

IMITATION DYNAMICS WITH PAYOFF SHOCKS

PANAYOTIS MERTIKOPOULOS AND YANNICK VIOSSAT

ABSTRACT. We investigate the impact of payoff shocks on the evolution of large populations of myopic players that employ simple strategy revision protocols such as the “imitation of success”. In the noiseless case, this process is governed by the standard (deterministic) replicator dynamics; in the presence of noise however, the induced stochastic dynamics are different from previous versions of the stochastic replicator dynamics (such as the aggregate-shocks model of Fudenberg and Harris, 1992). In this context, we show that strict equilibria are always stochastically asymptotically stable, irrespective of the magnitude of the shocks; on the other hand, in the high-noise regime, non-equilibrium states may also become stochastically asymptotically stable and dominated strategies may survive in perpetuity (they become extinct if the noise is low). Such behavior is eliminated if players are less myopic and revise their strategies based on their cumulative payoffs. In this case, we obtain a second order stochastic dynamical system whose attracting states coincide with the game’s strict equilibria and where dominated strategies become extinct (a.s.), no matter the noise level.

CONTENTS

1. Introduction	1
2. The model	3
3. Long-term rationality analysis	9
4. The effect of aggregating payoffs	15
5. Discussion	19
Appendix A. Auxiliary results from stochastic analysis	22
References	23

1. INTRODUCTION

Evolutionary game dynamics study the evolution of behavior in populations of boundedly rational agents that interact strategically. The most widely studied dynamical model in this context is the replicator dynamics: introduced in biology as a model of natural selection (Taylor and Jonker, 1978), the replicator dynamics also arise from models of imitation of successful individuals (Björnerstedt and Weibull, 1996; Schlag, 1998; Weibull, 1995) and from models of learning in games (Hofbauer et al., 2009; Mertikopoulos and Moustakas, 2010; Rustichini, 1999). Mathematically, they stipulate that the growth rate of the frequency of a strategy is proportional to the difference between the payoff of individuals playing this strategy and

Supported in part by the French National Research Agency under grant no. GAGA-13-JS01-0004-01 and the French National Center for Scientific Research (CNRS) under grant no. PEPS-GATHERING-2014.

the mean payoff in the population. These payoffs are usually assumed deterministic: this is typically motivated by a large population assumption and the premise that, owing to the law of large numbers, the resulting mean field provides a good approximation of a more realistic but less tractable stochastic model. This approach makes sense when the stochasticity affecting payoffs is independent across individuals playing the same strategies, but it fails when the payoff shocks are aggregate, that is, when they affect all individuals playing a given strategy in a similar way.

Such aggregate shocks are not uncommon. Bergstrom (2014) recounts the story of squirrels stocking nuts for the winter months: squirrels may stock a few or a lot of nuts, the latter leading to a higher probability of surviving a long winter but a higher exposure to predation. The unpredictable mildness or harshness of the ensuing winter will then favor one of these strategies in an aggregate way (see also Robson and Samuelson, 2011, Sec. 3.1.1, and references therein). In traffic engineering, one might think of a choice of itinerary to go to work: fluctuations of traffic on some roads affect all those who chose them in a similar way. Likewise, in data networks, a major challenge occurs when trying to minimize network latencies in the presence of stochastic disturbances: in this setting, the travel time of a packet in the network does not depend only on the load of each link it traverses, but also on unpredictable factors such as random packet drops and retransmissions, fluctuations in link quality, excessive backlog queues, etc. (Bertsekas and Gallager, 1992).

Incorporating such aggregate payoff shocks in the biological derivation of the replicator dynamics leads to the stochastic replicator dynamics of Fudenberg and Harris (1992), later studied by (among others) Cabrales (2000), Imhof (2005) and Hofbauer and Imhof (2009). To study the long-run behavior of these dynamics, Imhof (2005) introduced a modified game where the expected payoff of a strategy is penalized by a term which increases with the variance of the noise affecting this strategy's payoff (see also Hofbauer and Imhof, 2009). Among other results, it was then shown that *a*) strategies that are iteratively (strictly) dominated in this modified game become extinct almost surely; and *b*) strict equilibria of the modified game are stochastically asymptotically stable.

In this biological model, noise is detrimental to the long-term survival of strategies: a strategy which is strictly dominant on average (i.e. in the original, unmodified game) but which is affected by shocks of substantially higher intensity becomes extinct almost surely. By contrast, in the learning derivation of the replicator dynamics, noise leads to a stochastic exponential learning model where only iteratively undominated strategies survive, irrespective of the intensity of the noise (Mertikopoulos and Moustakas, 2010); as a result the frequency of a strictly dominant strategy converges to 1 almost surely. Moreover, strict Nash equilibria of the original game remain stochastically asymptotically stable (again, independently of the level of the noise), so the impact of the noise in the stochastic replicator dynamics of exponential learning is minimal when compared to the stochastic replicator dynamics with aggregate shocks.

In this paper, we study the effect of payoff shocks when the replicator equation is seen as a model of imitation of successful agents. As in the case of Imhof (2005) and Hofbauer and Imhof (2009), it is convenient to introduce a noise-adjusted game which is reduced to the original game in the noiseless, deterministic regime. We show that: *a*) strategies that are iteratively strictly dominated in the modified

game become extinct almost surely; and *b*) strict equilibria of the modified game are stochastically asymptotically stable. However, despite the formal similarity, our results are qualitatively different from those of Imhof (2005) and Hofbauer and Imhof (2009): in the modified game induced by imitation of success in the presence of noise, noise is not detrimental per se. In fact, in the absence of differences in expected payoffs, a strategy survives with a probability that does not depend on the random variance of its payoffs: a strategy's survival probability is simply its initial frequency. Similarly, even if a strategy which is strictly dominant in expectation is subject to arbitrarily high noise, it will always survive with positive probability; by contrast, such strategies become extinct (a.s.) in the aggregate shocks model of Fudenberg and Harris (1992).

That said, the dynamics' long-term properties change dramatically if players are less "myopic" and, instead of imitating strategies based on their instantaneous payoffs, they base their decisions on the cumulative payoffs of their strategies over time. In this case, we obtain a second-order stochastic replicator equation which can be seen as a noisy version of the higher order game dynamics of Laraki and Mertikopoulos (2013). Thanks to this payoff aggregation mechanism, the noise averages out in the long run and we recover results that are similar to those of Mertikopoulos and Moustakas (2010): strategies that are dominated in the original game become extinct (a.s.) and strict Nash equilibria attract nearby initial conditions with arbitrarily high probability.

1.1. Paper Outline. The remainder of our paper is structured as follows: in Section 2, we present our model and we derive the stochastic replicator dynamics induced by imitation of success in the presence of noise. Our long-term rationality analysis begins in Section 3 where we introduce the noise-adjusted game discussed above and we state our elimination and stability results in terms of this modified game. In Section 4, we consider the case where players imitate strategies based on their cumulative payoffs and we show that the adjustment due to noise is no longer relevant. Finally, in Section 5, we discuss some variants of our core model related to different noise processes.

1.2. Notational conventions. The real space spanned by a finite set $\mathcal{S} = \{s_\alpha\}_{\alpha=1}^{d+1}$ will be denoted by $\mathbb{R}^{\mathcal{S}}$ and we will write $\{e_s\}_{s \in \mathcal{S}}$ for its canonical basis; in a slight abuse of notation, we will also use α to refer interchangeably to either s_α or e_α and we will write $\delta_{\alpha\beta}$ for the Kronecker delta symbols on \mathcal{S} . The set $\Delta(\mathcal{S})$ of probability measures on \mathcal{S} will be identified with the d -dimensional simplex $\Delta = \{x \in \mathbb{R}^{\mathcal{S}} : \sum_{\alpha} x_\alpha = 1 \text{ and } x_\alpha \geq 0\}$ of $\mathbb{R}^{\mathcal{S}}$ and the relative interior of Δ will be denoted by Δ° ; also, the support of $p \in \Delta(\mathcal{S})$ will be written $\text{supp}(p) = \{\alpha \in \mathcal{S} : p_\alpha > 0\}$. For simplicity, if $\{\mathcal{S}_k\}_{k \in \mathcal{N}}$ is a finite family of finite sets, we use the shorthand $(\alpha_k; \alpha_{-k})$ for the tuple $(\dots, \alpha_{k-1}, \alpha_k, \alpha_{k+1}, \dots)$ and we write \sum_{α}^k instead of $\sum_{\alpha \in \mathcal{S}_k}$. Unless mentioned otherwise, deterministic processes will be represented by lowercase letters, while their stochastic counterparts will be denoted by the corresponding uppercase letter. Finally, we will suppress the dependence of the law of a process $X(t)$ on its initial condition $X(0) = x$, and we will write \mathbb{P} instead of \mathbb{P}_x .

2. THE MODEL

In this section, we recall a few preliminaries from the theory of population games and evolutionary dynamics, and we introduce the stochastic game dynamics under study.

2.1. Population games. Our main focus will be games played by populations of nonatomic players. Formally, such games consist of a finite set of player *populations* $\mathcal{N} = \{1, \dots, N\}$ (assumed for simplicity to have unit mass), each with a finite set of *pure strategies* (or *types*) $\mathcal{A}_k = \{\alpha_{k,1}, \alpha_{k,2}, \dots\}$, $k \in \mathcal{N}$. During play, each player chooses a strategy and the *state* of each population is given by the distribution $x_k = (x_{k\alpha})_{\alpha \in \mathcal{A}_k}$ of players employing each strategy $\alpha \in \mathcal{A}_k$. Accordingly, the *state space* of the k -th population is the simplex $\mathcal{X}_k \equiv \Delta(\mathcal{A}_k)$ and the state space of the game is the product $\mathcal{X} \equiv \prod_k \mathcal{X}_k$.

The payoff to a player of population $k \in \mathcal{N}$ playing $\alpha \in \mathcal{A}_k$ is determined by the corresponding *payoff function* $v_{k\alpha}: \mathcal{X} \rightarrow \mathbb{R}$ (assumed Lipschitz).¹ Thus, given a population state $x \in \mathcal{X}$, the average payoff to population k will be

$$\sum_{\alpha}^k x_{k\alpha} v_{k\alpha}(x) = \langle v_k(x) | x_k \rangle, \quad (2.1)$$

where $v_k(x) \equiv (v_{k\alpha}(x))_{\alpha \in \mathcal{A}_k}$ denotes the *payoff vector* of the k -th population in the state $x \in \mathcal{X}$. Putting all this together, a *population game* is then defined as a tuple $\mathfrak{G} \equiv \mathfrak{G}(\mathcal{N}, \mathcal{A}, v)$ of nonatomic player populations $k \in \mathcal{N}$, their pure strategies $\alpha \in \mathcal{A}_k$ and the associated payoff functions $v_{k\alpha}: \mathcal{X} \rightarrow \mathbb{R}$.

In this context, we say that a pure strategy $\alpha \in \mathcal{A}_k$ is *dominated* by $\beta \in \mathcal{A}_k$ if

$$v_{k\alpha}(x) < v_{k\beta}(x) \quad \text{for all } x \in \mathcal{X}, \quad (2.2)$$

i.e. the payoff of an α -strategist is always inferior to that of a β -strategist. More generally (and in a slight abuse of terminology), we will say that $p_k \in \mathcal{X}_k$ is dominated by $p'_k \in \mathcal{X}_k$ if

$$\langle v_k(x) | p_k \rangle < \langle v_k(x) | p'_k \rangle \quad \text{for all } x \in \mathcal{X}, \quad (2.3)$$

i.e. when the average payoff of a small influx of mutants in population k is always greater when they are distributed according to p'_k rather than p_k (irrespective of the incumbent population state $x \in \mathcal{X}$).

