alpha_N)$. The set of *mixed strategies* of player k is denoted by $\mathcal{X}_k \equiv \Delta(\mathcal{A}_k)$ and the space $\mathcal{X} \equiv \prod_k \mathcal{X}_k$ of *mixed strategy profiles* $x = (x_1, \dots, x_N)$ will be called the game’s *strategy space*. In this mixed context, the expected payoff of player k in the strategy profile $x = (x_1, \dots, x_N) \in \mathcal{X}$ is

$$u_k(x) = \sum_{\alpha_1}^1 \cdots \sum_{\alpha_N}^N u_k(\alpha_1, \dots, \alpha_N) x_{1,\alpha_1} \cdots x_{N,\alpha_N}, \quad (2.1)$$

where, in a slight abuse of notation, $u_k(\alpha_1, \dots, \alpha_N)$ denotes the payoff of player k in the profile $(\alpha_1, \dots, \alpha_N) \in \mathcal{A}$. Accordingly, the payoff corresponding to $\alpha \in \mathcal{A}_k$ in the mixed strategy profile $x \in \mathcal{X}$ is

$$v_{k\alpha}(x) \equiv u_k(\alpha; x_{-k}) = \sum_{\alpha_1}^1 \cdots \sum_{\alpha_N}^N u_k(\alpha_1, \dots, \alpha_N) x_{1,\alpha_1} \cdots \delta_{\alpha_k, \alpha} \cdots x_{N,\alpha_N}, \quad (2.2)$$

leading to the more concise expression

$$u_k(x) = \sum_{\alpha}^k x_{k\alpha} v_{k\alpha}(x) = \langle v_k(x), x_k \rangle \quad (2.3)$$

where $v_k(x) = (v_{k\alpha}(x))_{\alpha \in \mathcal{A}_k}$ denotes the *payoff vector* of player k at $x \in \mathcal{X}$.

2.2. The deterministic model. Our learning model will be based on the following simple idea: the game’s players “score” their actions by keeping track of their cumulative payoffs and then, at each moment $t \geq 0$, they play an “approximate” best response to this vector of performance scores (thus reinforcing the probability of playing an action with a higher overall payoff). More precisely, following Mertikopoulos and Sandholm (2016), this process can be described by the dynamics

$$\begin{aligned} y_k(t) &= \int_0^t v_k(x(s)) ds, \\ x_k(t) &= Q_k(\eta_k(t)y_k(t)), \end{aligned} \quad (\text{RL})$$

where:

- (1) the *score vector* $y_k(t) = (y_{k\alpha}(t))_{\alpha \in \mathcal{A}_k}$ of player k ranks each strategy $\alpha \in \mathcal{A}_k$ based on its cumulative payoff $y_{k\alpha}(t) = \int_0^t v_{k\alpha}(x(s)) ds$ up to time t .
- (2) $Q_k(y_k)$ is a *choice map* which reinforces the strategies with the highest scores (see below for a rigorous definition).
- (3) $\eta_k(t) > 0$ is a *learning parameter* which can be tuned freely by each player.

Clearly, a natural choice for the “scores-to-strategies” map Q_k would be to take the arg max correspondence $y_k \mapsto \arg \max_{x_k \in \mathcal{X}_k} \langle y_k, x_k \rangle$, i.e. to greedily assign all weight to the strategy (or strategies) with the highest score. However, since the arg max operator is multi-valued (so each player would also need to employ some tie-breaking rule to resolve ambiguities), we will focus on *single-valued* choice maps of the general form

$$Q_k(y_k) = \arg \max_{x_k \in \mathcal{X}_k} \{ \langle y_k, x_k \rangle - h_k(x_k) \}, \quad (2.4)$$

where the *penalty function* $h_k: \mathcal{X}_k \rightarrow \mathbb{R}$ satisfies the following properties:

- a) h_k is continuous on \mathcal{X}_k .
- b) h_k is smooth on the relative interior of every face of \mathcal{X}_k .
- c) h_k is *strongly convex* on \mathcal{X}_k , i.e. there exists some $K > 0$ such that

$$h_k(tx_k + (1-t)x'_k) \leq th_k(x_k) + (1-t)h_k(x'_k) - \frac{1}{2}Kt(1-t)\|x'_k - x_k\|^2, \quad (2.5)$$

for all $x_k, x'_k \in \mathcal{X}_k$ and for all $t \in [0, 1]$.

This “softening” of the arg max operator has a long history in game theory and optimization, and the induced map Q_k is intimately related to the notion of *softmax* or *perturbed/regularized best response maps*; for an in-depth discussion, we refer the reader to Hofbauer and Sandholm (2002) and Mertikopoulos and Sandholm (2016). For our immediate purposes, the key observation is that (2.4) is a strictly concave problem, so it admits a unique solution for every input vector y_k . Therefore, when the penalty term h_k is small relative to y_k , $Q_k(y_k)$ can be seen as a single-valued approximation of the standard best response correspondence $y_k \mapsto \arg \max_{x_k \in \mathcal{X}_k} \langle y_k, x_k \rangle$.

Regarding the learning parameter η_k , its role in (RL) is to temper the growth of the cumulative payoff vector $y_k(t)$ so as to allow the player to better explore his strategies (instead of prematurely reinforcing one or another). In other words,

η_k can be interpreted as an extrinsic weight that the player assigns to cumulative observations of past payoffs.⁴ Accordingly, given that $y(t)$ grows (at most) as $\mathcal{O}(t)$, we will assume throughout that:

Assumption 1. $\eta_k(t)$ is C^1 -smooth, nonincreasing and $\lim_{t \rightarrow \infty} t\eta_k(t) = +\infty$.

The decay rate of $\eta(t)$ will play a crucial role when the players' payoff observations are subject to stochastic perturbations, an issue that we will explore in detail in later sections. For now, we turn to two representative examples of (RL):

Example 2.1. The prototype penalty function on the simplex is the Gibbs (negative) entropy $h(x) = \sum_{\alpha} x_{\alpha} \log x_{\alpha}$ which, after a standard calculation, yields the so-called *logit map*:

$$G_{\alpha}(y) = \frac{\exp(y_{\alpha})}{\sum_{\beta} \exp(y_{\beta})}. \quad (2.6)$$

An easy differentiation then yields

$$\dot{x}_{k\alpha} = \frac{e^{y_{k\alpha}} \dot{y}_{k\alpha}}{\sum_{\beta} e^{y_{k\beta}}} - \frac{e^{y_{k\alpha}} \sum_{\beta} e^{y_{k\beta}} \dot{y}_{k\beta}}{\left(\sum_{\beta} e^{y_{k\beta}}\right)^2} = x_{k\alpha} \left[v_{k\alpha}(x) - \sum_{\beta} x_{k\beta} v_{k\beta}(x) \right], \quad (\text{RD})$$

which is simply the (multi-population) replicator equation of Taylor and Jonker (1978) for population evolution under natural selection. For a more thorough treatment of the links between logit choice and the replicator dynamics, see Hofbauer et al. (2009), Mertikopoulos and Moustakas (2010), and references therein.

Example 2.2. As another example, consider the penalty function $h(x) = \frac{1}{2} \sum_{\alpha} x_{\alpha}^2$. This quadratic penalty function leads to the *projected choice map*

$$\Pi(y) = \arg \min_{x \in \Delta} \left\{ \langle y, x \rangle - \frac{1}{2} \|x\|^2 \right\} = \arg \min_{x \in \Delta} \|y - x\|^2, \quad (2.7)$$

and, as was shown by Mertikopoulos and Sandholm (2016), the induced trajectories $x(t) = \Pi(y(t))$ of (RL) satisfy the so-called *projection dynamics*

$$\dot{x}_{k\alpha} = \begin{cases} v_{k\alpha}(x) - |\text{supp}(x_k)|^{-1} \sum_{\beta \in \text{supp}(x_k)} v_{k\beta}(x) & \text{if } x_{k\alpha} > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{PD})$$

on an open dense set of times (in particular, except when the support of $x(t)$ changes). The dynamics (PD) were introduced in game theory by Friedman (1991) as a geometric model of the evolution of play in population games; for a closely related model (but with different long-term properties), see Nagurney and Zhang (1997) and Lahkar and Sandholm (2008).

⁴The role of the learning parameter $\eta(t)$ can also be linked to the vanishing step-size rules that are used in the theory of stochastic approximation – see e.g. Benaïm (1999), Lambertson, Pagès, and Tarrès (2004), Oyarzun and Ruf (2014) and references therein. The difference between the two is that, in the theory of stochastic approximation, a variable step-size means that new information enters the algorithm with decreasing weight; on the other hand, in our context, all information is weighted evenly, but the *aggregate* information is weighted by $\eta(t)$ to avoid extreme behaviors.

Related models. The reinforcement mechanism of (RL) is seen quite clearly in the work of Vovk (1990) and Littlestone and Warmuth (1994) on multi-armed bandits. Therein, the agent is faced with a repeated decision process (e.g. choosing a slot machine in the eponymous problem) and, at each stage, he selects an action with probability exponentially proportional to an estimate of said action’s cumulative payoff up to that time. In this way, the mean dynamics of the agent’s learning process boil down to the exponential learning scheme (EW) of Example 2.1 – itself a special case of (RL).

Leslie and Collins (2005) and Coucheney, Gaujal, and Mertikopoulos (2015) also considered a reinforcement learning process where players discount past observations by a constant multiplicative factor and then play a perturbed best response to the resulting cumulative payoff vector – or estimate thereof. The mean dynamics that describe this process in continuous time are then given by (RL) with constant η and an additional adjustment term that accounts for the exponential discounting of past observations. For a more detailed survey of the surrounding literature, we refer the reader to Fudenberg and Levine (1998) and Mertikopoulos and Sandholm (2016).

2.3. Learning in the presence of noise. A key assumption underlying the reinforcement learning scheme (RL) is that the players’ payoff observations are impervious to any sort of exogenous random noise. However, this assumption is rarely met in practical applications of game-theoretic learning: for instance, in telecommunication networks and traffic engineering, signal strength and latency measurements are constantly subject to stochastic fluctuations which introduce noise to the input of any learning algorithm (Kang, Kelly, Lee, and Williams, 2009; Kelly, Maulloo, and Tan, 1998). Thus, to account for the lack of accurate payoff information in settings where uncertainty is an issue, we will consider the *stochastically perturbed reinforcement learning* process:

$$\begin{aligned} dY_{k\alpha} &= v_{k\alpha}(X) dt + \sigma_{k\alpha}(X) dW_{k\alpha}, \\ X_k &= Q_k(\eta_k Y_k), \end{aligned} \tag{SRL}$$

where $W_{k\alpha}$ is a family of standard Wiener processes (assumed independent across players and strategies) and the diffusion coefficients $\sigma_{k\alpha}: \mathcal{X} \rightarrow \mathbb{R}$ (assumed Lipschitz) measure the strength of the noise process.

Intuitively, the random component of (SRL) means that the players’ observed payoffs are only accurate up to an error term of the order of $\sigma_{k\alpha}$, possibly depending on the players’ mixed strategies but otherwise uncorrelated over players and strategies.⁵ As such, (SRL) should be viewed as a special instance of the more general case where observation errors exhibit correlations between different actions: for instance, if each player’s action corresponds to a choice of route in a congestion game, any two routes that overlap will exhibit such correlations. The impact of including such correlations in our model is discussed at length in Section 7; however, for simplicity, our analysis will be stated in the uncorrelated case.

⁵An alternative source of stochasticity could be any inherent randomness in the players’ payoffs – for instance, if there is a random component to the players’ payoffs. A detailed analysis of this case would require a careful reformulation of the underlying game which, for simplicity, we do not attempt here; for a related treatment in an evolutionary context, see Mertikopoulos and Viossat (2016).

Regarding the existence and uniqueness of solutions to (SRL), Proposition B.1 in Appendix B shows that the players' choice maps Q_k are Lipschitz continuous, so (SRL) admits a unique (strong) solution $Y(t)$ for every initial condition $Y(0)$. Standard arguments can then be used to show that these solutions exist for all time (a.s.), so the players' mixed strategy profile $X(t) = Q(\eta(t)Y(t))$ is also fully determined for all $t \geq 0$. However, since this is a somewhat indirect description of the evolution of $X(t)$, our first task will be to derive the governing dynamics of $X(t)$ in the form of a stochastic differential equation stated directly on \mathcal{X} .

For simplicity, we only present here the special case where each player's penalty function is of the *separable* form:

$$h_k(x_k) = \sum_{\alpha}^k \theta_k(x_{k\alpha}) \quad (2.8)$$

for some strongly convex *kernel function* $\theta_k \in C^0[0, 1] \cap C^3(0, 1]$. We then get:

Proposition 2.1. *Let $X(t)$ be an orbit of (SRL) in \mathcal{X} and let I be a random open interval such that $\text{supp}(X(t))$ remains constant over I . Then, for all $t \in I$, $X(t)$ satisfies the stochastic differential equation*

$$dX_{k\alpha} = \frac{\eta_k}{\theta''_{k\alpha}} \left[v_{k\alpha} - \Theta''_k \sum_{\beta} v_{k\beta} / \theta''_{k\beta} \right] dt \quad (2.9a)$$

$$+ \frac{\eta_k}{\theta''_{k\alpha}} \left[\sigma_{k\alpha} dW_{k\alpha} - \Theta''_k \sum_{\beta} \sigma_{k\beta} / \theta''_{k\beta} dW_{k\beta} \right] \quad (2.9b)$$

$$+ \frac{\dot{\eta}_k}{\eta_k} \frac{1}{\theta''_{k\alpha}} \left[\theta'_{k\alpha} - \Theta''_k \sum_{\beta} \theta'_{k\beta} / \theta''_{k\beta} \right] dt \quad (2.9c)$$

$$- \frac{1}{2} \frac{1}{\theta''_{k\alpha}} \left[\theta'''_{k\alpha} U_{k\alpha}^2 - \Theta''_k \sum_{\beta} \theta'''_{k\beta} / \theta''_{k\beta} U_{k\beta}^2 \right] dt, \quad (2.9d)$$

where all summations are taken over $\beta \in \text{supp}(X_k)$ and:

- a) $\theta'_{k\alpha} = \theta'_k(X_{k\alpha})$, $\theta''_{k\alpha} = \theta''_k(X_{k\alpha})$, $\theta'''_{k\alpha} = \theta'''_k(X_{k\alpha})$,
- b) $\Theta''_k = \left(\sum_{\beta} 1 / \theta''_{k\beta} \right)^{-1}$,
- c) $U_{k\alpha}^2 = \left(\frac{\eta_k}{\theta''_{k\alpha}} \right)^2 \left[\sigma_{k\alpha}^2 \left(1 - \Theta''_k / \theta''_{k\alpha} \right)^2 + \sum_{\beta \neq \alpha} \left(\Theta''_k / \theta''_{k\beta} \right)^2 \sigma_{k\beta}^2 \right]$.