Finally, we will say that the population state $x^* \in \mathcal{X}$ is at *Nash equilibrium* if

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

In particular, if x^* is pure (in the sense that $\text{supp}(x^*)$ is a singleton) and (NE) holds as a strict inequality for all $\beta \notin \text{supp}(x_k^*)$, x^* will be called a *strict equilibrium*.

Remark 2.1. Throughout this paper, we will be suppressing the population index $k \in \mathcal{N}$ for simplicity, essentially focusing in the single-population case. This is done only for notational clarity: all our results apply as stated to the multi-population model described in detail above.

¹Note that we are considering general payoff functions and not only multilinear (resp. linear) payoffs arising from asymmetric (resp. symmetric) random matching in finite N -person (resp. 2-person) games. This distinction is important as it allows our model to cover e.g. general traffic games as in Sandholm (2010).

2.2. Revision protocols. A fundamental evolutionary model in the context of population games is provided by the notion of a *revision protocol*. Following Sandholm (2010, Chapter 3), it is assumed that each nonatomic player receives an opportunity to switch strategies at every ring of an independent Poisson alarm clock, and this decision is based on the payoffs associated to each strategy and the current population state. The players' revision protocol is thus defined in terms of the *conditional switch rates* $\rho_{\alpha\beta} \equiv \rho_{\alpha\beta}(v, x)$ that determine the relative mass $dx_{\alpha\beta}$ of players switching from α to β over an infinitesimal time interval dt :²

$$dx_{\alpha\beta} = x_{\alpha}\rho_{\alpha\beta} dt. \quad (2.4)$$

The population shares x_{α} are then governed by the *revision protocol dynamics*:

$$\dot{x}_{\alpha} = \sum_{\beta} x_{\beta}\rho_{\beta\alpha} - x_{\alpha} \sum_{\beta} \rho_{\alpha\beta}, \quad (2.5)$$

with $\rho_{\alpha\alpha}$ defined arbitrarily.

In what follows, we will focus on revision protocols of the general “imitative” form

$$\rho_{\alpha\beta}(v, x) = x_{\beta}r_{\alpha\beta}(v, x), \quad (2.6)$$

corresponding to the case where a player imitates the strategy of a uniformly drawn opponent with probability proportional to the so-called *conditional imitation rate* $r_{\alpha\beta}$ (assumed Lipschitz). In particular, one of the most widely studied revision protocols of this type is the “imitation of success” protocol (Weibull, 1995) where the imitation rate of a given target strategy is proportional to its payoff,³ i.e.

$$r_{\alpha\beta}(v, x) = v_{\beta} \quad (2.7)$$

On account of (2.5), the mean evolutionary dynamics induced by (2.7) take the form:

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

which is simply the classical replicator equation of Taylor and Jonker (1978).

The replicator dynamics have attracted significant interest in the literature and their long-run behavior is relatively well understood. For instance, Akin (1980), Nachbar (1990) and Samuelson and Zhang (1992) showed that dominated strategies become extinct under (RD), whereas the (multi-population) “folk theorem” of evolutionary game theory (Hofbauer and Sigmund, 2003) states that *a)* (Lyapunov) stable states are Nash; *b)* limits of interior trajectories are Nash; and *c)* strict Nash equilibria are asymptotically stable under (RD).

2.3. Payoff shocks and the induced dynamics. Our main goal in this paper is to investigate the rationality properties of the replicator dynamics in a setting where the players' payoffs are subject to exogenous stochastic disturbances. To model these “payoff shocks”, we assume that the players' payoffs at time t are of the

²In other words, $\rho_{\alpha\beta}$ is the probability of an α -strategist becoming a β -strategist up to normalization by the alarm clocks' rate.

³Modulo an additive constant which ensures that ρ is positive but which cancels out when it comes to the dynamics.

form $\hat{v}_\alpha(t) = v_\alpha(x(t)) + \xi_\alpha(t)$ for some zero-mean “white noise” process ξ_α . Then, in Langevin notation, the replicator dynamics (RD) become:

$$\begin{aligned} \frac{dX_\alpha}{dt} &= X_\alpha \left[\hat{v}_\alpha - \sum_\beta X_\beta \hat{v}_\beta \right] \\ &= X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] + X_\alpha \left[\xi_\alpha - \sum_\beta X_\beta \xi_\beta \right], \end{aligned} \quad (2.8)$$

or, in stochastic differential equation (SDE) form:

$$\begin{aligned} dX_\alpha &= X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] dt \\ &\quad + X_\alpha \left[\sigma_\alpha(X) dW_\alpha - \sum_\beta X_\beta \sigma_\beta(X) dW_\beta \right], \end{aligned} \quad (\text{SRD})$$

where the diffusion coefficients $\sigma_\alpha: \mathcal{X} \rightarrow \mathbb{R}$ (assumed Lipschitz) measure the intensity of the payoff shocks and the Wiener processes W_α are assumed independent.

The stochastic dynamics (SRD) will constitute the main focus of this paper, so some remarks are in order:

Remark 2.2. With v and σ assumed Lipschitz, it follows that (SRD) admits a unique (strong) solution $X(t)$ for every initial condition $X(0) \in \mathcal{X}$. Moreover, since the drift and diffusion terms of (SRD) all vanish at the boundary $\text{bd}(\mathcal{X})$ of \mathcal{X} , standard arguments can be used to show that these solutions exist (a.s.) for all time, and that $X(t) \in \mathcal{X}^\circ$ for all $t \geq 0$ if $X(0) \in \mathcal{X}^\circ$ (Khasminskii, 2012; Øksendal, 2007).

Remark 2.3. The independence assumption for the Wiener processes W_α can be relaxed without qualitatively affecting our analysis;⁴ in particular, as we shall see in the proofs of our results, the rationality properties of (SRD) can be formulated directly in terms of the quadratic (co)variation of the noise processes W_α . Doing so however would complicate the relevant expressions considerably, so, for clarity, we will retain this independence assumption throughout our paper.

Remark 2.4. The deterministic replicator dynamics (RD) are also the governing dynamics for the “pairwise proportional imitation” revision protocol (Schlag, 1998) where a revising agent imitates the strategy of a randomly chosen opponent only if the opponent’s payoff is higher than his own, and he does so with probability proportional to the payoff difference. Formally, the conditional switch rate $\rho_{\alpha\beta}$ under this revision protocol is:

$$\rho_{\alpha\beta} = x_\beta [v_\beta - v_\alpha]_+, \quad (2.9)$$

where $[x]_+ = \max\{x, 0\}$ denotes the positive part of x . Accordingly, if the game’s payoffs at time t are of the perturbed form $\hat{v}_\alpha(t) = v_\alpha(x(t)) + \xi_\alpha(t)$ as before, (2.5) leads to the master stochastic equation:

$$\begin{aligned} \dot{X}_\alpha &= \sum_\beta X_\beta X_\alpha [\hat{v}_\alpha - \hat{v}_\beta]_+ - X_\alpha \sum_\beta X_\beta [\hat{v}_\beta - \hat{v}_\alpha]_+ \\ &= X_\alpha \sum_\beta X_\beta \left\{ [\hat{v}_\alpha - \hat{v}_\beta]_+ - [\hat{v}_\beta - \hat{v}_\alpha]_+ \right\} = X_\alpha \sum_\beta X_\beta (\hat{v}_\alpha - \hat{v}_\beta) \\ &= X_\alpha \left[\hat{v}_\alpha - \sum_\beta X_\beta \hat{v}_\beta \right], \end{aligned} \quad (2.10)$$

⁴An important special case where it makes sense to consider correlated shocks is if the payoff functions $v_\alpha(x)$ are derived from random matchings in a finite game whose payoff matrix is subject to stochastic perturbations. This specific disturbance model is discussed in Section 5.

which is simply the stochastic replicator dynamics (2.8). In other words, (SRD) could also be interpreted as the mean dynamics of a pairwise imitation process with perturbed payoff comparisons as above.

2.4. Related stochastic models. The replicator dynamics were first introduced in biology, as a model of frequency-dependent selection. They arise from the geometric population growth equation:

$$\dot{z}_\alpha = z_\alpha v_\alpha \quad (2.11)$$

where z_α denotes the absolute population size of the α -th genotype of a given species.⁵ This biological model was also the starting point of Fudenberg and Harris (1992) who added aggregate payoff shocks to (2.11) based on the geometric Brownian model:

$$dZ_\alpha = Z_\alpha [v_\alpha dt + \sigma_\alpha dW_\alpha], \quad (2.12)$$

where the diffusion process $\sigma_\alpha dW_\alpha$ represents the impact of random, weather-like effects on the genotype's fitness (see also Cabrales, 2000; Hofbauer and Imhof, 2009; Imhof, 2005).⁶ Itô's lemma applied to the population shares $X_\alpha = Z_\alpha / \sum_\beta Z_\beta$ then yields the *replicator dynamics with aggregate shocks*:

$$\begin{aligned} dX_\alpha &= X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] dt \\ &+ X_\alpha \left[\sigma_\alpha dW_\alpha - \sum_\beta \sigma_\beta X_\beta dW_\beta \right] \\ &- X_\alpha \left[\sigma_\alpha^2 X_\alpha - \sum_\beta \sigma_\beta^2 X_\beta^2 \right] dt. \end{aligned} \quad (2.13)$$

In a repeated game context, the replicator dynamics also arise from a continuous-time variant of the exponential weight algorithm introduced by Vovk (1990) and Littlestone and Warmuth (1994) (see also Sorin, 2009). In particular, if players follow the exponential learning scheme:

$$\begin{aligned} dy_\alpha &= v_\alpha(x) dt, \\ x_\alpha &= \frac{\exp(y_\alpha)}{\sum_\beta \exp(y_\beta)}, \end{aligned} \quad (2.14)$$

that is, if they play a logit best response to the vector of their cumulative payoffs, then the frequencies x_α follow (RD).⁷ Building on this, Mertikopoulos and Moustakas (2009, 2010) considered the stochastically perturbed exponential learning scheme:

$$\begin{aligned} dY_\alpha &= v_\alpha(X) dt + \sigma_\alpha(X) dW_\alpha, \\ X_\alpha &= \frac{\exp(Y_\alpha)}{\sum_\beta \exp(Y_\beta)}, \end{aligned} \quad (2.15)$$

where the cumulative payoffs are perturbed by the observation noise process $\sigma_\alpha dW_\alpha$. By Itô's lemma, we then obtain the stochastic replicator dynamics of exponential

⁵The replicator equation (RD) is obtained simply by computing the evolution of the frequencies $x_\alpha = z_\alpha / \sum_\beta z_\beta$ under (2.11).

⁶Khasminskii and Potsepun (2006) also considered a related evolutionary model with Stratonovich-type perturbations while, more recently, Vlasic (2012) studied the effect of discontinuous semimartingale shocks incurred by catastrophic, earthquake-like events.