In particular, if $\lim_{z \rightarrow 0^+} \theta'_k(z) = -\infty$ for all $k \in \mathcal{N}$, $X(t)$ is an ordinary (strong) solution of (2.9); otherwise, $X(t)$ satisfies (2.9) on a random open dense subset of $[0, \infty)$.

The proof of Proposition 2.1 is a simple but interesting application of Itô's formula; to streamline our presentation, we relegate it to Appendix A and we focus here on some remarks and examples (see also Fig. 1 for some sample trajectories of the process):

Remark 2.1. Even though the dynamics (2.9) appear quite convoluted, each of the constituent terms (2.9a)–(2.9d) admits a relatively simple interpretation:

- a) The term (2.9a) drives the process in the case $\sigma = 0$, $\eta = \text{constant}$. These are the deterministic dynamics studied by Mertikopoulos and Sandholm (2016) and correspond to learning in settings with no uncertainty.
- b) The diffusion term (2.9b) reflects the direct impact of the noise on (SRL).

- c) The term (2.9c) is due to the variability of the players' learning rate η so its impact on (2.9) vanishes if $\eta(t) \rightarrow 0$ sufficiently fast. On the other hand, this term obviously persists even if there is no noise in the players' learning process.
- d) Finally, the term (2.9d) is the Itô correction induced on X_k through (SRL) due to the non-anticipative nature of the Itô integral.⁶ Importantly, this term depends on η but not $\dot{\eta}$, so it does not vanish for constant η .

Example 2.3. As we saw in Example 2.1, the replicator dynamics (RD) correspond to the entropic kernel $\theta(x) = x \log x$. In this case, (2.9) leads to the following stochastic variant of the replicator dynamics:

$$\begin{aligned}
dX_{k\alpha} &= \eta_k X_{k\alpha} \left[v_{k\alpha} - \sum_{\beta}^k X_{k\beta} v_{k\beta} \right] dt \\
&+ \eta_k X_{k\alpha} \left[\sigma_{k\alpha} dW_{k\alpha} - \sum_{\beta}^k \sigma_{k\beta} X_{k\beta} dW_{k\beta} \right] \\
&+ \frac{\dot{\eta}_k}{\eta_k} X_{k\alpha} \left[\log X_{k\alpha} - \sum_{\beta}^k X_{k\beta} \log X_{k\beta} \right] dt \\
&+ \frac{\dot{\eta}_k^2}{2} X_{k\alpha} \left[\sigma_{k\alpha}^2 (1 - 2X_{k\alpha}) - \sum_{\beta}^k \sigma_{k\beta}^2 X_{k\beta} (1 - 2X_{k\beta}) \right] dt.
\end{aligned} \tag{SRD}$$

For constant η , (SRD) is simply the stochastic replicator dynamics of exponential learning studied by Mertikopoulos and Moustakas (2010). As such, (SRD) should be contrasted to the evolutionary *replicator dynamics with aggregate shocks* of Fudenberg and Harris (1992):

$$\begin{aligned}
dX_{k\alpha} &= X_{k\alpha} \left[v_{k\alpha} - \sum_{\beta}^k X_{k\beta} v_{k\beta} \right] dt \\
&+ X_{k\alpha} \left[\sigma_{k\alpha} dW_{k\alpha} - \sum_{\beta}^k \sigma_{k\beta} X_{k\beta} dW_{k\beta} \right] \\
&- X_{k\alpha} \left[\sigma_{k\alpha}^2 X_{k\alpha} - \sum_{\beta}^k \sigma_{k\beta}^2 X_{k\beta}^2 \right] dt,
\end{aligned} \tag{ASRD}$$

where $X_{k\alpha}$ denotes the population share of the α -th genotype of species k in a multi-species environment, $v_{k\alpha}$ represents its reproductive fitness, and the noise coefficients $\sigma_{k\alpha}$ measure the impact of random weather-like effects on population evolution.⁷

Besides the absence of the learning rate η , the fundamental difference between (SRD) and (ASRD) is in their Itô correction: as we shall see in what follows, this term leads to a drastically different long-term behavior and highlights an important contrast between learning and evolution in the presence of noise.

⁶If (SRL) had been formulated as a Stratonovich equation, (2.9d) would vanish; however, the future-anticipating nature of the Stratonovich integral (van Kampen, 1981) is not well-suited for our purposes.

⁷For a comprehensive account of the literature surrounding (ASRD), see Hofbauer and Imhof (2009), Mertikopoulos and Viossat (2016), and references therein. In addition to the works mentioned above, Khasminskii and Potsepun (2006) study a Stratonovich-based formulation of (ASRD) while Vlasic (2012) also considers random jumps induced by discontinuous, earthquake-like events.




FIGURE 1. Evolution of play under (SRL) with logit (left) and choice maps (right) for a 2×2 congestion game (Nash equilibria are depicted in red; for the game's payoffs, see the vertex labels). For comparison purposes, we took $\sigma_{k\alpha} = 1$ for all $\alpha \in \mathcal{A}_k$, $k = 1, 2$, and we used the same Wiener process realization in both cases: the orbits of projected reinforcement learning attain the boundary of \mathcal{X} and converge to Nash equilibrium much faster than in the case of exponential learning.

Example 2.4. In the case of the projected reinforcement learning scheme (PD), substituting $\theta(x) = x^2/2$ in (2.9) yields the *stochastic projection dynamics*:

$$\begin{aligned}
 dX_{k\alpha} = & \left[v_{k\alpha} - |\text{supp}(X_k)|^{-1} \sum_{\beta \in \text{supp}(X_k)} v_{k\beta} \right] dt \\
 & + \left[\sigma_{k\alpha} dW_{k\alpha} - |\text{supp}(X_k)|^{-1} \sum_{\beta \in \text{supp}(X_k)} \sigma_{k\beta} dW_{k\beta} \right] \quad (\text{SPD}) \\
 & + \frac{\dot{\eta}_k}{\eta_k} \left[X_{k\alpha} - |\text{supp}(X_k)|^{-1} \right] dt.
 \end{aligned}$$

There are two important qualitative differences between (SRD) and (SPD): first, (SRD) holds for all $t \geq 0$ whereas (SPD) describes the solution orbits of (SRL) only on intervals over which the support of $X(t)$ remains constant. Second, the projection mapping Π of (2.7) is piecewise linear, so there is no Itô correction in (SPD); accordingly, the distinction between Itô and Stratonovich perturbations becomes void in the context of (SPD).

3. REGRET MINIMIZATION

We begin our rationality analysis in the case where there is a single player whose payoffs are determined by the state of his environment – which, in turn, may evolve *arbitrarily* over time (including adversarially if the player is playing against an opponent).

More precisely, consider a decision process where, at each $t \geq 0$, the player chooses an action from a finite set \mathcal{A} according to some mixed strategy $x(t) \in \mathcal{X} \equiv \Delta(\mathcal{A})$ and his expected payoff is determined by the (a priori unknown) payoff vector

$v(t) = (v_\alpha(t))_{\alpha \in \mathcal{A}}$ of stage t . In this context, the performance of a dynamic strategy $x(t)$ can be measured by comparing the player's (expected) cumulative payoff to the payoff that he could have obtained if the state of nature were known in advance and the player had best-responded to it; specifically, the player's *cumulative regret* at time t is defined as

$$\text{Reg}(t) = \max_{\alpha \in \mathcal{A}} \int_0^t v_\alpha(s) ds - \int_0^t \langle v(s), x(s) \rangle ds, \quad (3.1)$$

and we say that a strategy $x(t)$ is *consistent* if it leads to *no (average) regret*, i.e.

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \text{Reg}(t) \leq 0 \quad (\text{a.s.}), \quad (3.2)$$

or, equivalently:

$$\text{Reg}(t) = o(t) \quad \text{as } t \rightarrow \infty. \quad (3.3)$$

The notion of consistency presented above is commonly referred to as *external* or *universal* consistency and was originally introduced by Hannan (1957) in a discrete-time context (for a detailed overview, see Fudenberg and Levine (1998), Shalev-Shwartz (2011) and references therein). There, the agent is assumed to receive a payoff vector $v_n \in \mathbb{R}^{\mathcal{A}}$ at each stage $n \in \mathbb{N}$ and, observing the past realizations v_1, \dots, v_{n-1} , he chooses an action $\alpha_n \in \mathcal{A}$ according to some probability distribution $x_n \in \Delta(\mathcal{A})$. The strategy x_n is then said to be *externally consistent* if, on average, it earns more than any action (or expert suggestion) $\alpha \in \mathcal{A}$; in other words, the no-regret criterion (3.2) is simply the continuous-time analogue of the standard definition of external consistency.

With this in mind, the main question that we seek to address here is whether the perturbed reinforcement learning process (SRL) is consistent. To make this precise, we will focus on the *unilateral* process

$$\begin{aligned} dY_\alpha(t) &= v_\alpha(t) dt + \sigma_\alpha(t) dW_\alpha(t), \\ X(t) &= Q(\eta(t)Y(t)), \end{aligned} \quad (\text{SRL-U})$$

where $v(t)$ is a locally integrable stream of payoff vectors, $W(t)$ is a Wiener process in $\mathbb{R}^{\mathcal{A}}$ and the noise coefficients σ_α (assumed continuous and bounded) represent the error in the player's payoff observations. In the deterministic case ($\sigma = 0$), Sorin (2009) proved that the unilateral variant of (EW) is consistent, a result which was recently extended by Kwon and Mertikopoulos (2014) to the more general setting of (RL). Below, we show that this learning process remains consistent even when the player's payoff observations are subject to arbitrarily high observation noise:

Theorem 3.1. *Assume that (SRL-U) is run with learning parameter $\eta(t)$ satisfying $\lim_{t \rightarrow \infty} \eta(t) = 0$ (in addition to Assumption 1) and initial bias $Y(0) = 0$. Then, (SRL-U) is consistent and it enjoys the cumulative regret bound:*

$$\text{Reg}(t) \leq \frac{\Omega}{\eta(t)} + \sigma_{\max}^2 \frac{|\mathcal{A}|}{2K} \int_0^t \eta(s) ds + \mathcal{O}(\sigma_{\max} \sqrt{t \log \log t}), \quad (3.4)$$

where $\Omega = \max\{h(x) - h(x') : x, x' \in \mathcal{X}\}$, $\sigma_{\max} = \sup_t \max_\beta \sigma_\beta(t)$ and K is the strong convexity constant of the player's penalty function h .

Remark. The assumption $Y(0) = 0$ is only made for simplicity: the policy (SRL-U) is consistent for any initial score vector $Y(0)$ but the corresponding regret guarantee is a bit more cumbersome to write down.

The basic idea of our proof will be to compare the cumulative payoff of the policy (SRL-U) up to time t to that of an arbitrary test strategy $p \in \mathcal{X}$. To that end, we will examine how far the induced trajectory of play $X(t) = Q(\eta(t)Y(t))$ can stray from p ; however, since $X(t)$ is defined via the cumulative payoff process $Y(t)$, we will carry out this comparison between p and $Y(t)$ *directly* using the so-called *Fenchel coupling*:

$$F(x, y) = h(x) + h^*(y) - \langle y, x \rangle, \quad (3.5)$$

where

$$h^*(y) = \max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\} \quad (3.6)$$

denotes the *convex conjugate* of h (Rockafellar, 1970).⁸ By Fenchel’s inequality (Rockafellar, 1970), $F(x, y)$ is non-negative in both arguments, so it provides a “congruity” measure between x and y . With this in mind, our proof strategy will be to express the player’s regret with respect to a test strategy $p \in \mathcal{X}$ in terms of $F(p, Y(t))$ and then show that the latter grows sublinearly in t .

The details of the proof are presented in Appendix A and rely on Itô’s lemma and the law of the iterated logarithm (which is used to control the impact of the noise on the agent’s learning process). For now, we focus on some aspects of Theorem 3.1 and the regret bound (3.4):

Remark 3.1 (The role of η). It is important to note that the second term of (3.4) becomes linear in t when the player’s learning parameter η is constant, explaining in this way the requirement that $\eta(t) \rightarrow 0$ as $t \rightarrow \infty$. On the other hand, this requirement can be dropped in the noiseless case: when $\sigma = 0$, the guarantee (3.4) reduces to $\Omega/\eta(t)$, so (SRL-U) remains consistent even for constant η – a fact which was first observed by Sorin (2009) in the context of (EW) and Kwon and Mertikopoulos (2014) for (RL). In the presence of noise, we conjecture that (SRL-U) may lead to positive regret with positive probability for constant η but we have not been able to prove this.

The form of the bound (3.4) also highlights a trade-off between more aggressive learning rates (slowly decaying η) and the noise affecting the player’s payoff observations. Specifically, the first (deterministic) term of (3.4) is decreasing in η while the second one (which is due to all the noise) is increasing in η ;⁹ as a result, in the absence of noise ($\sigma = 0$), it is better to use a large, constant η rather than letting $\eta(t) \rightarrow 0$, but this might lead to disastrous results under uncertainty. Shalev-Shwartz (2011) draws a similar conclusion for the discrete-time analogue of (RL) known as online mirror descent (OMD); in this way, discretization and random payoff disturbances seem to have comparably adverse effects on the agent’s learning process.

The above considerations can be illustrated more explicitly by choosing $\eta(t) = t^{-\gamma}$ for some $\gamma \in (0, 1)$; in this case, Theorem 3.1 yields the regret bounds:

⁸The terminology “Fenchel coupling” is due to Mertikopoulos and Sandholm (2016) and reflects the fact that (3.5) collects all the terms of Fenchel’s inequality.

⁹The last term of (3.4) is also due to the noise, but it does not otherwise depend on $\eta(t)$.

Corollary 3.2. *Assume that (SRL-U) is run with learning rate $\eta(t) \sim t^{-\gamma}$ for some $\gamma \in (0, 1)$. Then:*

$$\text{Reg}(t) = \begin{cases} \mathcal{O}(t^{1-\gamma}) & \text{if } 0 < \gamma < \frac{1}{2}, \\ \mathcal{O}(\sqrt{t \log \log t}) & \text{if } \gamma = \frac{1}{2}, \\ \mathcal{O}(t^\gamma) & \text{if } \frac{1}{2} < \gamma < 1. \end{cases} \quad (3.7)$$

Proof. Simply note that $t^{\max\{\gamma, 1-\gamma\}}$ is the dominant term in (3.4) for all $\gamma \neq 1/2$. Otherwise, for $\gamma = 1/2$, the first two terms of (3.4) are both $\mathcal{O}(\sqrt{t})$ and are dominated by the third. \square

Remark 3.2 (Links with vanishingly smooth fictitious play). We close this section by discussing the links of (SRL-U) with the vanishingly smooth fictitious play process that was recently examined by Benaïm and Faure (2013) in a discrete-time setting. Specifically, by interpreting $t^{-1} \int_0^t v(s) ds$ as the payoff that a player obtains in a 2-player game against his opponent’s empirical frequency of play (cf. Section 6), the strategy

$$x(t) = Q \left(\eta(t) \int_0^t v(s) ds \right) = Q \left(t\eta(t) \cdot t^{-1} \int_0^t v(s) ds \right) \quad (3.8)$$

can be viewed as a “vanishingly smooth” best response to the empirical distribution of play of one’s opponent.¹⁰ As such, (SRL-U) can be seen as a stochastically perturbed variant of vanishingly smooth fictitious play in continuous time, in which case Corollary 3.2 provides the continuous-time, stochastic analogue of Theorem 1.8 of Benaïm and Faure (2013).

4. EXTINCTION OF DOMINATED STRATEGIES

A fundamental rationality requirement for any game-theoretic learning process is the elimination of suboptimal, dominated strategies. Formally, given a finite game $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \mathcal{A}, u)$, we say that $p_k \in \mathcal{X}_k$ is dominated by $p'_k \in \mathcal{X}_k$ (and we write $p_k \prec p'_k$) if

$$\langle v_k(x), p_k \rangle < \langle v_k(x), p'_k \rangle \quad \text{for all } x \in \mathcal{X}. \quad (4.1)$$

Strategies that are *iteratively dominated*, *undominated*, or *iteratively undominated* are defined similarly; also, for pure strategies $\alpha, \beta \in \mathcal{A}_k$, we obviously have $\alpha \prec \beta$ if and only if

$$v_{k\alpha}(x) < v_{k\beta}(x) \quad \text{for all } x \in \mathcal{X}. \quad (4.2)$$

With this in mind, given a trajectory of play $x(t)$, $t \geq 0$, we will say that a pure strategy $\alpha \in \mathcal{A}_k$ becomes extinct along $x(t)$ if $x_{k\alpha}(t) \rightarrow 0$ as $t \rightarrow \infty$. More generally, following Samuelson and Zhang (1992), we will say that the mixed strategy $p_k \in \mathcal{X}_k$ becomes extinct along $x(t)$ if $\min\{x_{k\alpha}(t) : \alpha \in \text{supp}(p_k)\} \rightarrow 0$; otherwise, p_k is said to *survive*.

In the case of perfect payoff observations, Mertikopoulos and Sandholm (2016) showed that dominated strategies become extinct under the general reinforcement learning dynamics (RL), thus extending earlier results for the replicator dynamics (for an overview, see Viossat, 2015). Under noise and uncertainty however, the situation is a bit more complex: in the replicator dynamics with aggregate shocks

¹⁰It is “smooth” because the player is employing a regularized best response map (as opposed to a hard arg max correspondence), and it is “vanishingly smooth” because the factor $t\eta(t)$ hardens to ∞ as $t \rightarrow \infty$ (for a more detailed discussion, see Benaïm and Faure, 2013).

(ASRD), Cabrales (2000), Imhof (2005) and Hofbauer and Imhof (2009) provided a set of sufficient conditions on the intensity of the noise that guarantee the elimination of dominated strategies; on the other hand, Mertikopoulos and Moustakas (2010) showed that the noisy replicator dynamics (SRD) eliminate all strategies that are not iteratively undominated, with no conditions on the noise level.

As we show below, this unconditional elimination result is a consequence of the core reinforcement principle behind (RL) and it extends to the entire class of learning dynamics covered by (SRL):

Theorem 4.1. *Let $X(t)$ be a solution orbit of (SRL). If $p_k \in \mathcal{X}_k$ is dominated (even iteratively), then it becomes extinct along $X(t)$ almost surely.*

The basic idea of our proof is to show that the Itô process $X(t)$ has a dominant drift term that pushes it away from dominated strategies. However, given the complicated form of the (non-autonomous) dynamics (2.9), we do so by using the Fenchel coupling (3.5) to relate the evolution of the generating process $Y_k(t)$ to p_k .

Proof of Theorem 4.1. Suppose $p_k \prec p'_k$ so that $\langle v_k(x), p'_k - p_k \rangle \geq m_k$ for some $m_k > 0$ and for all $x \in \mathcal{X}$. Then, if $Y(t)$ is a solution orbit of (SRL), we get:

$$\begin{aligned} d\langle Y_k, p'_k - p_k \rangle &= \langle dY_k, p'_k - p_k \rangle = \langle v_k, p'_k - p_k \rangle dt + \sum_{\beta}^k (p'_{k\beta} - p_{k\beta}) \sigma_{k\beta} dW_{k\beta} \\ &\geq m_k dt + \sum_{\beta}^k (p'_{k\beta} - p_{k\beta}) \sigma_{k\beta} dW_{k\beta}, \end{aligned} \quad (4.3)$$

or, equivalently:

$$\langle Y_k(t), p'_k - p_k \rangle \geq c_k + m_k t + \xi_k(t), \quad (4.4)$$

where $c_k = \langle Y_k(0), p_k - p'_k \rangle$ and

$$\xi_k(t) = \sum_{\beta}^k (p_{k\beta} - p'_{k\beta}) \int_0^t \sigma_{k\beta}(X(s)) dW_{k\beta}(s). \quad (4.5)$$

Consider now the rate-adjusted ‘‘cross-coupling’’

$$\begin{aligned} V_k(y_k) &= \eta_k^{-1} [F_k(p_k, \eta_k y_k) - F_k(p'_k, \eta_k y_k)] \\ &= \eta_k^{-1} [h_k(p_k) - h_k(p'_k)] - \langle y_k, p_k - p'_k \rangle, \end{aligned} \quad (4.6)$$

with F_k defined as in (3.5). Then, by substituting (4.4) in (4.6) and recalling that $F_k(p'_k, y_k) \geq 0$, we obtain:

$$F_k(p_k, \eta_k Y_k) \geq h_k(p_k) - h_k(p'_k) + \eta_k \cdot [c_k + m_k t + \xi_k(t)]. \quad (4.7)$$

To proceed, Lemma B.4 shows that $m_k t + \xi_k(t) \sim m_k t$ (a.s.), so the RHS of (4.7) tends to infinity as $t \rightarrow \infty$ on account of the fact that $t\eta_k(t) \rightarrow +\infty$ (cf. Assumption 1). In turn, this gives $F_k(p_k, \eta_k(t)Y_k(t)) \rightarrow +\infty$, so p_k becomes extinct along $X(t) = Q(\eta(t)Y(t))$ by virtue of Proposition B.1. Finally, our claim for iteratively dominated strategies follows by induction on the rounds of elimination of dominated strategies – cf. Cabrales (2000, Proposition 1A). \square

Theorem 4.1 shows that dominated strategies become extinct under (SRL) but it does not give any information on the rate of extinction – or how probable it is to observe a dominated strategy above a given level at some $t \geq 0$. To address this issue, we provide two results below: Proposition 4.2 describes the long-term decay rate of dominated strategies and provides a ‘‘large deviations’’ bound for the probability of observing a dominated strategy above a given level at time $t \geq$

0. Subsequently, in Proposition 4.3 we estimate the expected time it takes for a dominated strategy to drop below a given level. For simplicity, we present our results in the case where the players' choice maps are derived from separable penalty functions as in (2.8); again, proofs are relegated to Appendix A:

Proposition 4.2. *Let $\alpha \in \mathcal{A}_k$ be dominated by $\beta \in \mathcal{A}_k$ and assume that the choice map of player k is generated by a separable penalty function of the form (2.8) with $\lim_{x \rightarrow 0^+} \theta'_k(x) = -\infty$. Then, for all $\delta, \varepsilon > 0$ and for all large enough t , we have:*

$$X_{k\alpha}(t) \leq \phi_k \left[C_k - \eta_k(t) \left(m_k t - 2(1 + \varepsilon) \sigma_{\alpha\beta} \sqrt{t \log \log t} \right) \right] \quad (\text{a.s.}) \quad (4.8)$$

and

$$\mathbb{P}(X_{k\alpha}(t) > \delta) \leq \frac{1}{2} \operatorname{erfc} \left[\frac{1}{2\sigma_{\alpha\beta}} \left(m_k \sqrt{t} - \frac{C_k - \theta'_k(\delta)}{\eta_k(t) \sqrt{t}} \right) \right] \quad (4.9)$$

where:

- a) $\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-t^2} dt$ is the complementary error function.
- b) $\phi_k = (\theta'_k)^{-1}$ (note that $\lim_{z \rightarrow -\infty} \phi_k(z) = 0$ by assumption).
- c) $m_k = \min_{x \in \mathcal{X}} \{v_{k\beta}(x) - v_{k\alpha}(x)\} > 0$ is the minimum payoff difference between α and β .
- d) $\sigma_{\alpha\beta}^2 = \frac{1}{2} \max_{x \in \mathcal{X}} \{ \sigma_{k\alpha}^2(x) + \sigma_{k\beta}^2(x) \} > 0$.
- e) C_k is a constant that depends only on the initial conditions of (SRL).

Proposition 4.3. *With notation as in Proposition 4.2, assume that (SRL) is run with constant learning rates η_k and noisy observations with constant variance. If $\tau_\delta = \inf\{t > 0 : X_{k\alpha}(t) \leq \delta\}$, then:*

$$\mathbb{E}[\tau_\delta] \leq \frac{[C_k - \theta'_k(\delta)]_+}{\eta_k m_k}. \quad (4.10)$$

Propositions 4.2 and 4.3 are our main results concerning the rate of elimination of dominated strategies under uncertainty, so a few remarks are in order:

Remark 4.1 (Asymptotics versus mean behavior). Propositions 4.2 and 4.3 capture complementary aspects of the statistics of the extinction rate of dominated strategies under (SRL). For instance, (4.8) describes the asymptotic rate of elimination of dominated strategies but it does not provide an estimate of how much time must pass until this rate becomes relevant. On the other hand, the bound (4.10) estimates the mean time it takes for a dominated strategy to fall below a given level; however, in contrast to (4.9), it does not describe how probable it is to observe deviations from this mean.

Remark 4.2 (Comparison with the noiseless case). When $\sigma = 0$, the decay estimate (4.8) boils down to $\phi_k[C_k - m_k \eta_k(t)t]$, a bound which recovers the results of Mertikopoulos and Sandholm (2016) for the deterministic dynamics (RL). Thus, even though the noise may initially mask the fact that a strategy is dominated, Proposition 4.2 shows that the long-run behavior of (SRL) and (RL) is the same as far as dominated strategies are concerned. This is in contrast with the aggregate-shocks dynamics (ASRD) where, even if a dominated strategy becomes extinct, its rate of elimination is different to leading order than in the noiseless case (Imhof, 2005, Theorem 3.1).

Remark 4.3 (The role of η). The estimate (4.8) shows that running (SRL) with large, constant η leads to a much faster rate of extinction of dominated strategies. In particular, recalling that $\operatorname{erfc}(t) \sim \pi^{-1/2}t^{-1}e^{-t^2}$ as $t \rightarrow \infty$, the bound (4.9) becomes:

$$\mathbb{P}(X_{k\alpha}(t) > \delta) = \mathcal{O}\left(t^{-1/2} \exp\left(-\frac{m_k^2 t}{2\sigma_{\alpha\beta}^2}\right)\right), \quad (4.11)$$

up to a subleading term of the order of $\mathcal{O}(1/\eta_k^2(t))$ in the exponent. Thus, even though the leading behavior of (4.9) is not affected by the player's choice of learning parameter, the subleading term is minimized for large, constant η so the asymptotic bound (4.11) becomes tighter in that case.

We thus observe an important contrast between regret minimization and the elimination of dominated strategies: whereas the optimal regret guarantee of Theorem 3.1 is achieved for $\eta(t) \propto 1/\sqrt{t}$, the asymptotic extinction rate of dominated strategies is much faster for constant η . The reason for this disparity is that higher values of η reinforce consistent payoff differences and therefore eliminate dominated strategies faster (independently of the noise level). On the other hand, to attain a no-regret state, players must be careful not to make too many mistakes in the presence of noise, so a more conservative choice of η is warranted.

5. LONG-TERM STABILITY AND CONVERGENCE ANALYSIS

We now turn to the long-term stability and convergence properties of the dynamics (SRL) with respect to equilibrium play. To that end, recall first that $x^* \in \mathcal{X}$ is a *Nash equilibrium* of $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \mathcal{A}, u)$ if it is unilaterally stable for all players, i.e.

$$u_k(x^*) \geq u_k(x_k, x_{-k}^*) \quad \text{for all } x_k \in \mathcal{X}_k, k \in \mathcal{N}, \quad (5.1)$$

or, equivalently:

$$v_{k\alpha}(x^*) \geq v_{k\beta}(x^*) \quad \text{for all } k \in \mathcal{N} \text{ and for all } \alpha \in \operatorname{supp}(x_k^*), \beta \in \mathcal{A}_k. \quad (5.2)$$

Strict equilibria are defined by requiring that (5.1) hold as a strict inequality for all $x_k \neq x_k^*$; obviously, such equilibria are also *pure* in the sense that they correspond to pure strategy profiles in $\mathcal{A} = \prod_k \mathcal{A}_k$ (i.e. vertices of \mathcal{X}).

In the noiseless case ($\sigma = 0$) with constant learning rates ($\dot{\eta} = 0$), Mertikopoulos and Sandholm (2016) recently showed that the deterministic dynamics (RL) exhibit the following properties with respect to Nash equilibria of \mathcal{G} :

- a) If a solution orbit of (RL) converges to x^* , then x^* is a Nash equilibrium.
- b) If $x^* \in \mathcal{X}$ is (Lyapunov) stable, then it is also a Nash equilibrium.
- c) Strict Nash equilibria are asymptotically stable in (RL).