⁷The intermediate variable y_α should be thought of as an evaluation of how good the strategy α is, and the formula for x_α as a way of transforming these evaluations into a strategy.

learning:

$$\begin{aligned} dX_\alpha &= X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] dt \\ &\quad + X_\alpha \left[\sigma_\alpha dW_\alpha - \sum_\beta \sigma_\beta X_\beta dW_\beta \right] \\ &\quad + \frac{X_\alpha}{2} \left[\sigma_\alpha^2 (1 - 2X_\alpha) - \sum_\beta \sigma_\beta^2 X_\beta (1 - 2X_\beta) \right] dt. \end{aligned} \quad (2.16)$$

Besides their very distinct origins, a key difference between the stochastic replicator dynamics (SRD) and the stochastic models (2.13)/(2.16) is that there is no Itô correction term in the former. The reason for this is that in (2.13) and (2.16), the noise affects primarily the evolution of an intermediary variable (the absolute population sizes Z_α and the players' cumulative payoffs Y_α respectively) before being carried over to the evolution of the strategy shares X_α . By contrast, the payoff shocks that impact the players' revision protocol in (SRD) affect the corresponding strategy shares directly, so there is no intervening Itô correction.

The pure noise case. To better understand the differences between our model and previous models of stochastic replicator dynamics, it is useful to consider the case of pure noise, that is, when the expected payoff of each strategy is equal to one and the same constant C : $v_\alpha(x) = C$ for all $\alpha \in \mathcal{A}$ and for all $x \in \mathcal{X}$.

For simplicity, let us also assume that $\sigma_\alpha(x)$ is independent of the state of the population x . Eq. (2.12) then becomes a simple geometric Brownian motion of the form:

$$dZ_\alpha = Z_\alpha [C dt + \sigma_\alpha dW_\alpha], \quad (2.17)$$

which readily yields $Z_\alpha(t) = Z_\alpha(0) \exp((C - \sigma_\alpha^2/2)t + \sigma_\alpha W_\alpha(t))$. The corresponding frequency $X_\alpha = Z_\alpha / \sum_\beta Z_\beta$ will then be:

$$X_\alpha(t) = \frac{X_\alpha(0) \exp(-\frac{1}{2}\sigma_\alpha^2 t + \sigma_\alpha W_\alpha(t))}{\sum_\beta X_\beta(0) \exp(-\frac{1}{2}\sigma_\beta^2 t + \sigma_\beta W_\beta(t))}. \quad (2.18)$$

If $\sigma_\alpha \neq 0$, the law of large numbers yields $-\frac{1}{2}\sigma_\alpha^2 t + \sigma_\alpha W_\alpha(t) \sim -\frac{1}{2}\sigma_\alpha^2 t$ (a.s.). Therefore, letting $\sigma_{\min} = \min_{\alpha \in \mathcal{A}} \sigma_\alpha$, it follows from (2.18) that strategy $\alpha \in \mathcal{A}$ is eliminated if $\sigma_\alpha > \sigma_{\min}$ and survives if $\sigma_\alpha = \sigma_{\min}$ (a.s.).⁸ In particular, if all intensities are equal ($\sigma_\alpha = \sigma_{\min}$ for all $\alpha \in \mathcal{A}$), then all strategies survive and the share of each strategy oscillates for ever, occasionally taking values arbitrarily close to 0 and arbitrarily close to 1. On the other hand, under the stochastic replicator dynamics of exponential learning for the pure noise case, (2.16) readily yields:

$$X_\alpha(t) = \frac{X_\alpha(0) \exp(\sigma_\alpha W_\alpha(t))}{\sum_\beta X_\beta(0) \exp(\sigma_\beta W_\beta(t))}. \quad (2.19)$$

Therefore, for any value of the diffusion coefficients σ_α (and, in particular, even if some strategies are affected by noise much more than others), all pure strategies survive.

Our model behaves differently from both (2.13) and (2.16): in the pure noise case, for any value of the noise coefficients σ_α (as long as $\sigma_\alpha > 0$ for all α), only a single strategy survives (a.s.), and strategy α survives with probability equal to

⁸Elimination is obvious; for survival, simply add $\frac{1}{2}\sigma_{\min}^2 t$ to the exponents of (2.18) and recall that any Wiener process has $\limsup_t W(t) > 0$ and $\liminf_t W(t) < 0$ (a.s.).

$X_\alpha(0)$. To see this, consider first the model with pure noise and only two strategies, α and β . Then, letting $X(t) = X_\alpha(t)$ (so $X_\beta(t) = 1 - X(t)$), we get:

$$dX(t) = X(t)(1 - X(t)) [\sigma_\alpha dW_\alpha - \sigma_\beta dW_\beta] = X(t)(1 - X(t)) \sigma dW(t), \quad (2.20)$$

where $\sigma^2 = \sigma_\alpha^2 + \sigma_\beta^2$ and we have used the time-change theorem for martingales to write $\sigma dW = \sigma_\alpha dW_\alpha - \sigma_\beta dW_\beta$ for some Wiener process $W(t)$. This diffusion process can be seen as a continuous-time random walk on $[0, 1]$ with step sizes that get smaller as X approaches $\{0, 1\}$. Thus, at a heuristic level, when $X(t)$ starts close to $X = 1$ and takes one step to the left followed by one step to the right (or the opposite), the walk does not return to its initial position, but will approach 1 (of course, the same phenomenon occurs near 0). This suggests that the process should eventually converge to one of the vertices: indeed, letting $f(x) = \log x(1 - x)$, Itô's lemma yields

$$df(X) = (1 - 2X) \sigma dW - \frac{1}{2} [(1 - X)^2 + X^2] \sigma^2 dt \leq (1 - 2X) \sigma dW - \frac{1}{4} \sigma^2 dt, \quad (2.21)$$

so, by Lemma A.1, we get $\lim_{t \rightarrow \infty} f(X(t)) = 0$ (a.s.), that is, $\lim_{t \rightarrow \infty} X(t) \in \{0, 1\}$.

More generally, consider the model with pure noise and n strategies. Then, computing $d[\log X_\alpha(1 - X_\alpha)]$ as above, we readily obtain $\lim_{t \rightarrow \infty} X_\alpha(t) \in \{0, 1\}$ (a.s.), for every strategy $\alpha \in \mathcal{A}$ with $\sigma_\alpha > 0$. Since X_α is a martingale, we will have $\mathbb{E}[X_\alpha(t)] = X_\alpha(0)$ for all $t \geq 0$,⁹ so $X_\alpha \rightarrow 1$ with probability $X_\alpha(0)$ and $X_\alpha(t) \rightarrow 0$ with probability $1 - X_\alpha(0)$.¹⁰

The above highlights two important differences between our model and the stochastic replicator dynamics of Fudenberg and Harris (1992). First, in our model, noise is not detrimental in itself: in the pure noise case, the expected frequency of a strategy remains constant, irrespective of the noise level; by contrast, in the model of Fudenberg and Harris (1992), the expected frequency of strategies affected by strong payoff noise decreases.¹¹ Second, our model behaves in a somewhat more “unpredictable” way: for instance, in the model of Fudenberg and Harris (1992), when there are only two strategies with the same expected payoff, and if one of the strategies is affected by a stronger payoff noise, then it will be eliminated (a.s.); in our model, we cannot say in advance whether it will be eliminated or not.

⁹We are implicitly assuming here deterministic initial conditions, i.e. $X(0) = x$ (a.s.) for some $x \in \mathcal{X}$.

¹⁰If several strategies are unaffected by noise, that is, are such that $\sigma_\alpha = 0$, then their relative shares remain constant (that is, if α and β are two such strategies, then $X_\alpha(t)/X_\beta(t) = X_\alpha(0)/X_\beta(0)$ for all $t \geq 0$). It follows from this observation and the above result that, almost surely, all these strategies are eliminated or all these strategies survive (and only them).

¹¹In the pure noise case of the model of Fudenberg and Harris (1992), what remains constant is the expected number of individuals playing a strategy. A crucial point here is that this number may grow to infinity. What happens to strategies affected by large aggregate shocks is that with small probability, the total number of individuals playing this strategy gets huge, but with a large probability (going to 1), it gets small (at least compared to the number of individuals playing other strategies). This can be seen as a gambler's ruin phenomenon, which explains that even with a higher expected payoff than others (hence a higher expected subpopulation size), the frequency of a strategy may go to zero almost surely (see e.g. Robson and Samuelson, 2011, Sec 3.1.1). This cannot happen in our model since noise is added directly to the frequencies (which are bounded).

3. LONG-TERM RATIONALITY ANALYSIS

In this section, we investigate the long-run rationality properties of the stochastic dynamics (SRD); in particular, we focus on the elimination of dominated strategies and the stability of equilibrium play.

3.1. Elimination of dominated strategies. We begin with the elimination of dominated strategies. Formally, given a trajectory of play $x(t) \in \mathcal{X}$, we say that a pure strategy $\alpha \in \mathcal{A}$ becomes extinct along $x(t)$ if $x_\alpha(t) \rightarrow 0$ as $t \rightarrow \infty$. More generally, following Samuelson and Zhang (1992), we will say that the mixed strategy $p \in \mathcal{X}$ becomes extinct along $x(t)$ if $\min\{x_\alpha(t) : \alpha \in \text{supp}(p)\} \rightarrow 0$ as $t \rightarrow \infty$; otherwise, we say that p survives.

Now, with a fair degree of hindsight, it will be convenient to introduce a modified game $\mathfrak{G}^\sigma \equiv \mathfrak{G}^\sigma(\mathcal{N}, \mathcal{A}, v^\sigma)$ with payoff functions v_α^σ adjusted for noise as follows:

$$v_\alpha^\sigma(x) = v_\alpha(x) - \frac{1}{2}(1 - 2x_\alpha)\sigma_\alpha^2(x). \quad (3.1)$$

Imhof (2005) introduced a similar modified game to study the long-term convergence and stability properties of the stochastic replicator dynamics with aggregate shocks (2.13) and showed that strategies that are dominated in this modified game are eliminated (a.s.) – cf. Remark 3.7 below. Our main result concerning the elimination of dominated strategies under (SRD) is of a similar nature:

Theorem 3.1. *Let $X(t)$ be an interior solution orbit of the stochastic replicator dynamics (SRD). Assume further that $p \in \mathcal{X}$ is dominated by $p' \in \mathcal{X}$ in the modified game \mathfrak{G}^σ . Then, p becomes extinct along $X(t)$ (a.s.).*

Remark 3.1. As a special case, if the (pure) strategy $\alpha \in \mathcal{A}$ is dominated by the (pure) strategy $\beta \in \mathcal{A}$, Theorem 3.1 shows that α becomes extinct under (SRD) as long as

$$v_\beta(x) - v_\alpha(x) > \frac{1}{2} [\sigma_\alpha^2(x) + \sigma_\beta^2(x)] \quad \text{for all } x \in \mathcal{X}. \quad (3.2)$$

In terms of the original game, this condition can be interpreted as saying that α is dominated by β by a margin no less than $\frac{1}{2} \max_x (\sigma_\alpha^2(x) + \sigma_\beta^2(x))$. Put differently, Theorem 3.1 shows that dominated strategies in the original, unmodified game become extinct provided that the payoff shocks are mild enough.

Proof of Theorem 3.1. Following Cabrales (2000), we will show that p becomes extinct along $X(t)$ by studying the “cross-entropy” function:

$$V(x) = D_{\text{KL}}(p, x) - D_{\text{KL}}(p', x) = \sum_\alpha (p_\alpha \log p_\alpha - p'_\alpha \log p'_\alpha) + \sum_\alpha (p'_\alpha - p_\alpha) \log x_\alpha, \quad (3.3)$$

where $D_{\text{KL}}(p, x) = \sum_\alpha p_\alpha \log(p_\alpha/x_\alpha)$ denotes the Kullback–Leibler (KL) divergence of x with respect to p . By a standard argument (Weibull, 1995), p becomes extinct along $X(t)$ if $\lim_{t \rightarrow \infty} D_{\text{KL}}(p, X(t)) = \infty$; thus, with $D_{\text{KL}}(p', x) \geq 0$, it suffices to show that $\lim_{t \rightarrow \infty} V(X(t)) = \infty$.