In turn, these properties are generalizations of the long-term stability and convergence properties of the (multi-population) replicator dynamics – sometimes referred to as the “folk theorem” of evolutionary game theory (Hofbauer and Sigmund, 1998, 2003). That being said, the situation is quite different in the presence of noise: for instance, interior Nash equilibria are not even traps (almost sure rest points) of the stochastic reinforcement learning dynamics (SRL), so the ordinary (deterministic) definitions of stability and convergence no longer apply. Instead, in the context of stochastic differential equations, Lyapunov and asymptotic stability are defined as follows (Khasminskii, 2012):

Definition 5.1. Let $x^* \in \mathcal{X}$. We will say that:

- (1) x^* is *stochastically (Lyapunov) stable* under (SRL) if, for every $\varepsilon > 0$ and for every neighborhood U_0 of x^* in \mathcal{X} , there exists a neighborhood $U \subseteq U_0$ of x^* such that

$$\mathbb{P}(X(t) \in U_0 \text{ for all } t \geq 0) \geq 1 - \varepsilon, \quad (5.3)$$

whenever $X(0) \in U$.

- (2) x^* is *stochastically asymptotically stable* under (SRL) if it is stochastically stable and attracting: for every $\varepsilon > 0$ and for every neighborhood U_0 of x^* in \mathcal{X} , there exists a neighborhood $U \subseteq U_0$ of x^* such that

$$\mathbb{P}(X(t) \in U_0 \text{ for all } t \geq 0 \text{ and } \lim_{t \rightarrow \infty} X(t) = x^*) \geq 1 - \varepsilon, \quad (5.4)$$

whenever $X(0) \in U$.

In the evolutionary setting of the stochastic replicator dynamics with aggregate shocks, Imhof (2005) and Hofbauer and Imhof (2009) showed that strict Nash equilibria are stochastically asymptotically stable under (ASRD) provided that the variability of the shocks across different strategies is small enough. More recently, in a learning context, Mertikopoulos and Moustakas (2010) showed that the same holds for the stochastic replicator dynamics (SRD) of exponential learning (with constant η), irrespective of the variance of the observation noise. However, this last result relies heavily on the specific properties of the logit map (2.6) and the infinitesimal generator of (SRD).

In our case, the convoluted (and non-autonomous) form of the stochastic dynamics (2.9) complicates things considerably, so such an approach is not possible – especially with regards to finding a local stochastic Lyapunov function for the dynamical system (2.9) that governs the evolution of $X(t)$. Nonetheless, by working *directly* with (SRL), we obtain the following general result:

Theorem 5.2. *Let $X(t)$ be a solution orbit of (SRL) and let $x^* \in \mathcal{X}$. Then:*

- (1) *If $\mathbb{P}(\lim_{t \rightarrow \infty} X(t) = x^*) > 0$, x^* is a Nash equilibrium of \mathcal{G} .*
- (2) *If x^* is stochastically (Lyapunov) stable, it is also Nash.*
- (3) *If x^* is a strict Nash equilibrium of \mathcal{G} , it is stochastically asymptotically stable under (SRL).*

The first ingredient of our proof (presented in detail in Appendix A) is to show that if the process $X(t)$ remains in the vicinity of $x^* \in \mathcal{X}$ for all $t \geq 0$ with positive probability, then x^* must be a Nash equilibrium. This is formalized in the following proposition (which is of independent interest):

Proposition 5.3. *With notation as in Theorem 5.2, assume that every neighborhood U of x^* in \mathcal{X} admits with positive probability a solution orbit $X(t)$ of (SRL) such that $X(t) \in U$ for all $t \geq 0$. Then, x^* is a Nash equilibrium.*

Proof. If x^* is not Nash, we must have $v_{k\alpha}(x^*) < v_{k\beta}(x^*)$ for some player $k \in \mathcal{N}$ and for some $\alpha \in \text{supp}(x_k^*)$, $\beta \in \mathcal{A}_k$. On that account, let U be a sufficiently small neighborhood of x^* in \mathcal{X} such that $v_{k\beta}(x) - v_{k\alpha}(x) \geq m_k$ for some $m_k > 0$ and for all $x \in U$. Then, conditioning on the positive probability event that there exists an orbit $X(t) = Q(\eta(t)Y(t))$ of (SRL) that is contained in U for all $t \geq 0$, we have:

$$\begin{aligned} dY_{k\alpha} - dY_{k\beta} &= (v_{k\alpha}(X) - v_{k\beta}(X)) dt + \sigma_{k\alpha} dW_{k\alpha} - \sigma_{k\beta} dW_{k\beta} \\ &\leq -m_k dt - d\xi_k, \end{aligned} \quad (5.5)$$

where $\xi_k(t) = \int_0^t \sigma_{k\beta}(X(s)) dW_{k\beta}(s) - \int_0^t \sigma_{k\alpha}(X(s)) dW_{k\alpha}(s)$. Since the diffusion coefficients σ are bounded, the process $\xi_k(t)$ will grow at most as $\mathcal{O}(\sqrt{t \log \log t})$, so we get $\xi_k(t) + m_k t \sim m_k t$ (for a precise statement, see Lemma B.4). Thus, on account of Assumption 1 for $\eta_k(t)$, we obtain:

$$\eta_k(t) \cdot [Y_{k\alpha}(t) - Y_{k\beta}(t)] \leq -\frac{m_k t + \xi_k(t)}{t} \cdot \eta_k(t) t \rightarrow -\infty, \quad (5.6)$$

The above shows that the score difference between α and β grows to $-\infty$, so $X_{k\alpha}(t)$ tends to zero as $t \rightarrow \infty$ (Proposition B.1). This contradicts the assumption that $X(t) \in U$ for all t (recall that $\alpha \in \text{supp}(x_k^*)$), so x^* must be a Nash equilibrium of \mathcal{G} . \square

With Proposition 5.3 at hand, we are left to show that strict equilibria are stochastically asymptotically stable under (SRL). This is carried out in Appendix A by showing that *a*) trajectories that remain close to a strict equilibrium are eventually attracted to it (because it is locally dominant); and *b*) that this occurs with controllably high probability (a result which is harder to establish and which relies on an application of Girsanov's theorem to estimate the probability that a Brownian motion with positive drift attains a given negative level in finite time).

We close this section with a few remarks on Theorem 5.2:

Remark 5.1 (Convergence rates and the role of η). By reasoning as in Proposition 4.2, it is also possible to show that the rate of convergence of (SRL) to strict equilibria increases with the players' learning parameter η . As in the case of dominated strategies, the reason for this is that players experience a consistent payoff difference in favor of their equilibrium strategy when close to a strict equilibrium, and this difference (which is reflected in the score process $Y(t)$) is heightened by η . In this regard, convergence to a strict equilibrium seems to be somewhat antagonistic to attaining a no-regret state in the presence of noise: the regret guarantee (3.4) is optimized for $\eta(t) \propto t^{-1/2}$ while the dynamics' rate of convergence to strict equilibria is faster for large, constant η .

Remark 5.2. By considering the modified game $\tilde{\mathcal{G}}$ with noise-adjusted payoff functions $\tilde{v}_{k\alpha}(x) = v_{k\alpha}(x) - \frac{1}{2}\sigma_{k\alpha}^2$, the reasoning in the proof of Theorem 5.2 yields an alternative proof of the stability and convergence results of Hofbauer and Imhof (2009) for the replicator dynamics with aggregate shocks (ASRD).

6. TIME AVERAGES IN 2-PLAYER GAMES

In this section, we analyze the asymptotic behavior of the time averages

$$\bar{X}(t) = \frac{1}{t} \int_0^t X(s) ds, \quad (6.1)$$

and we establish an averaging principle for (SRL) in 2-player games. Our analysis is motivated by the original deterministic results of Hofbauer and Sigmund (1998) and Hofbauer et al. (2009) who showed that $\bar{X}(t)$ converges to Nash equilibrium under the replicator dynamics whenever $\liminf_{t \rightarrow \infty} X_{k\alpha}(t) > 0$ for all $\alpha \in \mathcal{A}_k$, $k = 1, 2$. More recently, Mertikopoulos and Sandholm (2016) proved a version of this result for arbitrary regularized best response maps (always in a deterministic setting), while Hofbauer and Imhof (2009) showed that time-averages of the aggregate-shocks replicator dynamics (ASRD) converge to the Nash set of a modified game.

Our main result here is that the averaging principle of Hofbauer and Sigmund (1998) extends to the perturbed reinforcement learning scheme (SRL), even in the presence of arbitrarily large payoff observation errors:

Theorem 6.1. *Let \mathcal{G} be a 2-player game and let $X(t)$ be a solution orbit of the stochastic dynamics (SRL). If the players' score differences $Y_{k\alpha}(t) - Y_{k\beta}(t)$ grow sublinearly with t for all $\alpha, \beta \in \mathcal{A}_k$, $k = 1, 2$, then the time-averaged process $\bar{X}(t)$ converges almost surely to the set of Nash equilibria of \mathcal{G} .*

Proof of Theorem 6.1. Pick $\alpha, \beta \in \mathcal{A}_k$. Then, by the definition of (SRL), we get:

$$\begin{aligned} Y_{k\alpha}(t) - Y_{k\beta}(t) &= Y_{k\alpha}(0) - Y_{k\beta}(0) + \int_0^t [v_{k\alpha}(X(s)) - v_{k\beta}(X(s))] ds \\ &\quad + \int_0^t \sigma_{k\alpha}(X(s)) dW_{k\alpha}(s) - \int_0^t \sigma_{k\beta}(X(s)) dW_{k\beta}(s) \\ &= c_k + \xi_k(t) + t [v_{k\alpha}(\bar{X}(t)) - v_{k\beta}(\bar{X}(t))], \end{aligned} \quad (6.2)$$

where $c_k = Y_{k\alpha}(0) - Y_{k\beta}(0)$, $\xi_k(t)$ denotes the martingale part of (6.2) and we have used the fact that $v_k(x)$ is linear in x (and not only multilinear). Since $\xi_k(t)$ is subleading in (6.2) with respect to t (cf. Lemma B.4 for a precise statement), dividing by t yields:

$$\lim_{t \rightarrow \infty} [v_{k\alpha}(\bar{X}(t)) - v_{k\beta}(\bar{X}(t))] = 0 \quad (\text{a.s.}), \quad (6.3)$$

so $v_{k\alpha}(x^*) = v_{k\beta}(x^*)$ whenever x^* is an ω -limit of $\bar{X}(t)$. This shows that any ω -limit of $\bar{X}(t)$ is a Nash equilibrium, and since the ω -set of $\bar{X}(t)$ is nonempty (by compactness of \mathcal{X}), our claim follows. \square

In the case of (EW)/(RD), the sublinear growth requirement for $Y_{k\alpha} - Y_{k\beta}$ boils down to the permanency condition $\liminf_{t \rightarrow \infty} X_{k\alpha}(t) > 0$, so we recover the original result of Hofbauer and Sigmund (1998). Other than that however, the applicability of Theorem 6.1 is limited by the growth requirement for $Y_{k\alpha}(t) - Y_{k\beta}(t)$. The following proposition (whose proof can be found in Appendix A) shows that this condition always holds in 2-player zero-sum games under the additional assumption that $\lim_{t \rightarrow \infty} \eta_k(t) = 0$:

Proposition 6.2. *Let \mathcal{G} be a 2-player zero-sum game with an interior Nash equilibrium, and assume that (SRL) is run with $\eta_k(t)$ satisfying $\lim_{t \rightarrow \infty} \eta_k(t) = 0$ (in addition to Assumption 1). Then, the time average of (SRL) converges to the set of Nash equilibria of \mathcal{G} (a.s.).*

Proposition 6.2 is reminiscent of Theorem 3.1 on the no-regret properties of (SRL) in that both results require a vanishing learning parameter. In the noiseless regime however ($\sigma = 0$), this is not the case: in the replicator dynamics (which correspond to (EW) with $\eta = 1$), the conclusions of Proposition 6.2 still hold essentially verbatim (Hofbauer and Sigmund, 1998). Intuitively, the reason for this is that time-averaged replicator orbits exhibit the same long-term behavior as the (deterministic) best response dynamics of Gilboa and Matsui (1991):

$$\dot{x}_k \in \text{br}_k(x) - x_k, \quad (\text{BRD})$$

where $\text{br}_k(x) \equiv \arg \max_{x'_k \in \mathcal{X}_k} \langle v_k(x), x'_k \rangle$ denotes the standard (non-regularized) best response correspondence of player k . Thus, given that the dynamics (BRD)




FIGURE 2. Time averages of (SRL) with logit best responses in a game of Matching Pennies (as in Fig. 1, Nash equilibria are depicted in red and the game’s payoff’s are displayed inline; for benchmarking purposes, we also took $\sigma_{k\alpha} = 1$ for all $\alpha \in \mathcal{A}_k$, $k = 1, 2$). Fig. 2(a) shows the evolution of a sample trajectory and its time average; in Fig. 2(b), we show a density plot of the distribution of 10^4 time-averaged trajectories for different values of the integration horizon T . In tune with Proposition 6.2, we see that time averages converge to the game’s Nash equilibrium.

converge to equilibrium in zero-sum games, the same must also hold for the time averages of (RD).

More precisely, Hofbauer et al. (2009) showed that the ω -limit set Ω of time-averaged solutions of (RD) is *internally chain transitive* (ICT) under (BRD), i.e. any two points in Ω may be joined by a piecewise continuous “chain” of arbitrarily long orbit segments of (BRD) in Ω broken by arbitrarily small jump discontinuities.¹¹ Mertikopoulos and Sandholm (2016) subsequently established a similar averaging principle linking the deterministic reinforcement learning scheme (RL) to (BRD). Importantly, as we show below, the *stochastic* dynamics (SRL) share the same connection to the *deterministic* best response dynamics (BRD), despite all the noise:

Theorem 6.3. *Let $X(t)$ be a solution orbit of (SRL) for a 2-player game \mathcal{G} . Then, the ω -limit set of $\bar{X}(t)$ is internally chain transitive under the deterministic best response dynamics (BRD).*

Proof. Our proof follows the approach of Hofbauer et al. (2009) and is presented in detail in Appendix A. \square

Thanks to Theorem 6.3, several conclusions of Hofbauer et al. (2009) for 2-player games can be readily generalized to the full *stochastic* setting of (SRL) simply by exploiting the properties of the *deterministic* dynamics (BRD):

¹¹In particular, ICT sets are invariant, connected and have no proper attractors; for a full development, see Benaïm, Hofbauer, and Sorin, (2005).

- (1) If $\bar{X}(t)$ converges under (SRL), its limit is a Nash equilibrium.
- (2) Any global attractor of (BRD) also attracts time averages of (SRL), independently of the noise level. In particular, since the set of Nash equilibria is globally attracting under (BRD) in zero-sum games, this observation extends Proposition 6.2 to the constant η case.
- (3) The only ICT sets of (BRD) in potential games consist of (isolated) components of Nash equilibria, so $\bar{X}(t)$ converges to one such component.

7. DISCUSSION

7.1. Correlated observation noise. An important assumption underlying the noisy learning model (SRL) is that the perturbations to the players' payoff observations are stochastically independent across players and strategies. However, in many cases of practical interest, such observations are actually *correlated*: in congestion games for example, a choice of action corresponds to a choice of route in a road network, and two routes that overlap would have the same uncertainty characteristics over common edges/road segments (and similarly for correlations across different players).

To account for correlations of this sort, we can consider the more general learning model

$$\begin{aligned} dY_{k\alpha} &= v_{k\alpha}(X) dt + dZ_{k\alpha}, \\ X_k &= Q_k(\eta_k Y_k), \end{aligned} \tag{SRL_c}$$

where $Z_{k\alpha}$ is a family of Itô martingale processes with quadratic variation

$$d[Z_{k\alpha}, Z_{j\beta}] = dZ_{k\alpha} \cdot dZ_{j\beta} = \Sigma_{\alpha\beta}^{kj} dt \quad k, j \in \mathcal{N}, \alpha \in \mathcal{A}_k, \beta \in \mathcal{A}_j, \tag{7.1}$$

and such that $\sup_t \|\Sigma(t)\| < 0$. Obviously, if Σ is block-diagonal (i.e. $\Sigma_{\alpha\beta}^{kj} \propto \delta_{\alpha\beta} \delta_{kj}$), we recover the original learning model (SRL) with uncorrelated uncertainty across strategies and payoffs; otherwise, in the non-diagonal case, the covariance matrix Σ simply captures the possible correlations between the players' observations.

A first example of this type is provided by Cabrales (2000) where $Z_{k\alpha}$ is assumed to be of the form $dZ_{k\alpha} = \sum_{r=1}^d \sigma_{k\alpha}^r dW_r$ with $W(t)$ denoting a standard Wiener process in \mathbb{R}^d . Such correlations also occur naturally in the study of congestion games: indeed, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a (non-oriented) graph with vertex set \mathcal{V} and edge set \mathcal{E} , and let $\mathcal{N} \subseteq \mathcal{V} \times \mathcal{V}$ denote a set of origin-destination pairs in the network. In (finite, nonsplittable) congestion games, each player k is typically identified with an origin-destination pair $(v, w) \in \mathcal{N} \subseteq \mathcal{V} \times \mathcal{V}$, and his actions are drawn from a set \mathcal{A}_k of (acyclic) paths in \mathcal{G} that join v to w . The cost associated to a choice of path $\alpha \in \mathcal{A}_k$ is then given by the total time taken to traverse path α , i.e. $\ell_{k\alpha} = \sum_{r \in \alpha} \ell_r = \sum_{r \in \mathcal{E}} P_{r\alpha} \ell_r$, where ℓ_r denotes the corresponding delay function of edge $r \in \mathcal{E}$ (a function of the number of users who use $r \in \mathcal{E}$) and $P_{r\alpha} = \mathbf{1}\{r \in \alpha\}$ denotes the edge-path incidence matrix of the network. As a result, if the delay across each edge is only observable up to a random error, the players' score updates would take the form:

$$dY_{k\alpha} = - \sum_{r \in \alpha} (\ell_r(X) dt + \sigma_r dW_r) = -\ell_{k\alpha}(X) dt + dZ_{k\alpha}, \tag{7.2}$$

with $dZ_{k\alpha} = -\sum_{r \in \mathcal{E}} P_{r\alpha} \sigma_r dW_r$. In turn, this leads to the covariance matrix

$$\begin{aligned} d[Z_{k\alpha}, Z_{j\beta}] &= \sum_{r, r' \in \mathcal{E}} P_{r\alpha} P_{r'\beta} \sigma_r \sigma_{r'} dW_r \cdot dW_{r'} \\ &= \sum_{r \in \mathcal{E}} P_{r\alpha} P_{r\beta} \sigma_r^2 dt = \sum_{r \in \alpha \cap \beta} \sigma_r^2 dt, \end{aligned} \quad (7.3)$$

i.e. stochastic fluctuations across different paths α, β are correlated precisely over the paths' common edges.

The question that arises in this context is whether the rationality analysis of the previous sections still applies to the more general model (SRL_c) with correlated uncertainty. The short answer to this is that the qualitative part of our analysis goes through essentially unchanged:

Proposition 7.1. *Assume that (SRL_c) is run with a vanishing learning parameter $\eta(t)$. Then, with notation as in the rest of the paper:*

1. (SRL_c) leads to no regret.
2. Dominated strategies become extinct under (SRL_c).
3. Strict Nash equilibria are stochastically asymptotically stable under (SRL_c).
4. In 2-player zero-sum games with an interior equilibrium, the time-averaged process $\bar{X}(t) = t^{-1} \int_0^t X(s) ds$ converges to the game's set of Nash equilibria.

Finally, properties (2) and (3) above hold even without the assumption $\eta(t) \rightarrow 0$.

The proof of Proposition 7.1 is similar to that of Theorems 3.1, 4.1, 5.2, and Proposition 6.2, so we only offer a short sketch in Appendix A. On the other hand, the quantitative bounds that we derived (e.g. for the player's regret or for the rate of elimination of dominated strategies) clearly cannot hold as stated because of the non-diagonal components of the quadratic variation matrix Σ . It is still possible to calculate explicitly the corresponding expressions under the correlated noise model (7.1) but the end result is fairly messy so we do not present it here.

7.2. Abrupt shocks. Another important assumption underlying the noisy learning model (SRL) (or its correlated version above) is that uncertainty enters the process in the form of a continuous, non-anticipative diffusion process. This model is sufficiently wide to cover the case of noise processes whose aggregate impact scales as a function of the observation window but it does not account for discontinuous jump processes that represent very abrupt perturbations (such as defaults and stock price crashes, momentary link failures in data networks, etc.).

Processes of this type are usually modeled in probability theory as *semimartingales*, i.e. stochastic processes that can be decomposed as a local martingale (e.g. a Wiener process) plus an adapted process with finite variation (e.g. any piecewise continuous martingale). Given that the Itô integral can also be defined for semimartingales, it would also be possible to extend the scope of (SRL) even further by replacing the Brownian noise terms dW by an arbitrary (adapted) semimartingale. In this case however, Itô's lemma would also carry additional terms stemming from the discontinuous component of the noise process, thus complicating the analysis considerably; because of this, we relegate a detailed discussion of this issue to a future paper.

APPENDIX

APPENDIX A. TECHNICAL PROOFS

Caveat. In this appendix (and, in particular, in all applications of Itô's lemma), we will be treating the players' choice map Q as if it were C^1 . This is always the case if the players' penalty functions are steep (i.e. $\|dh(x_n)\| \rightarrow \infty$ whenever $x_n \rightarrow \text{bd}(\mathcal{X})$), but not otherwise: for instance, the closest-point projection map induced by the quadratic penalty function $h(x) = \frac{1}{2}\|x\|^2$ is piecewise smooth but not C^1 . The use of Itô's lemma in the non-steep case is justified by the fact that Q is absolutely continuous (Rockafellar, 1970, Chap. 26) and the work of Alsmeyer and Jaeger (2005, Theorem 2.1) who provide a nonsmooth variant of Itô's lemma that allows us to treat h^* as if it were C^2 whenever it appears in a stochastic integral.

A.1. Proof of Proposition 2.1. For simplicity, the player index k is suppressed in this proof; also, all summation indices are assumed to run over the (constant) support \mathcal{A}' of $X(t)$.

By the Karush–Kuhn–Tucker (KKT) conditions for the convex problem (2.4), we get $Y_\alpha - \eta^{-1}\theta'(X_\alpha) = \zeta$ where ζ is the Lagrange multiplier corresponding to the probability constraint $\sum_{\alpha \in \text{supp}(X)} X_\alpha = 1$. Itô's formula then gives

$$d\zeta = dY_\alpha + \frac{\dot{\eta}}{\eta^2}\theta'_\alpha dt - \eta^{-1}\theta''_\alpha dX_\alpha - \frac{1}{2}\eta^{-1}\theta'''_\alpha (dX_\alpha)^2 \quad (\text{A.1})$$

and hence:

$$dX_\alpha = \frac{\eta}{\theta''_\alpha} \left[dY_\alpha + \frac{\dot{\eta}}{\eta^2}\theta'_\alpha dt - \frac{1}{2\eta}\theta'''_\alpha (dX_\alpha)^2 - d\zeta \right]. \quad (\text{A.2})$$

Now, write dX_α in the general form:

$$dX_\alpha = b_\alpha dt + \sum_\beta c_{\alpha\beta} dW_\beta, \quad (\text{A.3})$$

where the drift and diffusion coefficients b_α and $c_{\alpha\beta}$ are to be determined. To do so, recall that the Wiener processes W_α are independent (in the sense that $dW_\alpha \cdot dW_\beta = \delta_{\alpha\beta}$), so we get $(dX_\alpha)^2 = \sum_\beta c_{\alpha\beta}^2 dt$; therefore, summing (A.2) over $\alpha \in \mathcal{A}'$ and solving for $d\zeta$ yields:

$$\frac{1}{\Theta''} d\zeta = \sum_\alpha \frac{1}{\theta''_\alpha} \left[dY_\alpha + \frac{\dot{\eta}}{\eta^2}\theta'_\alpha dt - \frac{1}{2\eta}\theta'''_\alpha \sum_\beta c_{\alpha\beta}^2 dt \right]. \quad (\text{A.4})$$

Substituting this last expression in (A.2) leads to (2.9) with $U_\alpha^2 = \sum_\beta c_{\alpha\beta}^2$, so we are left to show that U_α^2 has the prescribed form. To that end, by comparing the diffusion terms of (2.9) and (A.3), we get

$$c_{\alpha\beta} = \frac{\eta}{\theta''_\alpha} \left[\sigma_\alpha \delta_{\alpha\beta} - \frac{\Theta''}{\theta''_\beta} \sigma_\beta \right], \quad (\text{A.5})$$

and hence:

$$\begin{aligned} U_\alpha^2 &= \sum_\beta c_{\alpha\beta}^2 = \left(\frac{\eta}{\theta''_\alpha} \right)^2 \sum_\beta (\sigma_\alpha \delta_{\alpha\beta} - \Theta'' \sigma_\beta / \theta''_\beta)^2 \\ &= \left(\frac{\eta}{\theta''_\alpha} \right)^2 \left[\sigma_\alpha^2 (1 - \Theta'' / \theta''_\alpha)^2 + \sum_{\beta \neq \alpha} (\Theta'' / \theta''_\beta)^2 \sigma_\beta^2 \right], \end{aligned} \quad (\text{A.6})$$

which concludes our derivation of (2.9).

To show that the support of $X(t)$ is piecewise constant on a dense open subset of $[0, \infty)$, fix some $\alpha \in \mathcal{A}$ and let $A = \{t : X_\alpha(t) > 0\}$ so that $A^c = X_\alpha^{-1}(0)$. Then, A is open because $X(t)$ is continuous (a.s.), so it suffices to show that $A \cup \text{int}(A^c)$ is dense in $[0, \infty)$; however, this is trivially true because of the identity $\text{cl}(A) \cup \text{int}(A^c) = [0, \infty)$. Finally, if $\lim_{x \rightarrow 0^+} \theta'(x) = -\infty$, standard convex analysis arguments show that the (necessarily unique) solution of (2.4) lies in the relative interior of \mathcal{X} (Rockafellar, 1970, Chap. 26), so we conclude that $X(t) \in \mathcal{X}^\circ$ for all $t \geq 0$ by the well-posedness of (SRL).

A.2. Proof of Theorem 3.1. Our proof hinges on the rate-adjusted Fenchel coupling

$$H_p \equiv \frac{1}{\eta} F(p, \eta Y) = \frac{1}{\eta} \cdot [h(p) + h^*(\eta Y) - \langle \eta Y, p \rangle] \quad (\text{A.7})$$

between a benchmark strategy $p \in \mathcal{X}$ and the generating process ηY . Specifically, to begin with, the Itô formula of Lemma B.3 gives:

$$\begin{aligned} dH_p &= -\frac{\dot{\eta}}{\eta} H_p dt + \frac{1}{\eta} \langle d(\eta Y), X - p \rangle + \frac{1}{2\eta} \sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \eta^2 \sigma_{\beta}^2 dt \\ &= -\frac{\dot{\eta}}{\eta} H_p dt + \frac{\dot{\eta}}{\eta} \langle Y, X - p \rangle dt + \langle dY, X - p \rangle + \frac{\eta}{2} \sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2 dt, \end{aligned} \quad (\text{A.8})$$

so, by combining the definition of H_p with (SRL-U), we get:

$$\begin{aligned} dH_p &= -\frac{\dot{\eta}}{\eta^2} [h(p) - h(X)] dt + \langle v, X - p \rangle dt \\ &\quad + \sum_{\beta} (X_{\beta} - p_{\beta}) \sigma_{\beta} dW_{\beta} + \frac{\eta}{2} \sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2 dt, \end{aligned} \quad (\text{A.9})$$

where we used the fact that $h^*(\eta Y) = \langle \eta Y, X \rangle - h(X)$ in the first line. Therefore, the player's cumulative regret for not playing p up to time t will be:

$$\int_0^t \langle v(s), p - X(s) \rangle ds = H_p(0) - H_p(t) \quad (\text{A.10a})$$

$$- \int_0^t \frac{\dot{\eta}(s)}{\eta^2(s)} [h(p) - h(X(s))] ds \quad (\text{A.10b})$$

$$+ \sum_{\beta} \int_0^t (X_{\beta}(s) - p) \sigma_{\beta}(s) dW_{\beta}(s) \quad (\text{A.10c})$$

$$+ \frac{1}{2} \sum_{\beta} \int_0^t \eta(s) \frac{\partial^2 h^*}{\partial^2 y_{\beta}^2} \sigma_{\beta}^2(s) ds. \quad (\text{A.10d})$$

We now proceed to bound each term of (A.10) by a sublinear function:

a) Since $H_p \geq 0$, the term (A.10a) is bounded from above as follows:

$$H_p(0) \leq \frac{h(p) + h^*(0)}{\eta(0)} = \frac{h(p) - \min_{x \in \mathcal{X}} h(x)}{\eta(0)} \leq \frac{\Omega}{\eta(0)}, \quad (\text{A.11})$$

where we have used the fact that $Y(0) = 0$.

b) For (A.10b), we have $h(p) - h(X(s)) \leq \Omega$ by definition; hence, with $\dot{\eta} \leq 0$, we get:

$$(\text{A.10b}) \leq -\Omega \int_0^t \frac{\dot{\eta}(s)}{\eta^2(s)} ds = \frac{\Omega}{\eta(t)} - \frac{\Omega}{\eta(0)} = o(t). \quad (\text{A.12})$$

- c) By Lemma B.4, we readily obtain that the term (A.10c) is $\mathcal{O}(\sigma_{\max}\sqrt{t\log\log t})$.
d) Finally, for (A.10d), recall that the Hessian of $h^*(y)$ is equal to the inverse Hessian of $h(x)$, suitably restricted to the affine hull of the face of the simplex spanned by $x = Q(y)$ – see e.g. Chap. 26 in Rockafellar (1970). Thus, given that h is K -strongly convex, (2.5) readily yields $\partial^2 h^*/\partial y_\beta^2 \leq K^{-1}$, and hence:

$$(A.10d) \leq \frac{|A|}{2K} \sigma_{\max}^2 \int_0^t \eta(s) ds, \quad (A.13)$$

where $\sigma_{\max}^2 = \sup_{t \geq 0} \max_\beta \sigma_\beta^2(t)$. L'Hôpital's rule then yields $t^{-1} \int_0^t \eta(s) ds \sim \eta(t) \rightarrow 0$, so (A.10d) is also sublinear in t .