To that end, let $Y_\alpha = \log X_\alpha$ so that

$$dY_\alpha = \frac{dX_\alpha}{X_\alpha} - \frac{1}{2} \frac{1}{X_\alpha^2} (dX_\alpha)^2, \quad (3.4)$$

by Itô's lemma. Then, writing $dS_\alpha = X_\alpha [\sigma_\alpha dW_\alpha - \sum_\beta X_\beta \sigma_\beta dW_\beta]$ for the martingale term of (SRD), we readily obtain:

$$\begin{aligned} (dS_\alpha)^2 &= X_\alpha^2 \left[\sigma_\alpha dW_\alpha - \sum_\beta X_\beta \sigma_\beta dW_\beta \right] \cdot \left[\sigma_\alpha dW_\alpha - \sum_\gamma X_\gamma \sigma_\gamma dW_\gamma \right] \\ &= X_\alpha^2 \left[(1 - 2X_\alpha) \sigma_\alpha^2 + \sum_\beta \sigma_\beta^2 X_\beta^2 \right] dt, \end{aligned} \quad (3.5)$$

where we have used the orthogonality conditions $dW_\beta \cdot dW_\gamma = \delta_{\beta\gamma} dt$. By the same token, we also get $(dX_\alpha)^2 = (dS_\alpha)^2$, and hence:

$$\begin{aligned} dY_\alpha &= (v_\alpha - \langle v|X \rangle) dt - \frac{1}{2} \left[(1 - 2X_\alpha) \sigma_\alpha^2 + \sum_\beta \sigma_\beta^2 X_\beta^2 \right] dt \\ &\quad + \sigma_\alpha dW_\alpha - \sum_\beta X_\beta \sigma_\beta dW_\beta. \end{aligned} \quad (3.6)$$

Therefore, after some easy algebra, we obtain:

$$\begin{aligned} dV &= \sum_\alpha (p'_\alpha - p_\alpha) dY_\alpha \\ &= \langle v(X) | p' - p \rangle dt \\ &\quad - \frac{1}{2} \sum_\alpha (p'_\alpha - p_\alpha) (1 - 2X_\alpha) \sigma_\alpha^2(X) dt + \sum_\alpha (p'_\alpha - p_\alpha) \sigma_\alpha(X) dW_\alpha \\ &= \langle v^\sigma(X) | p' - p \rangle dt + \sum_\alpha (p'_\alpha - p_\alpha) \sigma_\alpha(X) dW_\alpha \end{aligned} \quad (3.7)$$

where we have used the fact that $\sum_\alpha (p'_\alpha - p_\alpha) = 0$.

Now, since p is dominated by p' in \mathfrak{G}^σ , we will have $\langle v^\sigma(x) | p' - p \rangle \geq m$ for some positive constant $m > 0$ and for all $x \in \mathcal{X}$. Eq. (4.6) then yields:

$$V(X(t)) \geq V(X(0)) + mt + \xi(t), \quad (3.8)$$

where ξ denotes the martingale part of (4.6), viz.

$$\xi(t) = \sum_\alpha (p'_\alpha - p_\alpha) \int_0^t \sigma_\alpha(X(s)) dW_\alpha(s). \quad (3.9)$$

Since the $\sigma(X(t))$ is bounded and continuous (a.s.), Lemma A.1 shows that $mt + \xi(t) \sim mt$ as $t \rightarrow \infty$, so the RHS of (3.8) escapes to ∞ as $t \rightarrow \infty$. This implies $\lim_{t \rightarrow \infty} V(X(t)) = \infty$ and our proof is complete. \square

Theorem 3.1 is our main result concerning the extinction of dominated strategies under (SRD) so a few remarks are in order:

Remark 3.2. Theorem 3.1 is analogous to the elimination results of Imhof (2005, Theorem 3.1) and Cabrales (2000, Prop. 1A) who show that dominated strategies become extinct under the replicator dynamics with aggregate shocks (2.13) if the shocks satisfy certain ‘‘tameness’’ requirements. On the other hand, Theorem 3.1 should be contrasted to the corresponding results of Mertikopoulos and Moustakas (2010) who showed that dominated strategies become extinct under the stochastic replicator dynamics of exponential learning (2.16) *irrespective* of the noise level (for a related elimination result, see also Bravo and Mertikopoulos, 2014). The crucial qualitative difference here lies in the Itô correction term that appears in the drift of the stochastic replicator dynamics: the Itô correction in (2.16) is ‘‘just right’’ with respect to the logarithmic variables $Y_\alpha = \log X_\alpha$ and this is what leads to the unconditional elimination of dominated strategies. On the other hand, even though there is no additional drift term in (SRD) except for the one driven by the game's

payoffs, the logarithmic transformation $Y_\alpha = \log X_\alpha$ incurs an Itô correction which is reflected in the definition of the modified payoff functions (3.1).

Remark 3.3. A standard induction argument based on the rounds of elimination of iteratively dominated strategies (see e.g. Cabrales, 2000 or Mertikopoulos and Moustakas, 2010) can be used to show that the only strategies that survive under the stochastic replicator dynamics (SRD) must be iteratively undominated in the modified game \mathfrak{G}^σ .

Remark 3.4. Finally, it is worth mentioning that Imhof (2005) also establishes an exponential rate of extinction of dominated strategies under the stochastic replicator dynamics with aggregate shocks (2.13). Specifically, if $\alpha \in \mathcal{A}$ is dominated, Imhof (2005) showed that there exist constants $A, B > 0$ and $A', B' > 0$ such that

$$X_\alpha(t) = o\left(\exp\left(-At + B\sqrt{t \log \log t}\right)\right) \quad (\text{a.s.}), \quad (3.10)$$

and

$$\mathbb{P}[X_\alpha(t) > \varepsilon] \leq \frac{1}{2} \operatorname{erfc}\left[A't^{1/2} + B' \log \varepsilon \cdot t^{-1/2}\right], \quad (3.11)$$

provided that the noise coefficients of (2.13) satisfy a certain ‘‘tameness’’ condition. Following the same reasoning, it is possible to establish similar exponential decay rates for the elimination of dominated strategies under (SRD), but the exact expressions for the constants in (3.10) and (3.11) are more complicated, so we do not present them here.

3.2. Stability analysis of equilibrium play. In this section, our goal will be to investigate the stability and convergence properties of the stochastic replicator dynamics (SRD) with respect to equilibrium play. Motivated by a collection of stability results that is sometimes called the ‘‘folk theorem’’ of evolutionary game theory (Hofbauer and Sigmund, 2003), we will focus on the following three properties of the deterministic replicator dynamics (RD):

- (1) Limits of interior orbits are Nash equilibria.
- (2) Lyapunov stable states are Nash equilibria.
- (3) Strict Nash equilibria are asymptotically stable under (RD).

Of course, given the stochastic character of the dynamics (SRD), the notions of Lyapunov and asymptotic stability must be suitably modified. In this SDE context, we have:

Definition 3.2. Let $x^* \in \mathcal{X}$. We will say that:

- (1) x^* is *stochastically Lyapunov stable* under (SRD) 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

$$\mathbb{P}(X(t) \in U_0 \text{ for all } t \geq 0) \geq 1 - \varepsilon \quad \text{whenever } X(0) \in U. \quad (3.12)$$

- (2) x^* is *stochastically asymptotically stable* under (SRD) 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

$$\mathbb{P}\left(X(t) \in U_0 \text{ for all } t \geq 0 \text{ and } \lim_{t \rightarrow \infty} X(t) = x^*\right) \geq 1 - \varepsilon \quad \text{whenever } X(0) \in U. \quad (3.13)$$

For (SRD), we have:

Theorem 3.3. *Let $X(t)$ be an interior solution orbit of the stochastic replicator dynamics (SRD) and let $x^* \in \mathcal{X}$.*

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

Remark 3.5. By the nature of the modified payoff functions (3.1), strict equilibria of the original game \mathfrak{G} are also strict equilibria of \mathfrak{G}^σ , so Theorem 3.3 implies that strict equilibria of \mathfrak{G} are also stochastically asymptotically stable under the stochastic dynamics (SRD). The converse does not hold: if the noise coefficients σ_α are sufficiently large, (SRD) possesses stochastically asymptotically stable states that are not Nash equilibria of \mathfrak{G} . This is consistent with the behavior of (SRD) in the pure noise case that we discussed in the previous section: if $X(t)$ starts within ε of a vertex of \mathcal{X} and there are no payoff differences, then $X(t)$ converges to this vertex with probability at least $1 - \varepsilon$.

Remark 3.6. The condition for α to be a strict equilibrium of the modified game is that

$$v_\beta - v_\alpha < \frac{1}{2} (\sigma_\alpha^2 + \sigma_\beta^2) \quad \text{for all } \beta \neq \alpha, \quad (3.14)$$

where the payoffs and the noise coefficients are evaluated at the vertex e_α of \mathcal{X} (note the similarity with (3.2)). To provide some intuition for this condition, consider the case of only two pure strategies, α and β , and assume constant noise coefficients. Letting $X(t) = X_\beta(t)$ and proceeding as in (2.20), we get $dX = X(1 - X)[(v_\beta - v_\alpha) dt - \sigma dW]$ where $\sigma^2 = \sigma_\alpha^2 + \sigma_\beta^2$ and W is a rescaled Wiener process. Heuristically, a discrete-time counterpart of $X(t)$ is then provided by the random walk:

$$X(n+1) - X(n) = X(n)(1 - X(n)) \left[(v_\beta - v_\alpha)\delta + \sigma \xi_n \sqrt{\delta} \right] \quad (3.15)$$

where $\xi_n \in \{+1, -1\}$ is a zero-mean Bernoulli process, and the noise term is multiplied by $\sqrt{\delta}$ instead of δ because $dW \cdot dW = dt$. For small X and δ , a simple computation then shows that, in the event $\xi_{n+1} = -\xi_n$, we have:

$$X(n+2) - X(n) = 2\delta X(n) \left[v_\beta - v_\alpha - \frac{1}{2}\sigma^2 \right] + o(\delta) + o(X(n)). \quad (3.16)$$

Since $\sigma^2 = \sigma_\alpha^2 + \sigma_\beta^2$, the bracket is negative (so $X_\alpha = 1 - X$ increases) if and only if condition (3.14) is satisfied. Thus, (3.14) may be interpreted as saying that when the discrete-time process $X(n)$ is close to e_α and the random noise term ξ_n takes two successive steps in opposite direction, then the process ends up even closer to e_α .¹² On the other hand, if the opposite strict inequality holds, then this interpretation suggests that β should successfully invade a population where most individuals play α – which, in turn, explains (3.2).

¹²Put differently, it's more probable for $X(n)$ to decrease rather than increase: $X(n+2) > X(n)$ with probability 1/4 (i.e. if and only if ξ_n takes two positive steps), while $X(n+2) < X(n)$ with probability 3/4.

Proof of Theorem 3.3. Contrary to the approach of Hofbauer and Imhof (2009), we will not employ the stochastic Lyapunov method (see e.g. Khasminskii, 2012) which requires calculating the infinitesimal generator of (SRD). Instead, motivated by the recent analysis of Bravo and Mertikopoulos (2014), our proof will rely on the “dual” variables $Y_\alpha = \log X_\alpha$ that were already used in the proof of Theorem 3.1.