The bound (3.4) then follows by combining all of the above.

A.3. Proof of Propositions 4.2 and 4.3. We begin with some groundwork that will be used in the proof of both propositions. Specifically, let $X(t)$ be a solution of (SRL) with $X(0) = x \in \mathcal{X}$. Then, suppressing the player index k for simplicity, we obtain:

$$\begin{aligned} dY_\alpha - dY_\beta &= (v_\alpha(X) - v_\beta(X)) dt + \sigma_\alpha(X) dW_\alpha - \sigma_\beta(X) dW_\beta \\ &\leq -m dt - d\xi, \end{aligned} \quad (A.14)$$

where $\xi(t) = \int_0^t \sigma_\beta(X(s)) dW_\beta(s) - \int_0^t \sigma_\alpha(X(s)) dW_\alpha(s)$ and m is defined as in the statement of the proposition. We thus get:

$$Y_\alpha(t) - Y_\beta(t) \leq Y_\alpha(0) - Y_\beta(0) - mt - \xi(t), \quad (A.15)$$

so the KKT conditions for the problem (2.4) give:

$$\theta'(X_\alpha(t)) - \theta'(X_\beta(t)) \leq \eta(t) \cdot [Y_\alpha(0) - Y_\beta(0) - mt - \xi(t)] \quad (A.16)$$

(recall that θ is steep so the solutions of (2.4) are interior). Thus, with $\eta(t)$ decreasing, we finally get:

$$\theta'(X_\alpha(t)) \leq C - \eta(t) \cdot [mt + \xi(t)], \quad (A.17)$$

where $C = \theta'(1) + \eta(0) \cdot |Y_\alpha(0) - Y_\beta(0)|$ depends only on the initial conditions of (SRL). This bound will be central in both proofs.

Proof of Proposition 4.2. Let $\rho = [\xi, \xi]$ denote the quadratic variation of ξ (cf. the proof of Lemma B.4). Then, the independence of W_α and W_β yields $d\rho = d\xi \cdot d\xi = (\sigma_\alpha^2 + \sigma_\beta^2) dt$ and hence:

$$\rho(t) = \int_0^t [\sigma_\alpha^2(X(s)) + \sigma_\beta^2(X(s))] ds \leq 2\sigma_{\alpha\beta}^2 t. \quad (A.18)$$

Consequently, invoking the time-change theorem for martingales (Øksendal, 2007), let \widetilde{W} be a Wiener process such that $\xi(t) = \widetilde{W}(\rho(t))$. Then, if $\lim_{t \rightarrow \infty} \rho(t) = \infty$, by the law of the iterated logarithm, we get:

$$\xi(t) = \widetilde{W}(\rho(t)) \leq (1 + \varepsilon) \sqrt{2\rho(t) \log \log \rho(t)} \leq 2(1 + \varepsilon) \sigma_{\alpha\beta} \sqrt{t \log \log t} \quad (a.s.), \quad (A.19)$$

for all $\varepsilon > 0$ and for all large enough t . Otherwise, if $\lim_{t \rightarrow \infty} \rho(t) < \infty$, $\widetilde{W}(\rho(t))$ will be bounded for all $t \geq 0$ (a.s.), so we trivially get $\xi(t) \geq 2(1 + \varepsilon) \sigma_{\alpha\beta} \sqrt{t \log \log t}$ for all large enough $t > 0$. The bound (4.8) is then obtained by combining (A.17) with (A.19).

Likewise, combining (A.17) and (A.18), we obtain:

$$\begin{aligned} \mathbb{P}(X_\alpha(t) > \delta) &= \mathbb{P}(\theta'(X_\alpha(t)) > \theta'(\delta)) \\ &\leq \mathbb{P}\left(\widetilde{W}(\rho(t)) < -mt + \frac{C - \theta'(\delta)}{\eta(t)}\right) \\ &= \frac{1}{2} \operatorname{erfc}\left[\frac{1}{\sqrt{2\rho(t)}}\left(mt - \frac{C - \theta'(\delta)}{\eta(t)}\right)\right]. \end{aligned} \quad (\text{A.20})$$

Since $m > 0$, Assumption 1 for the rate of decay of $\eta(t)$ guarantees that $mt\eta(t) > C - \theta'(\delta)$ for large t ; (4.9) then follows by substituting (A.18) in (A.20). \square

Proof of Proposition 4.3. Let $Z_{\alpha\beta}$ denote the RHS of (A.14), viz.:

$$dZ_{\alpha\beta} = m dt + \sigma_\beta dW_\beta - \sigma_\alpha dW_\alpha. \quad (\text{A.21})$$

Then, by (A.17), we readily obtain

$$\tau_\delta \leq \tilde{\tau}_a \equiv \inf\{t > 0 : Z_{\alpha\beta}(t) \leq a\}, \quad (\text{A.22})$$

with $a = \eta^{-1}[C - \theta'(\delta)]_+$ (obviously, $\tau_\delta = 0$ if $C - \theta'(\delta) < 0$). Since $Z_{\alpha\beta}$ is simply a Wiener process with drift m (recall that the diffusion coefficients σ are assumed constant), a standard argument based on Dynkin's formula yields $\mathbb{E}[\tilde{\tau}_a] = a/m$ and (4.10) follows. \square

A.4. Proof of Theorem 5.2. Parts 1 and 2 of Theorem 5.2 follow readily from Proposition 5.3 – simply note that the hypothesis of Proposition 5.3 is satisfied in both cases. As such, we only need to show that strict Nash equilibria are stochastically asymptotically stable under (SRL).

To that end, let $x^* = (\alpha_1^*, \dots, \alpha_N^*)$ be a strict equilibrium of \mathcal{G} and let $\mathcal{A}_k^* \equiv \mathcal{A}_k \setminus \{\alpha_k^*\}$. Moreover, for all $\alpha \in \mathcal{A}_k^*$, set

$$Z_{k\alpha} = \eta_k (Y_{k\alpha} - Y_{k\alpha_k^*}), \quad (\text{A.23})$$

so that $X(t) \rightarrow x^*$ if and only if $Z_{k\alpha}(t) \rightarrow -\infty$ for all $\alpha \in \mathcal{A}_k^*$, $k \in \mathcal{N}$ (cf. Proposition B.1).

Now, fix some tolerance $\varepsilon > 0$ and a neighborhood U_0 of x^* in \mathcal{X} . Since x^* is a strict equilibrium of \mathcal{G} , there exist $m_k > 0$ and a neighborhood $U \subseteq U_0$ of x^* such that

$$v_{k\alpha_k^*}(x) - v_{k\alpha}(x) \geq m_k \quad \text{for all } x \in U \text{ and for all } \alpha \in \mathcal{A}_k^*, k \in \mathcal{N}. \quad (\text{A.24})$$

With this in mind, let $M > 0$ be sufficiently large so that $X(t) \in U$ if $Z_{k\alpha}(t) \leq -M$ for all $\alpha \in \mathcal{A}_k^*$, $k \in \mathcal{N}$ (that such an M exists is again a consequence of Proposition B.1); furthermore, with a fair degree of hindsight, assume also that

$$M > m_k^{-1} \eta_k(0) \sigma_{k,\max}^2 \log(N/\varepsilon) \quad \text{for all } k = 1, \dots, N, \quad (\text{A.25})$$

where $\sigma_{k,\max}^2 = \max_{\alpha,x} \sigma_{k\alpha}^2(x)$. We will show that if $Z_{k\alpha}(0) \leq -2M$, then $X(t) \in U$ for all $t \geq 0$ and $Z_{k\alpha}(t) \rightarrow -\infty$ with probability at least $1 - \varepsilon$.

Indeed, assume that $Z_{k\alpha}(0) \leq -2M$ in (A.23) and define the escape time:

$$\tau_U = \inf\{t > 0 : X(t) \notin U\}. \quad (\text{A.26})$$

The stochastic dynamics (SRL) then give:

$$d(Y_{k\alpha} - Y_{k\alpha_k^*}) = [v_{k\alpha} - v_{k\alpha_k^*}] dt + \sigma_{k\alpha} dW_{k\alpha} - \sigma_{k\alpha_k^*} dW_{k\alpha_k^*}, \quad (\text{A.27})$$

so, for all $t \leq \tau_U$, we will have:

$$\begin{aligned} Z_{k\alpha}(t) &= Z_{k\alpha}(0) + \eta_k(t) \int_0^t [v_{k\alpha}(X(s)) - v_{k\alpha_k^*}(X(s))] ds + \eta_k(t)\xi_k(t) \\ &\leq -2M - \eta_k(t) [mt - \xi_k(t)], \end{aligned} \quad (\text{A.28})$$

where we have set

$$\xi_k(t) = \int_0^t \sigma_{k\alpha}(X(s)) dW_{k\alpha}(s) - \int_0^t \sigma_{k\alpha_k^*}(X(s)) dW_{k\alpha_k^*}(s). \quad (\text{A.29})$$

We will first show that $\mathbb{P}(\tau_U < \infty) \leq \varepsilon$. To that end, note that the time-change theorem for martingales (Øksendal, 2007, Cor. 8.5.4) provides a standard Wiener process $\widetilde{W}_k(t)$ such that $\xi_k(t) = \widetilde{W}_k(\rho_k(t))$ where $\rho_k = [\xi_k, \xi_k]$ is the quadratic variation of ξ_k . Then, from the fact that η is nonincreasing, we conclude that $Z_{k\alpha}(t) \leq -M$ whenever $m_k t - \widetilde{W}_k(\rho_k(t)) \geq -M/\eta_k(0)$. Accordingly, with $\rho_k(t) \leq 2\sigma_{k,\max}^2 t$ (cf. the proof of Proposition 4.2), it suffices to show that the hitting time

$$\tau_0 = \inf \left\{ t > 0 : \widetilde{W}_k(t) = \frac{m_k t}{2\sigma_{k,\max}^2} + \frac{M}{\eta_k(0)} \text{ for some } k \in \mathcal{N} \right\} \quad (\text{A.30})$$

is finite with probability not exceeding ε . However, if a trajectory of $\widetilde{W}_k(t)$ has $\widetilde{W}_k(t) \leq m_k t / (2\sigma_{k,\max}^2) + M/\eta_k(0)$ for all $t \geq 0$, we will also have

$$\widetilde{W}_k(\rho_k(t)) \leq \frac{m_k \rho_k(t)}{2\sigma_{k,\max}^2} + \frac{M}{\eta_k(0)} \leq m_k t + \frac{M}{\eta_k(0)}, \quad (\text{A.31})$$

so τ_U must be infinite for every trajectory of $\widetilde{W} = (\widetilde{W}_1, \dots, \widetilde{W}_N)$ with infinite τ_0 , i.e. $\mathbb{P}(\tau_U < +\infty) \leq \mathbb{P}(\tau_0 < +\infty)$. Thus, if we write E_k for the event that $\widetilde{W}_k(t) \geq m_k t / (2\sigma_{k,\max}^2) + M/\eta_k(0)$ for some finite $t \geq 0$, the hitting-time analysis of Karatzas and Shreve (1998, p. 197) for a Brownian motion with drift yields $\mathbb{P}(E_k) = e^{-\lambda_k M}$ with $\lambda_k = m_k / (\eta_k(0)\sigma_{k,\max}^2)$. Thus, by construction of M , we obtain:

$$\mathbb{P}(\tau_0 < +\infty) = \mathbb{P}\left(\bigcup_k E_k\right) \leq \sum_k \mathbb{P}(E_k) = \sum_k e^{-\lambda_k M} \leq \varepsilon. \quad (\text{A.32})$$

Thus, by conditioning on the event $\tau_U = +\infty$ and invoking Assumption 1 and Lemma B.4, Eq. (A.28) finally yields

$$Z_{k\alpha}(t) \leq -2M - \eta_k(t) \cdot [mt - \xi_k(t)] \sim -\eta_k(t) \cdot mt \rightarrow -\infty \quad (\text{a.s.}). \quad (\text{A.33})$$

We conclude that $Z_{k\alpha}(t) \rightarrow -\infty$ for all $\alpha \in \mathcal{A}_k^*$, $k \in \mathcal{N}$, so $\lim_{t \rightarrow \infty} X(t) = x^*$ (conditionally a.s.) and our proof is complete.