Part 1. We argue by contradiction. Indeed, assume that $\mathbb{P}(\lim_{t \rightarrow \infty} X(t) = x^*) > 0$ but that x^* is not Nash for the noise-adjusted game \mathfrak{G}^σ , so $v_\alpha^\sigma(x^*) < v_\beta^\sigma(x^*)$ for some $\alpha \in \text{supp}(x^*)$, $\beta \in \mathcal{A}$. On that account, let U be a sufficiently small neighborhood of x^* in \mathcal{X} such that $v_\beta^\sigma(x) - v_\alpha^\sigma(x) \geq m$ for some $m > 0$ and for all $x \in U$. Then, by (3.6), we get:

$$\begin{aligned} dY_\alpha - dY_\beta &= [v_\alpha - v_\beta] dt - \frac{1}{2} [(1 - 2X_\alpha)\sigma_\alpha^2 - (1 - 2X_\beta)\sigma_\beta^2] dt \\ &\quad + \sigma_\alpha dW_\alpha - \sigma_\beta dW_\beta, \end{aligned} \quad (3.17)$$

so, if $X(t)$ is an interior orbit of (SRD) that converges to x^* , we will have:

$$dY_\alpha - dY_\beta \leq -m dt - d\xi \quad \text{for all large enough } t > 0, \quad (3.18)$$

where ξ denoting the martingale part of (3.17). Since the diffusion coefficients of (3.17) are bounded, Lemma A.1 shows that $mt + \xi(t) \sim mt$ for large t (a.s.), so

$$\log \frac{X_\alpha(t)}{X_\beta(t)} \leq \log \frac{X_\alpha(0)}{X_\beta(0)} - mt - \xi(t) \sim -mt \rightarrow -\infty \quad (\text{a.s.}) \quad (3.19)$$

as $t \rightarrow \infty$. This implies that $\lim_{t \rightarrow \infty} X_\alpha(t) = 0$, contradicting our original assumption that $X(t)$ stays in a small enough neighborhood of x^* with positive probability (recall that $x_\alpha^* > 0$); we thus conclude that x^* is a Nash equilibrium of the noise-adjusted game \mathfrak{G}^σ , as claimed.

Part 2. Assume that x^* is stochastically Lyapunov stable. Then, every neighborhood U of x^* admits an interior trajectory $X(t)$ that stays in U for all time with positive probability. The proof of Part 1 shows that this is only possible if x^* is a Nash equilibrium of the modified \mathfrak{G}^σ , so our claim follows.

Part 3. To show that strict Nash equilibria of \mathfrak{G}^σ are stochastically asymptotically stable, let $x^* = (\alpha_1^*, \dots, \alpha_N^*) \in \mathcal{X}$ be a strict equilibrium of \mathfrak{G}^σ . Then, suppressing the population index k as before, let

$$Z_\alpha = Y_\alpha - Y_{\alpha^*}, \quad (3.20)$$

so that $X(t) \rightarrow x^*$ if and only if $Z_\alpha(t) \rightarrow -\infty$ for all $\alpha \in \mathcal{A}^* \equiv \mathcal{A} \setminus \{\alpha^*\}$.¹³

To proceed, fix some probability threshold $\varepsilon > 0$ and a neighborhood U_0 of x^* in \mathcal{X} . Since x^* is a strict equilibrium of \mathfrak{G}^σ , there exists a neighborhood $U \subseteq U_0$ of x^* and some $m > 0$ such that

$$v_{\alpha^*}^\sigma(x) - v_\alpha^\sigma(x) \geq m \quad \text{for all } x \in U \text{ and for all } \alpha \in \mathcal{A}^*. \quad (3.21)$$

Let $M > 0$ be sufficiently large so that $X(t) \in U$ if $Z_\alpha(t) \leq -M$ for all $\alpha \in \mathcal{A}^*$; we will show that if M is chosen suitably (in terms of ε) and $Z_\alpha(0) < -2M$, then $X(t) \in U$ for all $t \geq 0$ and $Z_\alpha(t) \rightarrow -\infty$ with probability at least $1 - \varepsilon$, i.e. x^* is stochastically asymptotically stable.

To that end, take $Z_\alpha(0) \leq -2M$ in (3.20) and define the first exit time:

$$\tau_U = \inf\{t > 0 : X(t) \notin U\}. \quad (3.22)$$

¹³Simply note that $X_{\alpha^*} = (1 + \sum_{\beta \in \mathcal{A}^*} \exp(Z_\beta))^{-1}$.

By applying (3.17), we then get:

$$dZ_\alpha = dY_\alpha - dY_{\alpha^*} = [v_\alpha^\sigma - v_{\alpha^*}^\sigma] dt - d\xi, \quad (3.23)$$

where the martingale term $d\xi$ is defined as in (3.17), taking $\beta = \alpha^*$. Hence, for all $t \leq \tau_U$, we will have:

$$Z_\alpha(t) = Z_\alpha(0) + \int_0^t [v_\alpha^\sigma(X(s)) - v_{\alpha^*}^\sigma(X(s))] ds - \xi(t) \leq -2M - mt - \xi(t). \quad (3.24)$$

By the time-change theorem for martingales (Øksendal, 2007, Cor. 8.5.4), there exists a standard Wiener process $\widetilde{W}(t)$ such that $\xi(t) = \widetilde{W}(\rho(t))$ where $\rho = [\xi, \xi]$ denotes the quadratic variation of ξ ; as such, we will have $Z_\alpha(t) \leq -M$ whenever $\widetilde{W}(\rho(t)) \geq -M - mt$. However, with σ Lipschitz over \mathcal{X} , we readily get $\rho(t) \leq Kt$ for some positive constant $K > 0$, so it suffices to show that the hitting time

$$\tau_0 = \inf \{t > 0 : \widetilde{W}(t) = -M - mt/K\} \quad (3.25)$$

is finite with probability not exceeding ε . Indeed, if a trajectory of $\widetilde{W}(t)$ has $\widetilde{W}(t) \geq -M - mt/K$ for all $t \geq 0$, we will also have

$$\widetilde{W}(\rho(t)) \geq -M - m\rho(t)/K \geq -M - mt, \quad (3.26)$$

so τ_U is infinite for every trajectory of \widetilde{W} with infinite τ_0 , hence $\mathbb{P}(\tau_U < +\infty) \leq \mathbb{P}(\tau_0 < +\infty)$. Lemma A.2 then shows that $\mathbb{P}(\tau_0 < +\infty) = e^{-2Mm/K}$, so, if we take $M > -(2m)^{-1}K \log \varepsilon$, we get $\mathbb{P}(\tau_U = \infty) \geq 1 - \varepsilon$. Conditioning on the event $\tau_U = +\infty$, Lemma A.1 applied to (3.24) yields

$$Z_\alpha(t) \leq -2M - mt - \xi(t) \sim -mt \rightarrow -\infty \quad (\text{a.s.}) \quad (3.27)$$

so $X(t) \rightarrow x^*$ with probability at least $1 - \varepsilon$, as was to be shown. \square

Remark 3.7. As mentioned before, Hofbauer and Imhof (2009) state a similar “evolutionary folk theorem” in the context of single-population random matching games under the stochastic replicator dynamics with aggregate shocks (2.13). In particular, Hofbauer and Imhof (2009) consider the modified game:

$$v_\alpha^\sigma(x) = v_\alpha(x) - \frac{1}{2}\sigma_\alpha^2, \quad (3.28)$$

where σ_α denotes the intensity of the aggregate shocks in (2.13), and they show that strict Nash equilibria of this noise-adjusted game are stochastically asymptotically stable under (2.13). It is interesting to note that the adjustments (3.1) and (3.28) do not coincide: the payoff shocks affect the deterministic replicator equation (RD) in a different way than the aggregate shocks of (2.13). Heuristically, in the model of Fudenberg and Harris (1992), noise is detrimental because for a given expected growth rate, noise almost surely lowers the long-term average geometric growth rate of the total number of individuals playing α by the quantity $\frac{1}{2}\sigma_\alpha^2$. In a geometric growth process, the quantities that matter (the proper fitness measures) are these long-term geometric growth rates, so the relevant payoffs are those of this modified game.¹⁴ In our model, noise is not detrimental, but if it is strong enough compared to the deterministic drift, then, with positive probability, it may lead to other outcomes than the deterministic model. Instead, the assumptions of Theorems 3.1

¹⁴In a discrete time setting, if $Z(n+1) = g(n)Z_n$ and $g(n) = k_i$ with probability p_i , what we mean is that the quantity that a.s. governs the long-term growth of Z is not $E(g) = \sum_i p_i k_i$, but $\exp(E(\ln g)) = \prod_i k_i^{p_i}$.

and 3.3 should be interpreted as guaranteeing that the deterministic drift prevails. One way to see this is to note that if β strictly dominates α in the original game and both strategies are affected by the same noise intensity ($\sigma_\alpha^2 = \sigma_\beta^2 = \sigma^2$), then β need not dominate α in the modified game defined by (3.1), unless the payoff margin in the original game is always greater than σ^2 .

Remark 3.8. It is also worth contrasting Theorem 3.3 to the unconditional convergence and stability results of Mertikopoulos and Moustakas (2010) for the stochastic replicator dynamics of exponential learning (2.16). As in the case of dominated strategies, the reason for this qualitative difference is the distinct origins of the perturbation process: the Itô correction in (2.16) is “just right” with respect to the dual variables $Y_\alpha = \log X_\alpha$, so a state $x^* \in \mathcal{X}$ is stochastically asymptotically stable under the (2.16) if and only if it is a strict equilibrium of the original game \mathfrak{G} .

4. THE EFFECT OF AGGREGATING PAYOFFS

In this section, we examine the case where players are less “myopic” and, instead of using revision protocols driven by their instantaneous payoffs, they base their decisions on the cumulative payoffs of their strategies over time. Formally, focusing for concreteness on the “imitation of success” revision protocol (2.7), this amounts to considering conditional switch rates of the form:

$$\tilde{\rho}_{\alpha\beta} = x_\beta U_\beta, \quad (4.1)$$

where

$$U_\beta(t) = \int_0^t v_\beta(x(s)) ds \quad (4.2)$$

denotes the cumulative payoff of strategy β up to time t . In this case, (RD) becomes:

$$\dot{x}_\alpha = x_\alpha \left[U_\alpha - \sum_\beta x_\beta U_\beta \right], \quad (4.3)$$

and, as was shown by Laraki and Mertikopoulos (2013), the evolution of mixed strategy shares is governed by the (deterministic) *second order replicator dynamics*:

$$\ddot{x}_\alpha = x_\alpha \left[v_\alpha(x) - \sum_\beta x_\beta u_\beta(x) \right] + x_\alpha \left[\dot{x}_\alpha^2/x_\alpha^2 - \sum_\beta \dot{x}_\beta^2/x_\beta \right]. \quad (\text{RD}_2)$$

As in the previous section, we are interested in the effects of random payoff shocks on the dynamics (RD₂). If the game’s payoff functions are subject to random shocks at each instant in time, then these shocks will also be aggregated over time, leading to the perturbed cumulative payoff process:

$$\hat{U}_\alpha(t) = \int_0^t v_\alpha(X(s)) ds + \int_0^t \sigma_\alpha(X(s)) dW_\alpha(s). \quad (4.4)$$

Since \hat{U}_α is continuous (a.s.), we obtain the stochastic integro-differential dynamics:

$$\begin{aligned} \dot{X}_\alpha &= X_\alpha \left[U_\alpha(t) - \sum_\beta X_\beta(t) U_\beta(t) \right] \\ &+ X_\alpha \left[\int_0^t \sigma_\alpha(X(s)) dW_\alpha(s) - \sum_\beta \int_0^t X_\beta(s) \sigma_\beta(X(s)) dW_\beta(s) \right], \end{aligned} \quad (4.5)$$

where, as in (SRD), we assume that the Brownian disturbances $W_\alpha(t)$ are independent.

To obtain an autonomous SDE from (4.5), let $V_\alpha = \dot{X}_\alpha$ denote the growth rate of strategy α . Then, differentiating (4.5) yields:

$$dV_\alpha = X_\alpha \left[\dot{U}_\alpha - \sum_\beta X_\beta \dot{U}_\beta \right] dt \quad (4.6a)$$

$$+ V_\alpha \left[U_\alpha - \sum_\beta X_\beta U_\beta \right] dt - X_\alpha \sum_\beta U_\beta V_\beta dt \quad (4.6b)$$

$$+ V_\alpha \left[\int_0^t \sigma_\alpha(X(s)) dW_\alpha(s) - \sum_\beta \int_0^t X_\beta(s) \sigma_\beta(X(s)) dW_\beta(s) \right] dt \quad (4.6c)$$

$$+ X_\alpha \left[\sigma_\alpha(X) dW_\alpha - \sum_\beta \sigma_\beta(X) X_\beta dW_\beta \right]. \quad (4.6d)$$

By (4.5), the sum of the first term of (4.6b) and (4.6c) is equal to V_α^2/X_α . Thus, using (4.2) we obtain:

$$\begin{aligned} dV_\alpha &= X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] dt + \frac{V_\alpha^2}{X_\alpha} dt - X_\alpha \sum_\beta U_\beta V_\beta dt \\ &\quad + X_\alpha \left[\sigma_\alpha(X) dW_\alpha - \sum_\beta \sigma_\beta(X) X_\beta dW_\beta \right], \end{aligned} \quad (4.7)$$

and, after summing over all α and solving for $X_\alpha \sum_\beta U_\beta V_\beta dt$, we get the second order SDE system:¹⁵

$$\begin{aligned} dX_\alpha &= V_\alpha dt \\ dV_\alpha &= X_\alpha \left[v_\alpha(X) - \sum_\beta x_\beta v_\beta(X) \right] dt + X_\alpha \left[V_\alpha^2/X_\alpha^2 - \sum_\beta V_\beta^2/X_\beta \right] dt \\ &\quad + X_\alpha \left[\sigma_\alpha(X) dW_\alpha - \sum_\beta \sigma_\beta(X) X_\beta dW_\beta \right]. \end{aligned} \quad (4.8)$$

By comparing the second order system (4.8) to (RD₂), we see that there is no Itô correction, just as in the first order case.¹⁶ By using similar arguments as in Laraki and Mertikopoulos (2013), it is then possible to show that the system (4.8) is well-posed, i.e. it admits a unique (strong) solution $X(t)$ for every interior initial condition $X(0) \in \mathcal{X}^\circ$, $V(0) \in \mathbb{R}^A$ and this solution remains in \mathcal{X}° for all time.

With this well-posedness result at hand, we begin by showing that (4.8) eliminates strategies that are dominated in the original game \mathfrak{G} (instead of the modified game \mathfrak{G}^σ):

Theorem 4.1. *Let $X(t)$ be a solution orbit of the dynamics (4.5) and assume that $\alpha \in \mathcal{A}$ is dominated by $\beta \in \mathcal{A}$. Then, α becomes extinct (a.s.).*

Proof. As in the proof of Theorem 3.1, let $Y_\alpha = \log X_\alpha$. Then, following the same string of calculations leading to (3.6), we get:¹⁷

$$dY_\alpha - dY_\beta = [U_\alpha - U_\beta] dt \quad (4.9a)$$

$$+ \left[\int_0^t \sigma_\alpha(X) dW_\alpha - \int_0^t \sigma_\beta(X) dW_\beta \right] dt \quad (4.9b)$$

¹⁵Recall that $\sum_\alpha dV_\alpha = 0$ since $\sum_\alpha X_\alpha = 1$.

¹⁶The reason however is different: in (SRD), there is no Itô correction because the noise is added directly to the dynamical system under study; in (4.8), there is no Itô correction because the noise is integrated over, so X_α is smooth (and, hence, obeys the rules of ordinary calculus).

¹⁷Recall that $\int_0^t \sigma_\alpha(X(s)) dW_\alpha(s)$ is continuous, so there is no Itô correction.

Since α is dominated by β , there exists some positive $m > 0$ such that $v_\alpha - v_\beta \leq -m$, and hence $U_\alpha(t) - U_\beta(t) \leq -mt$. Furthermore, with σ bounded and continuous on \mathcal{X} , Lemma A.1 readily yields:

$$-mt + \left[\int_0^t \sigma_\alpha(X) dW_\alpha - \int_0^t \sigma_\beta(X) dW_\beta \right] \sim -mt \quad (4.10)$$

as $t \rightarrow \infty$. Accordingly, (4.9) becomes:

$$dY_\alpha - dY_\beta \leq -mtdt + \theta(t)dt \quad (4.11)$$

where the remainder function $\theta(t)$ corresponding to the drift term (4.9b) is sublinear in t . By integrating and applying Lemma A.1 a second time, we then obtain:

$$Y_\alpha(t) - Y_\beta(t) \leq Y_\alpha(0) - Y_\beta(0) - \frac{1}{2}mt^2 + \int_0^t \theta(s) ds \sim -\frac{1}{2}mt^2 \quad (\text{a.s.}) \quad (4.12)$$

We infer that $\lim_{t \rightarrow \infty} Y_\alpha(t) = 0$ (a.s.), i.e. α becomes extinct along $X(t)$. \square

Remark 4.1. In view of Theorem 4.1, we see that the ‘‘imitation of long-term success’’ protocol (4.1) provides more robust elimination results than (2.7) in the presence of payoff shocks: contrary to Theorem 3.1, there are no ‘‘small noise’’ requirements in Theorem 4.1.¹⁸

Our next result provides the analogue of Theorem 3.3 regarding the stability of equilibrium play:

Theorem 4.2. *Let $X(t)$ be an interior solution orbit of the stochastic dynamics (4.5) 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 \mathfrak{G} .*

Moreover, for every neighborhood U_0 of x^ and for all $\varepsilon > 0$, we have:*

(2) *If $\mathbb{P}(X(t) \in U_0 \text{ for all } t \geq 0) \geq 1 - \varepsilon$ whenever $X(0) \in U$ for some neighborhood $U \subseteq U_0$ of x^* , then x^* is a Nash equilibrium of \mathfrak{G} .*

(3) *If x^* is a strict Nash equilibrium of \mathfrak{G} , there exists a neighborhood U of x^* such that:*

$$\mathbb{P}\left(X(t) \in U_0 \text{ for all } t \geq 0 \text{ and } \lim_{t \rightarrow \infty} X(t) = x^*\right) \geq 1 - \varepsilon, \quad (4.13)$$

whenever $X(0) \in U$.

Remark 4.2. Part 1 of Theorem 4.2 is in direct analogy with Part 1 of Theorem 3.3: the difference is that Theorem 4.2 shows that only Nash equilibria of the original game \mathfrak{G} can be ω -limits of interior orbits with positive probability. Put differently, if x^* is a strict equilibrium of \mathfrak{G}^σ but not of \mathfrak{G} ,¹⁹ there is zero probability that (4.5) converges to x^* .

On the other hand, Parts 2 and 3 are not tantamount to stochastic stability (Lyapunov or asymptotic) under the autonomous SDE system (4.8). The difference here is that (4.8) is only well-defined in the interior of \mathcal{X}° , so it is not straightforward how to define the notion of (stochastic) stability for boundary points $x^* \in \text{bd}(\mathcal{X})$; moreover, given that (4.8) is a second order system, stability should be stated in terms of the problem’s entire phase space, including initial velocities (for a relevant

¹⁸Theorem 4.1 actually applies to mixed dominated strategies as well (even iteratively dominated ones). The proof is a simple adaptation of the pure strategies case, so we omit it.

¹⁹Recall here that strict equilibria of \mathfrak{G} are also strict equilibria of \mathfrak{G}^σ , but the converse need not hold.

discussion, see Laraki and Mertikopoulos, 2013). Instead, the stated stability conditions simply reflect the fact that the integro-differential dynamics (4.5) always start with initial velocity $V(0) = 0$, so this added complexity is not relevant.

Proof of Theorem 4.2. We shadow the proof of Theorem 3.3.

Part 1. Assume that $\mathbb{P}(\lim_{t \rightarrow \infty} X(t) = x^*) > 0$ for some $x^* \in \mathcal{X}$. If x^* is not a Nash equilibrium of \mathfrak{G} , we will have $v_\alpha(x^*) < v_\beta(x^*)$ for some $\alpha \in \text{supp}(x^*)$, $\beta \in \mathcal{A}$. Accordingly, let U be a sufficiently small neighborhood of x^* in \mathcal{X} such that $v_\beta(x) - v_\alpha(x) \geq m$ for some $m > 0$ and for all $x \in U$. Since $X(t) \rightarrow x^*$ with positive probability, it also follows that $\mathbb{P}(X(t) \in U \text{ for all } t \geq 0) > 0$; hence, arguing as in (4.12) and conditioning on the positive probability event “ $X(t) \in U$ for all $t \geq 0$ ”, we get:

$$Y_\alpha(t) - Y_\beta(t) \sim -\frac{1}{2}mt^2 \quad (\text{conditionally a.s.}) \quad (4.14)$$

This implies $X_\alpha(t) \rightarrow 0$, contradicting our original assumption that $X(t)$ stays in a small neighborhood of x^* . We infer that x^* is a Nash equilibrium of \mathfrak{G} , as claimed.

Part 2. Simply note that the stability assumption of Part 2 implies that there exists a positive measure of interior trajectories $X(t)$ that remain in an arbitrarily small neighborhood of x^* with positive probability. The proof then follows as in Part 1.

Part 3. Let $Z_\alpha = Y_\alpha - Y_{\alpha^*}$ be defined as in (3.20) and let $m > 0$ be such that $v_{\alpha^*}(x) - v_\alpha(x) \geq m$ for all x in some sufficiently small neighborhood of x^* . Also, let $M > 0$ be sufficiently large so that $X(t) \in U$ if $Z_\alpha(t) \leq -M$ for all $\alpha \in \mathcal{A}^*$; as in the proof of Theorem 3.3, we will show that there is a suitable choice of M such that $Z_\alpha(0) < -2M$ for all $\alpha \neq \alpha^*$ implies that $X(t) \in U$ for all $t \geq 0$ and $Z_\alpha(t) \rightarrow -\infty$ for all $\alpha \neq \alpha^*$ with probability at least $1 - \varepsilon$.

Indeed, by setting $\beta = \alpha^*$ in (4.9), we obtain:

$$dZ_\alpha = [U_\alpha - U_{\alpha^*}] dt + \left[\int_0^t \sigma_\alpha(X) dW_\alpha - \int_0^t \sigma_{\alpha^*}(X) dW_{\alpha^*} \right] dt \quad (4.15)$$

so, recalling (eq:pay-cum-noise), for all $t \leq \tau_U = \inf\{t > 0 : X(t) \notin U\}$, we will have:

$$Z_\alpha(t) \leq -2M - \frac{1}{2}mt^2 + \int_0^t \theta(s) ds - \xi(t), \quad (4.16)$$

where ξ denotes the martingale part of (4.15) and $\theta(t)$ is defined as in (4.11).