A.5. Proof of Proposition 6.2. Let $F_p(t) = \sum_{k=1,2} F_k(p_k, \eta_k(t)Y_k(t))$ where $p = (p_1, p_2)$ is an interior Nash equilibrium of \mathcal{G} . Using the Itô formula of Lemma

B.3, we get:

$$\begin{aligned} dF_p &= \sum_{k,\beta} (X_{k\beta} - p_{k\beta}) d(\eta_k Y_k) + \frac{1}{2} \sum_{k,\beta} \frac{\partial^2 h_k^*}{\partial y_{k\beta}^2} \eta_k^2 \sigma_{k\beta}^2 dt \\ &= \sum_{k=1,2} \dot{\eta}_k \langle Y_k, X_k - p \rangle dt \end{aligned} \quad (\text{A.34a})$$

$$+ \sum_{k=1,2} \eta_k \langle v_k, X_k - p_k \rangle dt \quad (\text{A.34b})$$

$$+ \sum_{k,\beta} \eta_k (X_{k\beta} - p_{k\beta}) \sigma_{k\beta} dW_{k\beta} \quad (\text{A.34c})$$

$$+ \frac{1}{2} \sum_{k,\beta} \frac{\partial^2 h_k^*}{\partial y_{k\beta}^2} \eta_k^2 \sigma_{k\beta}^2 dt, \quad (\text{A.34d})$$

where we substituted dY from (SRL) to obtain (A.34b) and (A.34c).

We now claim that $F_p(t)$ grows sublinearly in t ; indeed:

a) Lemma B.4 shows that $Y(t) = \mathcal{O}(t)$, so there exist $M_k > 0$, $k = 1, 2$, such that

$$\sum_{k=1,2} \int_0^t \dot{\eta}_k(s) \langle X_k(s) - p, Y_k(s) \rangle ds \leq \sum_{k=1,2} M_k \int_0^t |\dot{\eta}_k(s)| s ds = o(t), \quad (\text{A.35})$$

where the sublinearity estimate follows from Assumption 1 and the fact that $\eta_k(t) \rightarrow 0$:

$$\int_0^t \dot{\eta}_k(s) s ds = \eta_k(t) t - \int_0^t \eta_k(s) ds = o(t). \quad (\text{A.36})$$

b) The term (A.34b) is identically zero because \mathcal{G} is zero-sum and p is an interior equilibrium of \mathcal{G} :

$$\begin{aligned} &\langle v_1(X), X_1 - p_1 \rangle + \langle v_2(X), X_2 - p_2 \rangle \\ &= u_1(X_1, X_2) - u_1(p_1, X_2) + u_2(X_1, X_2) - u_2(X_1, p_2) = 0. \end{aligned} \quad (\text{A.37})$$

c) The term (A.34c) is sublinear (a.s.) on account of (the proof of) Lemma B.4.

d) Finally, for (A.34d), the same reasoning as in the proof of Theorem 3.1 yields:

$$\frac{1}{2} \sum_{k,\beta} \int_0^t \frac{\partial^2 h_k^*}{\partial y_{k\beta}^2} \eta_k^2(s) \sigma_{k\beta}^2(X(s)) ds \leq \frac{1}{2K} \sum_{k,\beta} \sigma_{k,\max}^2 \int_0^t \eta_k^2(s) ds, \quad (\text{A.38})$$

with $\sigma_{k,\max}^2 = \max_{\alpha \in \mathcal{A}_k} \max_{x \in \mathcal{X}} \sigma_{k\alpha}^2(x)$ defined as in the proof of Theorem 5.2. Since $\eta_k(t) \rightarrow 0$, this last integral is also sublinear in t , as claimed.

Assume now that $\limsup_{t \rightarrow \infty} F_p(t) \geq m$ for some $m > 0$ (otherwise we would have the stronger result $X(t) \rightarrow p$). Then, Lemma B.2 gives $|Y_{k\alpha}(t) - Y_{k\beta}(t)| = \mathcal{O}(F_p(t)) = o(t)$, so our claim follows from Theorem 6.1.

A.6. Proof of Theorem 6.3. As we mentioned, the proof follows the ideas of Hofbauer et al. (2009). Specifically, from (SRL), we have:

$$\begin{aligned} Y_k(t) &= Y_k(0) + \int_0^t v_k(X(s)) ds + \int_0^t \sigma_k(X(s)) dW_k(s) \\ &= tv_k(\bar{X}(t)) + Y_k(0) + \xi_k(t), \end{aligned} \quad (\text{A.39})$$

where we have set $\xi_{k\alpha}(t) = \int_0^t \sigma_{k\alpha}(X(s)) dW_{k\alpha}(s)$ and we have used the fact that \mathcal{G} is a 2-player game in order to carry the integral inside the argument of v_k . Consequently, by the definition (2.4) of the players' regularized best response maps, $X_k(t) = Q_k(\eta_k(t)Y_k(t))$ solves the (strictly) concave maximization problem:

$$\text{maximize } \langle v_k(\bar{X}(t)), x_k \rangle + \frac{1}{t} \langle Y_k(0) + \xi_k(t), x_k \rangle - \frac{1}{t\eta_k(t)} h_k(x_k), \quad (\text{A.40})$$

subject to $x_k \in \mathcal{X}_k$.

By Lemma B.4 and Assumption 1, the last two terms of (A.40) vanish as $t \rightarrow 0$. Hence, by the maximum theorem of Berge (1997, p. 116) applied to the parametric optimization problem (A.40) with parameter t , it follows that $X_k(t)$ lies within a vanishing distance of $\arg \max_{x_k \in \mathcal{X}_k} \langle v_k(\bar{X}(t)), x_k \rangle \equiv \text{br}_k(\bar{X}(t))$.

On the other hand, differentiating $\bar{X}(t)$ yields:

$$\frac{d}{dt} \bar{X}(t) = t^{-1} X(t) - t^{-2} \int_0^t X(s) ds = t^{-1} [X(t) - \bar{X}(t)], \quad (\text{A.41})$$

and, after changing time to $\tau = \log t$, the expression above becomes $\frac{d}{dt} \bar{X}(t) = X - \bar{X}$. Combining all of the above, we conclude that $\bar{X}(t)$ tracks a perturbed version of the best reply dynamics (BRD) in the sense of Benaïm et al. (2005, Def. III), and our assertion follows from Theorem 3.6 in the same paper.

A.7. Correlated uncertainty. We provide below a short sketch of the proof of Proposition 7.1, highlighting the points of departure from the proofs of Theorems 3.1, 4.1, 5.2, and Proposition 6.2.

First, for the no-regret analysis of (SRL_c), note that (A.10) still holds essentially as stated, except for the term (A.10d) which should be replaced by the quantity

$$\frac{1}{2} \sum_{\beta} \int_0^t \eta(s) \text{tr}[\text{Hess}(h^*(Y(s))) \Sigma(s)] ds. \quad (\text{A.42})$$

Since the matrices $\text{Hess}(h^*)$ and Σ are both bounded in norm (by the strong convexity of h and by assumption, respectively), the rest of the proof of Theorem 3.1 goes through exactly as in the uncorrelated case, except that σ_{\max} must now be replaced by $\sup_t \|\Sigma(t)\|$.

For the elimination of dominated strategies and the stochastic asymptotic stability of strict equilibria, we first need to replace all terms of the form $\sigma_{k\alpha} dW_{k\alpha}$ in (4.5) and (A.29) by $dZ_{k\alpha}$. Then, it suffices to note that the rest of the proof of Theorems 4.1 and 5.2 only relies on the fact that the quadratic variation of ξ satisfies an inequality of the form $d[\xi, \xi] \leq M dt$ for some $M > 0$ (a consequence of the fact that Σ is assumed bounded).

Finally, for our result on time-averages, the only nontrivial change in the proof of Proposition 6.2 is that the term $\sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2 dt$ that appears in (A.34) must be replaced by the correlated variant $\text{tr}[\text{Hess}(h^*(Y(s))) \Sigma(s)] ds$. This term can be bounded as in the no-regret case above, so (A.38) still holds (but with a different constant than σ_{\max}^2/K); the rest of the proof then goes through as in the uncorrelated setting.

APPENDIX B. AUXILIARY RESULTS

B.1. Properties of the Fenchel coupling. In this appendix, we collect some basic properties of the Fenchel coupling and regularized best response maps. To state them, let Δ denote the d -dimensional simplex of \mathbb{R}^n and let $h: \Delta \rightarrow \mathbb{R}$ be a penalty function on Δ (i.e. h is continuous and strongly convex on Δ and smooth on the relative interior of any face of Δ). Also, recall that the *regularized best response map* $Q: \mathbb{R}^n \rightarrow \Delta$ induced by h is given by

$$Q(y) = \arg \max_{x \in \Delta} \{\langle y, x \rangle - h(x)\}, \quad y \in \mathbb{R}^n, \quad (\text{B.1})$$

while the *Fenchel coupling* between $p \in \Delta$ and $y \in \mathbb{R}^n$ is defined as:

$$F(p, y) = h(p) + h^*(y) - \langle y, p \rangle, \quad (\text{B.2})$$

where

$$h^*(y) = \max_{x \in \Delta} \{\langle y, x \rangle - h(x)\} \quad (\text{B.3})$$

denotes the convex conjugate of h . We then have:

Proposition B.1. *Let $h: \Delta \rightarrow \mathbb{R}$ be a penalty function on Δ with strong convexity constant K . Then:*

- (1) Q is $1/K$ -Lipschitz and $Q(y) = dh^*(y)$.
- (2) If $y_\alpha - y_\beta \rightarrow -\infty$ for some $\beta \neq \alpha$, then $Q_\alpha(y) \rightarrow 0$.
- (3) $F(p, y) \geq \frac{1}{2} \|Q(y) - p\|^2$ for all $p \in \Delta$ and for all y ; in particular, $F(p, y) \geq 0$ with equality if and only if $p = Q(y)$.
- (4) If $F(p, y_j) \rightarrow +\infty$ for some sequence $y_j \in \mathbb{R}^n$, the sequence $x_j = Q(y_j)$ converges to the union of faces of Δ that do not contain p ; more precisely: $\liminf_{j \rightarrow \infty} \{x_{j,\alpha} : \alpha \in \text{supp}(p)\} = 0$.

Proof. The first part of Proposition B.1 is well known (see e.g. Nesterov, 2009; Shalev-Shwartz, 2011). The rest of our claims are due to Mertikopoulos and Sandholm (2016); we only provide a proof below for completeness and ease of reference. Part 2. Let y_j be a sequence in \mathbb{R}^n such that $y_{j,\alpha} - y_{j,\beta} \rightarrow -\infty$, set $x_j = Q(y_j)$ and assume (by descending to a subsequence if necessary) that $x_{j,\alpha} > \varepsilon > 0$ for all n . By definition, we have $\langle y_j, x_j \rangle - h(x_j) \geq \langle y_j, x' \rangle - h(x')$ for all $x' \in \Delta$, so if we set $x'_j = x_j + \varepsilon(e_\beta - e_\alpha)$, we readily obtain

$$\varepsilon(y_{j,\alpha} - y_{j,\beta}) \geq h(x_j) - h(x'_j) \geq -(h_{\max} - h_{\min}). \quad (\text{B.4})$$

This contradicts our original assumption that $y_{j,\alpha} - y_{j,\beta} \rightarrow -\infty$; since Δ is compact, we get $x_\alpha^* = 0$ for any limit point x^* of x_j , i.e. $Q_\alpha(y_j) \rightarrow 0$.

Part 3. Let $x = Q(y)$ so that $h^*(y) = \langle y, x \rangle - h(x)$ and $F(p, y) = h(p) - h(x) - \langle y, p - x \rangle$. Since $x = dh^*(y)$ by Part 1, we will also have

$$\langle y, x' - x \rangle \leq h(x') - h(x) \quad \text{for all } x' \in \Delta, \quad (\text{B.5})$$

by standard convex analysis arguments (Rockafellar, 1970, Chap. 26). On the other hand, letting $z = p - x$, the definition (2.5) of strong convexity yields:

$$h(x + tz) \leq th(p) + (1 - t)h(x) - \frac{1}{2}Kt(1 - t)\|z\|^2 \quad \text{for all } t \in (0, 1). \quad (\text{B.6})$$

Hence, combining (B.6) with (B.5) for $x' = x + tz$, we get

$$h(x) + \langle y, tz \rangle \leq th(p) + (1 - t)h(x) - \frac{1}{2}Kt(1 - t)\|z\|^2, \quad (\text{B.7})$$

and, after rearranging and dividing by t , we obtain

$$F(p, y) = h(p) - h(x) - \langle y, z \rangle \geq \frac{1}{2}K(1-t)\|z\|^2. \quad (\text{B.8})$$

Our claim then follows by letting $t \rightarrow 0^+$ in (B.8).

Part 4. Let Λ_p denote the union of faces of Δ that do not contain p and set $\Delta_p \equiv \Delta \setminus \Lambda_p$. Then, if $x = Q(y) \in \Delta_p$ for some $y \in \mathbb{R}^n$, the function $h(x + t(p-x))$ will be finite and smooth for all t in a neighborhood of 0 (simply note that $x + t(p-x)$ lies in the relative interior of the same face of Δ for small t). This implies that h admits a two-sided derivative at x along $p-x$, so we will also have $\langle y, p-x \rangle = h'(x; p-x) < +\infty$ by Theorem 23.5 in Rockafellar (1970); in particular, this shows that $F(p, y) = h(p) - h(x) - h'(x; p-x) < +\infty$ whenever $x = Q(y) \in \Delta_p$.