Now, let $W(t)$ be a Wiener process starting at the origin. We will show that if M is chosen sufficiently large, then

$$\mathbb{P}\left(M + \frac{1}{2}mt^2 \geq \int_0^t W(s) ds \text{ for all } t \geq 0\right) \geq 1 - \varepsilon. \quad (4.17)$$

With a fair degree of hindsight, we note first that $\frac{1}{4}mt^2 + \frac{1}{2}M \geq at^2 + bt + c$ where $a = \frac{1}{4}m$, $b = \frac{1}{2}\sqrt{Mm}$ and $c = \frac{1}{2}M$, so it suffices to show that the hitting time $\tau = \inf\{t : \int_0^t W(s) ds = at^2 + bt + c\}$ is infinite with probability at least $1 - \varepsilon$. However, by the mean value theorem, there exists some (random) time τ_1 such that:

$$2a\tau_1 + b - W(\tau_1) = \frac{0 - c}{\tau} \leq 0. \quad (4.18)$$

Since $c/\tau > 0$, the hitting time $\tau' = \inf\{t > 0 : W(t) = 2at + b\}$ will satisfy $\tau'(\omega) < \tau(\omega)$ for every trajectory ω of W with $\tau(\omega) < \infty$. However, Lemma A.2 gives $\mathbb{P}[\tau' < \infty] = \exp(-2ab)$, hence:

$$\mathbb{P}(\tau < \infty) \leq \mathbb{P}(\tau' < \infty) = \exp(-2ab) = \exp(-Mm/4), \quad (4.19)$$

i.e. $\mathbb{P}(\tau < \infty)$ can be made arbitrarily small by choosing M large enough. We thus deduce that

$$\int_0^t W(s) ds \leq \frac{1}{2}mt^2 + M \quad \text{for all } t \geq 0 \quad (4.20)$$

with probability no less than $1 - \varepsilon$.

Going back to (4.16), we see that $\int_0^t \theta(s) ds - \xi(t) - \frac{1}{2}mt^2$ remains below M for all time with probability at least $1 - \varepsilon$ (simply use the probability estimate (4.17) and argue as in the proof of Theorem 3.3 recalling that the processes W_α are assumed independent). In turn, this shows that $\mathbb{P}(X(t) \in U \text{ for all } t \geq 0) \geq 1 - \varepsilon$, so, conditioning on this last event and letting $t \rightarrow \infty$ in (4.16), we obtain:

$$\mathbb{P}\left(\lim_{t \rightarrow \infty} Z_\alpha(t) = -\infty \mid X(t) \in U \text{ for all } t \geq 0\right) = 1 \quad \text{for all } \alpha \neq \alpha^*. \quad (4.21)$$

We conclude that $X(t)$ remains in U_0 and $X(t) \rightarrow x^*$ with probability at least $1 - \varepsilon$, as was to be shown. \square

5. DISCUSSION

In this section, we discuss some points that would have otherwise disrupted the flow of the main text:

5.1. Payoff shocks in bimatrix games. Throughout our paper, we have worked with generic population games, so we have not made any specific assumptions on the payoff shocks either. On the other hand, if the game's payoff functions are obtained from some common underlying structure, then the resulting payoff shocks may also end up having a likewise specific form.

For instance, consider a basic (symmetric) random matching model where pairs of players are drawn randomly from a nonatomic population to play a symmetric two-player game with payoff matrix $V_{\alpha\beta}$, $\alpha, \beta = 1, \dots, n$. In this case, the payoff to an α -strategist in the population state $x \in \mathcal{X}$ will be of the form:

$$v_\alpha(x) = \sum_\beta V_{\alpha\beta} x_\beta. \quad (5.1)$$

Thus, if the entries of V are disturbed at each $t \geq 0$ by some (otherwise independent) white noise process $\xi_{\alpha\beta}$, the perturbed payoff matrix $\hat{V}_{\alpha\beta} = V_{\alpha\beta} + \xi_{\alpha\beta}$ will result in the total payoff shock:

$$\xi_\alpha = \sum_\beta \xi_{\alpha\beta} x_\beta. \quad (5.2)$$

The stochastic dynamics (2.8) thus become:

$$\begin{aligned} dX_\alpha &= X_\alpha \left[\sum_\beta V_{\alpha\beta} X_\beta - \sum_{\beta,\gamma} V_{\beta\gamma} X_\beta X_\gamma \right] dt \\ &+ X_\alpha \left[\sum_\beta \sigma_{\alpha\beta} X_\beta dW_{\alpha\beta} - \sum_{\beta,\gamma} \sigma_{\beta\gamma} X_\beta X_\gamma dW_{\beta\gamma} \right], \end{aligned} \quad (5.3)$$

where the Wiener processes $W_{\alpha\beta}$ are assumed independent.

To compare (5.3) with the core model (SRD), the same string of calculations as in the proof of Theorems 3.1 and 3.3 leads to the modified payoff functions:

$$v_\alpha^\sigma(x) = v_\alpha(x) - \frac{1}{2}(1 - 2x_\alpha) \sum_\beta \sigma_{\alpha\beta}^2 x_\beta^2. \quad (5.4)$$

It is then trivial to see that Theorems 3.1 and 3.3 still apply with respect to the modified game \mathfrak{G}^σ with payoff functions defined as above; however, seeing as (5.4)

is cubic in x_α and considering the case of constant noise, these modified payoff functions no longer correspond to random matching in a modified bimatrix game.

5.2. Stratonovich-type perturbations. Depending on the origins of the payoff shock process (for instance, if there are nontrivial autocorrelation effects that do not vanish in the continuous-time regime), the perturbed dynamics (SRD) could instead be written as a Stratonovich-type SDE (Kuo, 2006):

$$\partial X_\alpha = X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] dt + X_\alpha \left[\sigma_\alpha \partial W_\alpha - \sum_\beta X_\beta \sigma_\beta \partial W_\beta \right], \quad (5.5)$$

where $\partial(\cdot)$ denotes Stratonovich integration.²⁰ In this case, if $M_{\alpha\beta} = X_\alpha(\delta_{\alpha\beta} - X_\beta)\sigma_\beta$ denotes the diffusion matrix of (5.5), the Itô equivalent SDE corresponding to (5.5) will be:

$$\begin{aligned} dX_\alpha &= X_\alpha \left[v_\alpha(X) - \sum_\beta X_\beta v_\beta(X) \right] dt + \frac{1}{2} \sum_{\beta,\gamma} \frac{\partial M_{\alpha\beta}}{\partial X_\gamma} M_{\gamma\beta} dt \\ &\quad + X_\alpha \left[\sigma_\alpha dW_\alpha - \sum_\beta X_\beta \sigma_\beta dW_\beta \right]. \end{aligned} \quad (5.6)$$

Then, assuming that the shock coefficients σ_β are constant, some algebra yields the following explicit expression for the Itô correction of (5.6):

$$\begin{aligned} &\frac{1}{2} \sum_{\beta,\gamma} \frac{\partial M_{\alpha\beta}}{\partial X_\gamma} M_{\gamma\beta} dt \\ &= \frac{1}{2} \sum_{\beta,\gamma} (\delta_{\alpha\beta\gamma} - \delta_{\alpha\gamma}x_\beta - \delta_{\beta\gamma}x_\alpha) \sigma_\beta \cdot X_\gamma (\delta_{\gamma\beta} - X_\beta) \sigma_\beta dt \\ &= \frac{1}{2} \sum_{\beta,\gamma} [\delta_{\alpha\beta\gamma}(1 - 2X_\beta) + \delta_{\alpha\gamma}X_\beta^2 - \delta_{\beta\gamma}X_\alpha + \delta_{\beta\gamma}X_\alpha X_\beta] X_\gamma \sigma_\beta^2 dt \\ &= \frac{1}{2} X_\alpha \left[(1 - 2X_\beta)\sigma_\beta^2 - \sum_\beta X_\beta(1 - 2X_\beta)\sigma_\beta^2 \right] dt. \end{aligned} \quad (5.7)$$

By substituting this correction back to (5.6), we see that the replicator dynamics with Stratonovich shocks (5.5) are equivalent to the (Itô) stochastic replicator dynamics of exponential learning (2.16). In this context, Mertikopoulos and Moustakas (2010) showed that the conclusions of Theorems 3.1 and 3.3 apply directly to the original, unmodified game \mathfrak{G} under (2.16), so dominated strategies become extinct and strict equilibria are stochastically asymptotically stable under (5.5) as well. Alternatively, this can also be seen directly from the correction term (5.7) which cancels with that of (3.1).

5.3. Random strategy switches. An alternative source of noise to the players' evolution under (RD) could come from random masses of players that switch strategies without following an underlying deterministic drift – as opposed to jumps with a well-defined direction induced by a revision protocol. To model this kind of “mutations”, we posit that the relative mass $dX_{\alpha\beta}$ of players switching from α to β over an infinitesimal time interval dt is governed by the SDE:

$$dX_{\alpha\beta} = X_\alpha (\rho_{\alpha\beta} dt + dM_{\alpha\beta}), \quad (5.8)$$

where $dM_{\alpha\beta}$ denotes the (conditional) mutation rate from α to β .

²⁰For a general overview of the differences between Itô and Stratonovich integration, see van Kampen (1981); for a more specific account in the context of stochastic population growth models, the reader is instead referred to Khasminskii and Potsepun (2006) and Hofbauer and Imhof (2009).

To account for randomness, we will assume that $M_{\alpha\beta}$ has unbounded variation over finite time intervals (contrary to the bounded variation drift term $X_\alpha \rho_{\alpha\beta} dt$). Moreover, for concreteness, we will focus on the imitative regime where $\rho_{\alpha\beta} = x_\beta r_{\alpha\beta}$ and the mutation processes $M_{\alpha\beta}$ follow a similar imitative pattern, namely $dM_{\alpha\beta} = X_\beta dR_{\alpha\beta}$. The net change in the population of α -strategists will then be

$$\sum_{\beta} X_\beta dM_{\beta\alpha} - X_\alpha \sum_{\beta} dM_{\alpha\beta} = X_\alpha \sum_{\beta} X_\beta dQ_{\beta\alpha}, \quad (5.9)$$

where $dQ_{\beta\alpha} = dR_{\beta\alpha} - dR_{\alpha\beta}$ describes the *net influx* of β -strategists in strategy α per unit population mass. Thus, assuming that the increments $dQ_{\beta\alpha}$ are zero-mean, we will model Q as an Itô process of the form:

$$dQ_{\alpha\beta} = \eta_{\alpha\beta}(X) dW_{\alpha\beta} \quad (5.10)$$

where $W_{\alpha\beta}$ is an ordinary Wiener process and the diffusion coefficients $\eta_{\alpha\beta}: \mathcal{X} \rightarrow \mathbb{R}$ reflect the intensity of the mutation process. In particular, the only assumptions that we need to make for W and η are that:

$$dW_{\alpha\beta} = -dW_{\beta\alpha} \quad \text{and} \quad \eta_{\alpha\beta} = \eta_{\beta\alpha} \quad \text{for all } \alpha, \beta \in \mathcal{A} \text{ and for all } k \in \mathcal{N}, \quad (5.11)$$

so that the net influx from α to β is minus the net influx from β to α ; except for this ‘‘conservation of mass’’ requirement, we will assume that the processes $dW_{\alpha\beta}$ are otherwise independent. Thus, in the special case of the ‘‘imitation of success’’ revision protocol (2.7), we obtain the *replicator dynamics with random mutations*:

$$dX_\alpha = X_\alpha \left[v_\alpha(X) - \sum_{\beta} X_\beta v_\beta(X) \right] dt + X_\alpha \sum_{\beta \neq \alpha} X_\beta \eta_{\beta\alpha}(X) dW_{\beta\alpha}. \quad (5.12)$$

This equation differs from (SRD) in that the martingale term of (SRD) cannot be recovered from that of (5.12) without violating the symmetry conditions (5.11) that guarantee that there is no net transfer of mass across any pair of strategies $\alpha, \beta \in \mathcal{A}$. Nonetheless, by repeating the same analysis as in the case of Theorems 3.1 and 3.3, we obtain the following proposition for the stochastic dynamics (5.12):

Proposition 5.1. *Let $X(t)$ be an interior solution orbit of the stochastic dynamics (5.12) and consider the noise-adjusted game \mathfrak{G}^η with modified payoff functions:*

$$v_\alpha^\eta(x) = v_\alpha(x) - \frac{1}{2} \sum_{\beta \neq \alpha} x_\beta^2 \eta_{\beta\alpha}^2(x). \quad (5.13)$$

We then have:

- (1) *If $p \in \mathcal{X}$ is dominated in \mathfrak{G}^η , then it becomes extinct under (5.12).*
- (2) *If $\mathbb{P}(\lim_{t \rightarrow \infty} X(t) = x^*) > 0$ for some $x^* \in \mathcal{X}$, then x^* is a Nash equilibrium of \mathfrak{G}^η .*
- (3) *If $x^* \in \mathcal{X}$ is stochastically Lyapunov stable, then it is a Nash equilibrium of \mathfrak{G}^η .*
- (4) *If $x^* \in \mathcal{X}$ is a strict Nash equilibrium of \mathfrak{G}^η , then it is stochastically asymptotically stable under (5.12).*

Proof. The proof is similar to that of Theorems 3.1 (for Part 1) and 3.3 (for Parts 2–4), so we omit it. \square

APPENDIX A. AUXILIARY RESULTS FROM STOCHASTIC ANALYSIS

In this appendix, we state and prove two auxiliary results from stochastic analysis that were used throughout the paper. Lemma A.1 is an asymptotic growth bound for Wiener processes relying on the law of the iterated logarithm, while Lemma A.2 is a calculation of the probability that a Wiener process starting at the origin hits the line $a + bt$ in finite time. Both lemmas appear in a similar context in Bravo and Mertikopoulos (2014); we provide a proof here only for completeness and ease of reference.

Lemma A.1. *Let $W(t) = (W_1(t), \dots, W_n(t))$, $t \geq 0$, be an n -dimensional Wiener processes and let $Z(t)$ be a bounded, continuous process in \mathbb{R}^n . Then:*

$$f(t) + \int_0^t Z(s) \cdot dW(s) \sim f(t) \quad \text{as } t \rightarrow \infty \text{ (a.s.)}, \quad (\text{A.1})$$

for any function $f: [0, \infty) \rightarrow \mathbb{R}$ such that $\lim_{t \rightarrow \infty} (t \log \log t)^{-1/2} f(t) = +\infty$.

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

$$d[\xi, \xi] = d\xi \cdot d\xi = \sum_{i=1}^n Z_i Z_j \delta_{ij} dt \leq M dt, \quad (\text{A.2})$$

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

$$\frac{f(t) + \xi(t)}{f(t)} = 1 + \frac{\widetilde{W}(\rho(t))}{f(t)}. \quad (\text{A.3})$$

Obviously, if $\lim_{t \rightarrow \infty} \rho(t) \equiv \rho(\infty) < +\infty$, $\widetilde{W}(\rho(\infty))$ is normally distributed so $\widetilde{W}(\rho(t))/f(t) \rightarrow 0$ and there is nothing to show. Otherwise, if $\lim_{t \rightarrow \infty} \rho(t) = +\infty$, the quadratic variation bound (A.2) and the law of the iterated logarithm yield:

$$\frac{|\widetilde{W}(\rho(t))|}{f(t)} \leq \frac{|\widetilde{W}(\rho(t))|}{\sqrt{2\rho(t) \log \log \rho(t)}} \times \frac{\sqrt{2Mt \log \log Mt}}{f(t)} \rightarrow 0 \quad \text{as } t \rightarrow \infty, \quad (\text{A.4})$$

and our claim follows. \square

Lemma A.2. *Let $W(t)$ be a standard one-dimensional Wiener process and consider the hitting time $\tau_{a,b} = \inf\{t > 0 : W(t) = a + bt\}$, $a, b \in \mathbb{R}$. Then:*

$$\mathbb{P}(\tau_{a,b} < \infty) = \exp(-ab - |ab|). \quad (\text{A.5})$$

Proof. Let $\overline{W}(t) = W(t) - bt$ so that $\tau_{a,b} = \inf\{t > 0 : \overline{W}(t) = a\}$. By Girsanov's theorem (see e.g. Øksendal, 2007, Chap. 8), there exists a probability measure \mathbb{Q} such that a) \overline{W} is a Brownian motion with respect to \mathbb{Q} ; and b) the Radon–Nikodym derivative of \mathbb{Q} with respect to \mathbb{P} satisfies

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = \exp(-b^2 t/2 + bW(t)) = \exp(b^2 t/2 - b\overline{W}(t)), \quad (\text{A.6})$$

where \mathcal{F}_t denotes the natural filtration of $W(t)$. We then get

$$\begin{aligned} \mathbb{P}(\tau_{a,b} < t) &= \mathbb{E}_{\mathbb{P}}[\mathbf{1}(\tau_{a,b} < t)] \\ &= \mathbb{E}_{\mathbb{Q}}[\mathbf{1}(\tau_{a,b} < t) \cdot \exp(-b^2 t/2 - b\overline{W}(t))] \\ &= \mathbb{E}_{\mathbb{Q}}[\mathbf{1}(\tau_{a,b} < t) \cdot \exp(-b^2 \tau_{a,b}/2 - b\overline{W}(\tau_{a,b}))] \\ &= \exp(-ab) \mathbb{E}_{\mathbb{Q}}[\mathbf{1}(\tau_{a,b} < t) \cdot \exp(-b^2 \tau_{a,b}/2)], \end{aligned} \quad (\text{A.7})$$

and hence:

$$\begin{aligned} \mathbb{P}(\tau_{a,b} < \infty) &= \lim_{t \rightarrow \infty} \mathbb{P}(\tau_{a,b} < t) \\ &= \lim_{t \rightarrow \infty} \exp(-ab) \mathbb{E}_{\mathbb{Q}}[\mathbf{1}(\tau_{a,b} < t) \cdot \exp(-b^2 \tau_{a,b}/2)] \\ &= \exp(-ab) \mathbb{E}_{\mathbb{Q}}[\exp(-b^2 \tau_{a,b}/2)] \\ &= \exp(-ab - |ab|), \end{aligned} \quad (\text{A.8})$$

where, in the last step, we used the expression $\mathbb{E}[\exp(-\lambda \tau_a)] = \exp(-a\sqrt{2\lambda})$ for the Laplace transform of the Brownian hitting time $\tau_a = \inf\{t > 0 : W(t) = a\}$ (Karatzas and Shreve, 1998). \square

REFERENCES

- Akin, E., 1980: Domination or equilibrium. *Mathematical Biosciences*, **50** (3-4), 239–250.
- Bergstrom, T. C., 2014: On the evolution of hoarding, risk-taking, and wealth distribution in nonhuman and human populations. *Proceedings of the National Academy of Sciences of the USA*, **111** (3), 10860–10867.
- Bertsekas, D. P. and R. Gallager, 1992: *Data Networks*. 2d ed., Prentice Hall, Englewood Cliffs, NJ.
- Björnerstedt, J. and J. W. Weibull, 1996: Nash equilibrium and evolution by imitation. *The Rational Foundations of Economic Behavior*, K. J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt, Eds., St. Martin's Press, New York, NY, 155–181.
- Bravo, M. and P. Mertikopoulos, 2014: On the robustness of learning in games with stochastically perturbed payoff observations. <http://arxiv.org/abs/1412.6565>.
- Cabrales, A., 2000: Stochastic replicator dynamics. *International Economic Review*, **41** (2), 451–81.
- Fudenberg, D. and C. Harris, 1992: Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, **57** (2), 420–441.
- Hofbauer, J. and L. A. Imhof, 2009: Time averages, recurrence and transience in the stochastic replicator dynamics. *The Annals of Applied Probability*, **19** (4), 1347–1368.
- Hofbauer, J. and K. Sigmund, 2003: Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, **40** (4), 479–519.
- Hofbauer, J., S. Sorin, and Y. Viossat, 2009: Time average replicator and best reply dynamics. *Mathematics of Operations Research*, **34** (2), 263–269.
- Imhof, L. A., 2005: The long-run behavior of the stochastic replicator dynamics. *The Annals of Applied Probability*, **15** (1B), 1019–1045.
- Karatzas, I. and S. E. Shreve, 1998: *Brownian Motion and Stochastic Calculus*. Springer-Verlag, Berlin.
- Khasminskii, R. Z., 2012: *Stochastic Stability of Differential Equations*. 2d ed., No. 66 in Stochastic Modelling and Applied Probability, Springer-Verlag, Berlin.
- Khasminskii, R. Z. and N. Potsepun, 2006: On the replicator dynamics behavior under Stratonovich type random perturbations. *Stochastic Dynamics*, **6**, 197–211.
- Kuo, H.-H., 2006: *Introduction to Stochastic Integration*. Springer, Berlin.
- Laraki, R. and P. Mertikopoulos, 2013: Higher order game dynamics. *Journal of Economic Theory*, **148** (6), 2666–2695.

- Littlestone, N. and M. K. Warmuth, 1994: The weighted majority algorithm. *Information and Computation*, **108** (2), 212–261.
- Mertikopoulos, P. and A. L. Moustakas, 2009: Learning in the presence of noise. *GameNets '09: Proceedings of the 1st International Conference on Game Theory for Networks*.
- Mertikopoulos, P. and A. L. Moustakas, 2010: The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, **20** (4), 1359–1388.
- Nachbar, J. H., 1990: Evolutionary selection dynamics in games. *International Journal of Game Theory*, **19**, 59–89.
- Øksendal, B., 2007: *Stochastic Differential Equations*. 6th ed., Springer-Verlag, Berlin.
- Robson, A. J. and L. Samuelson, 2011: The evolutionary foundations of preferences. *Handbook of Social Economics*, J. Benhabib, A. Bisin, and M. O. Jackson, Eds., North-Holland, Vol. 1, chap. 7, 221–310.
- Rustichini, A., 1999: Optimal properties of stimulus-response learning models. *Games and Economic Behavior*, **29**, 230–244.
- Samuelson, L. and J. Zhang, 1992: Evolutionary stability in asymmetric games. *Journal of Economic Theory*, **57**, 363–391.
- Sandholm, W. H., 2010: *Population Games and Evolutionary Dynamics*. Economic learning and social evolution, MIT Press, Cambridge, MA.
- Schlag, K. H., 1998: Why imitate, and if so, how? a boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, **78** (1), 130–156.
- Sorin, S., 2009: Exponential weight algorithm in continuous time. *Mathematical Programming*, **116** (1), 513–528.
- Taylor, P. D. and L. B. Jonker, 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.
- 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. *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*, 371–383.
- Weibull, J. W., 1995: *Evolutionary Game Theory*. MIT Press, Cambridge, MA.

(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://mescal.imag.fr/membres/panayotis.mertikopoulos>

(Y. Viossat) PSL, UNIVERSITÉ PARIS–DAUPHINE, CEREMADE UMR7534, PLACE DU MARÉCHAL DE LATTRE DE TASSIGNY, 75775 PARIS, FRANCE
E-mail address: viossat@ceremade.dauphine.fr