Assume now ad absurdum that $x_j = Q(y_j)$ has a limit point in Δ_p , so there exists a compact neighborhood K of p in Δ_p which is visited infinitely often by x_j . By continuity and compactness, this implies that there exists some $M > 0$ such that $F(p, y_j) = h(p) - h(x_j) - h'(x_j; p-x_j) < M$ infinitely often, a contradiction to our original assumption that $F(p, y_j) \rightarrow +\infty$. We conclude that every limit point of x_j lies in Δ_p , as claimed. \square

For the proof of Proposition 6.2, we also require the following lemma:

Lemma B.2. *Let $p \in \Delta^\circ$ and assume that the sequence y_j has $F(p, y_j) \geq m$ for some $m > 0$. Then, for all α, β , there exists some $K > 0$ such that $|y_{j,\alpha} - y_{j,\beta}| \leq KF(p, y_j)$.*

Proof. By descending to a subsequence if necessary, assume ad absurdum that $y_{j,\alpha} - y_{j,\beta} \geq K_j F(p, y_j)$ for some increasing sequence $K_j \rightarrow +\infty$. Then, by relabeling indices (and passing to a finer subsequence if needed), we may assume that a) $y_{j,\alpha} \geq y_{j,\gamma} \geq y_{j,\beta}$ for all $\gamma = 1, \dots, d$; and b) the index set $\{1, \dots, d\}$ can be decomposed into two nonempty sets, \mathcal{A}_0 and \mathcal{A}_∞ , such that $y_{j,\alpha} - y_{j,\gamma} \leq K' F(p, y_j)$ for some $K' > 0$ and for all $\gamma \in \mathcal{A}_0$, while $y_{j,\alpha} - y_{j,\gamma} > K'_j F(p, y_j)$ for some $K'_j \rightarrow +\infty$ and for all $\gamma \in \mathcal{A}_\infty$. Thus, letting $x_j = Q(y_j)$, we get:

$$\begin{aligned} \langle y_j, p - x_j \rangle &= \sum_{\gamma=1}^n y_{j,\gamma} (p_\gamma - x_{j,\gamma}) = \sum_{\gamma=1}^n (y_{j,\gamma} - y_{j,\alpha}) (p_\gamma - x_{j,\gamma}) \\ &= \sum_{\gamma \in \mathcal{A}_0} (y_{j,\gamma} - y_{j,\alpha}) (p_\gamma - x_{j,\gamma}) + \sum_{\gamma \in \mathcal{A}_\infty} (y_{j,\gamma} - y_{j,\alpha}) (p_\gamma - x_{j,\gamma}). \end{aligned} \quad (\text{B.9})$$

By construction, the first sum is bounded in absolute value by $2dK'F(p, y_j)$. As for the second sum, since $F(p, y_j) \geq M > 0$, we will also have $y_{j,\alpha} - y_{j,\gamma} \rightarrow +\infty$, so $x_{j,\gamma} \rightarrow 0$ by Proposition B.1; with $p \in \Delta^\circ$, we then get $\liminf_j (p_\gamma - x_{j,\gamma}) \geq \delta$ for some $\delta > 0$, so every \mathcal{A}_∞ summand in (B.9) is eventually bounded from above by $-\delta K'_j F(p, y_j)$. This shows that $\langle y_j, p - x_j \rangle \leq -K''_j F(p, y_j)$ for some positive sequence $K''_j \rightarrow +\infty$; however, $\langle y_j, p - x_j \rangle = h(p) - h(x_j) - F(p, y_j)$, a contradiction. \square

B.2. Stochastic analysis. From a stochastic analysis standpoint, the importance of the Fenchel coupling is captured by the following calculation:

Lemma B.3. *Let $dY_\alpha = \mu_\alpha dt + \sum_\beta \sigma_{\alpha\beta} dW_\beta$ be an Itô process with values in \mathbb{R}^n and set $X(t) = Q(Y(t))$. Then, for all $p \in \Delta$, we have:*

$$dF(p, Y) = \sum_\beta (X_\beta - p_\beta) dY_\beta + \frac{1}{2} \sum_\beta \frac{\partial^2 h^*}{\partial y_\beta^2} \sigma_\beta^2 dt. \quad (\text{B.10})$$

Proof. Let $H(t) = F(p, Y(t))$, $t \geq 0$. Then, Itô's formula yields:

$$\begin{aligned} dH &= \sum_{\beta} \frac{\partial F}{\partial y_{\beta}} dY_{\beta} + \frac{1}{2} \sum_{\beta, \gamma} \frac{\partial^2 F}{\partial y_{\beta} \partial y_{\gamma}} dY_{\beta} \cdot dY_{\gamma} \\ &= \sum_{\beta} \left(\frac{\partial h^*}{\partial y_{\beta}} - x_{\beta} \right) dY_{\beta} + \frac{1}{2} \sum_{\beta, \gamma} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta} \sigma_{\gamma} \delta_{\beta\gamma} dt \\ &= \sum_{\beta} (X_{\beta} - p_{\beta}) dY_{\beta} + \frac{1}{2} \sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2 dt, \end{aligned} \quad (\text{B.11})$$

where we used the definition of F and Proposition B.1 in the second and third lines. \square

Finally, Lemma B.4 is a useful order comparison result that we employ throughout:

Lemma B.4. *Let $W(t) = (W_1(t), \dots, W_n(t))$, $t \geq 0$, be a Wiener processes in \mathbb{R}^n and let $Z(t)$ be a bounded, continuous process in \mathbb{R}^n . Then, for every continuous function $f: [0, \infty) \rightarrow (0, \infty)$, we have*

$$f(t) + \int_0^t Z(s) \cdot dW(s) \sim f(t) \quad \text{as } t \rightarrow \infty \text{ (a.s.)}, \quad (\text{B.12})$$

whenever $\lim_{t \rightarrow \infty} (t \log \log t)^{-1/2} f(t) = +\infty$.

Proof. Let $\xi(t) = \int_0^t Z(s) \cdot dW(s) = \sum_{\alpha=1}^n \int_0^t Z_{\alpha}(s) dW_{\alpha}(s)$. Then, the quadratic variation $\rho = [\xi, \xi]$ of ξ satisfies:

$$d[\xi, \xi] = d\xi \cdot d\xi = \sum_{\alpha=1}^n Z_{\alpha} Z_{\beta} \delta_{\alpha\beta} dt \leq M dt, \quad (\text{B.13})$$

where $M = \sup_{t \geq 0} \|Z(t)\|^2 < +\infty$ (recall that $Z(t)$ is bounded by assumption). On the other hand, by the time-change theorem for martingales (Øksendal, 2007, Cor. 8.5.4), there exists a Wiener process $\widetilde{W}(t)$ such that $\xi(t) = \widetilde{W}(\rho(t))$, and hence:

$$\frac{f(t) + \xi(t)}{f(t)} = 1 + \frac{\widetilde{W}(\rho(t))}{f(t)}. \quad (\text{B.14})$$

Obviously, if $\lim_{t \rightarrow \infty} \rho(t) < \infty$, we will have $\limsup_{t \rightarrow \infty} |\widetilde{W}(\rho(t))| < \infty$ (a.s.) so there is nothing to show. Otherwise, if $\lim_{t \rightarrow \infty} \rho(t) = \infty$, the quadratic variation bound (B.13) and the law of the iterated logarithm yield:

$$\frac{|\widetilde{W}(\rho(t))|}{f(t)} \leq \frac{|\widetilde{W}(\rho(t))|}{\sqrt{2\rho(t) \log \log \rho(t)}} \times \frac{\sqrt{2Mt \log \log Mt}}{f(t)} \rightarrow 0 \quad \text{as } t \rightarrow \infty, \quad (\text{B.15})$$

and our claim follows. \square

REFERENCES

- Akin, E., 1980. Domination or equilibrium. *Mathematical Biosciences* 50 (3-4), 239–250.
- Alsmeyer, G., Jaeger, M., August 2005. A useful extension of itô's formula with applications to optimal stopping. *Acta Mathematica Sinica* 21 (4), 779–786.
- Benaïm, M., 1999. Dynamics of stochastic approximation algorithms. In: Azéma, J., Émery, M., Ledoux, M., Yor, M. (Eds.), *Séminaire de Probabilités XXXIII*. Vol. 1709 of *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, pp. 1–68.
- Benaïm, M., Faure, M., 2013. Consistency of vanishingly smooth fictitious play. *Mathematics of Operations Research* 38 (3), 437–450.

- Benaïm, M., Hofbauer, J., Sorin, S., 2005. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization* 44 (1), 328–348.
- Berge, C., 1997. *Topological Spaces*. Dover, New York.
- Börgers, T., Sarin, R., 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77 (1), 1–14.
- Cabrales, A., May 2000. Stochastic replicator dynamics. *International Economic Review* 41 (2), 451–81.
- Cominetti, R., Melo, E., Sorin, S., 2010. A payoff-based learning procedure and its application to traffic games. *Games and Economic Behavior* 70 (1), 71–83.
- Coucheney, P., Gaujal, B., Mertikopoulos, P., August 2015. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research* 40 (3), 611–633.
- Erev, I., Roth, A. E., 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88, 848–881.
- Friedman, D., 1991. Evolutionary games in economics. *Econometrica* 59 (3), 637–666.
- Fudenberg, D., Harris, C., August 1992. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory* 57 (2), 420–441.
- Fudenberg, D., Levine, D. K., 1998. *The Theory of Learning in Games*. Vol. 2 of *Economic learning and social evolution*. MIT Press, Cambridge, MA.
- Gilboa, I., Matsui, A., May 1991. Social stability and equilibrium. *Econometrica* 59 (3), 859–867.
- Hannan, J., 1957. Approximation to Bayes risk in repeated play. In: Dresher, M., Tucker, A. W., Wolfe, P. (Eds.), *Contributions to the Theory of Games, Volume III*. Vol. 39 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, pp. 97–139.
- Hofbauer, J., Imhof, L. A., 2009. Time averages, recurrence and transience in the stochastic replicator dynamics. *The Annals of Applied Probability* 19 (4), 1347–1368.
- Hofbauer, J., Sandholm, W. H., November 2002. On the global convergence of stochastic fictitious play. *Econometrica* 70 (6), 2265–2294.
- Hofbauer, J., Sigmund, K., 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.
- Hofbauer, J., Sigmund, K., July 2003. Evolutionary game dynamics. *Bulletin of the American Mathematical Society* 40 (4), 479–519.
- Hofbauer, J., Sorin, S., Viossat, Y., May 2009. Time average replicator and best reply dynamics. *Mathematics of Operations Research* 34 (2), 263–269.
- Hopkins, E., November 2002. Two competing models of how people learn in games. *Econometrica* 70 (6), 2141–2166.
- Imhof, L. A., 2005. The long-run behavior of the stochastic replicator dynamics. *The Annals of Applied Probability* 15 (1B), 1019–1045.
- Kang, W. N., Kelly, F. P., Lee, N. H., Williams, R. J., 2009. State space collapse and diffusion approximation for a network operating under a fair bandwidth-sharing policy. *Annals of Applied Probability* 19, 1719–1780.
- Karatzas, I., Shreve, S. E., 1998. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, Berlin.
- Kelly, F. P., Maulloo, A. K., Tan, D. K. H., 1998. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49, 237–252.
- Khasminskii, R. Z., 2012. *Stochastic Stability of Differential Equations*, 2nd Edition. No. 66 in *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin.
- Khasminskii, R. Z., Potsepun, N., 2006. On the replicator dynamics behavior under Stratonovich type random perturbations. *Stochastic Dynamics* 6, 197–211.
- Kwon, J., Mertikopoulos, P., 2014. A continuous-time approach to online optimization, <http://arxiv.org/abs/1401.6956>.

- Lahkar, R., Sandholm, W. H., 2008. The projection dynamic and the geometry of population games. *Games and Economic Behavior* 64, 565–590.
- Lamberton, D., Pagès, G., Tarrès, P., 2004. When can the two-armed bandit algorithm be trusted? *The Annals of Applied Probability* 14 (3), 1424–1454.
- Leslie, D. S., Collins, E. J., 2005. Individual Q -learning in normal form games. *SIAM Journal on Control and Optimization* 44 (2), 495–514.
- Littlestone, N., Warmuth, M. K., 1994. The weighted majority algorithm. *Information and Computation* 108 (2), 212–261.
- Mertikopoulos, P., Moustakas, A. L., 2009. Learning in the presence of noise. In: *GameNets '09: Proceedings of the 1st International Conference on Game Theory for Networks*.
- Mertikopoulos, P., Moustakas, A. L., 2010. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability* 20 (4), 1359–1388.
- Mertikopoulos, P., Sandholm, W. H., 2016. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*.
- Mertikopoulos, P., Viossat, Y., March 2016. Imitation dynamics with payoff shocks. *International Journal of Game Theory* 45 (1-2), 291–320.
- Nachbar, J. H., 1990. Evolutionary selection dynamics in games. *International Journal of Game Theory* 19, 59–89.
- Nagurney, A., Zhang, D., 1997. Projected dynamical systems in the formulation, stability analysis, and computation of fixed demand traffic network equilibria. *Transportation Science* 31, 147–158.
- Nesterov, Y., 2009. Primal-dual subgradient methods for convex problems. *Mathematical Programming* 120 (1), 221–259.
- Øksendal, B., 2007. *Stochastic Differential Equations*, 6th Edition. Springer-Verlag, Berlin.
- Oyarzun, C., Ruf, J., 2014. Convergence in models with bounded expected relative hazard rates. *Journal of Economic Theory* 154, 229–244.
- Rockafellar, R. T., 1970. *Convex Analysis*. Princeton University Press, Princeton, NJ.
- Rustichini, A., 1999. Optimal properties of stimulus-response learning models. *Games and Economic Behavior* 29 (1-2), 244–273.
- Samuelson, L., Zhang, J., 1992. Evolutionary stability in asymmetric games. *Journal of Economic Theory* 57, 363–391.
- Sandholm, W. H., 2010. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA.
- Shalev-Shwartz, S., 2011. Online learning and online convex optimization. *Foundations and Trends in Machine Learning* 4 (2), 107–194.
- Sorin, S., 2009. Exponential weight algorithm in continuous time. *Mathematical Programming* 116 (1), 513–528.
- Sutton, R. S., Barto, A. G., 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Taylor, P. D., Jonker, L. B., 1978. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* 40 (1-2), 145–156.
- van Kampen, N. G., 1981. Itô versus Stratonovich. *Journal of Statistical Physics* 24 (1), 175–187.
- Viossat, Y., April 2015. Evolutionary dynamics and dominated strategies. *Economic Theory Bulletin* 3 (1), 91–113.
- Vlasic, A., 2012. Long-run analysis of the stochastic replicator dynamics in the presence of random jumps. <http://arxiv.org/abs/1206.0344>.
- Vovk, V. G., 1990. Aggregating strategies. In: *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*. pp. 371–383.

(M. Bravo) UNIVERSIDAD DE SANTIAGO DE CHILE, DEPARTAMENTO DE MATEMÁTICA Y CIENCIA DE LA COMPUTACIÓN, AV. LIBERTADOR BERNARDO O'HIGGINS 3363, SANTIAGO, CHILE
E-mail address: mario.bravo.g@usach.cl

(P. Mertikopoulos) CNRS (FRENCH NATIONAL CENTER FOR SCIENTIFIC RESEARCH), LIG, F-38000 GRENOBLE, FRANCE, AND UNIV. GRENOBLE ALPES, LIG, F-38000 GRENOBLE, FRANCE
E-mail address: panayotis.mertikopoulos@imag.fr
URL: <http://mesca.lig.imag.fr/membres/panayotis.mertikopoulos>