

# BEATING PRICE OF ANARCHY AND GRADIENT DESCENT WITHOUT REGRET IN POTENTIAL GAMES

**Iosif Sakos**

Singapore University of Technology and Design  
iosif\_sakos@mymail.sutd.edu.sg

**Stefanos Leonardos**

King’s College London  
stefanos.leonardos@kcl.ac.uk

**Stelios Andrew Stavroulakis**

UC Irvine  
sstavrou@uci.edu

**William Overman**

Graduate School of Business, Stanford University  
wpo@stanford.edu

**Ioannis Panageas**

UC Irvine  
ipanagea@ics.uci.edu

**Georgios Piliouras**

Singapore University of Technology and Design  
georgios@sutd.edu.sg

## ABSTRACT

Arguably one of the thorniest problems in game theory is that of equilibrium selection. Specifically, in the presence of multiple equilibria do self-interested learning dynamics typically select the socially optimal ones? We study a rich class of continuous-time no-regret dynamics in potential games (PGs). Our class of dynamics, *Q-Replicator Dynamics* (QRD), include gradient descent (GD), log-barrier and replicator dynamics (RD) as special cases. We start by establishing *pointwise convergence* of all QRD to Nash equilibria in almost all PGs. In the case of GD, we show a tight average case performance within a factor of two of optimal, for a class of symmetric  $2 \times 2$  potential games with unbounded Price of Anarchy (PoA). Despite this positive result, we show that GD is not always the optimal choice even in this restricted setting. Specifically, GD outperforms RD, if and only if *risk-* and *payoff-dominance* equilibria coincide. Finally, we experimentally show how these insights extend to all QRD dynamics and that unbounded gaps between average case performance and PoA analysis are common even in larger settings.

## 1 INTRODUCTION

Multi-agent learning often occurs on highly non-convex landscapes even when agents have perfectly aligned or common interests (Bard et al., 2020; Dafoe et al., 2020; 2021). Thus, even if learning dynamics equilibrate, their fixed points may include saddle points or local optima of poor performance (Dauphin et al., 2014). A large stream of recent work has made considerable progress in showing convergence of optimization driven learning dynamics to locally attracting points (Ge et al., 2015; Lee et al., 2019; Mertikopoulos et al., 2019) or equilibria in cooperative games (Cohen et al., 2017; Palaiopoulos et al., 2017; Anagnostides et al., 2022; Leonardos & Piliouras, 2022).<sup>1</sup> However, non-trivial games routinely possess attracting points of vastly different performance, and this remains true, even if one restricts attention to refined and highly robust notions of equilibria, such as pure/strict *Nash equilibria* (NE) (Kleinberg et al., 2009; Flokas et al., 2020). Accordingly, these convergence results are not enough to reason about the *quality of the learning outcomes*.

This challenge, known as the *equilibrium selection problem*, is arguably one of the thorniest problems on the intersection of game theory and learning theory with multiple practical manifestation in ML systems (Dafoe et al., 2020). Standard game-theoretic approaches (Harsanyi, 1973; Harsanyi & Selten, 1988; van Damme, 1987) or *worst-case* metrics, such as the *Price of Anarchy* (PoA) Koutsoupias & Papadimitriou (1999); Christodoulou & Koutsoupias (2005); Roughgarden (2015), offer little

<sup>1</sup>For convergence to equilibria in competitive games see also (Daskalakis & Panageas, 2018; Bailey & Piliouras, 2019; Cai et al., 2022).

insight, often none at all, from a dynamic/learning perspective. The reason is that  $\mathcal{N}$  may be reachable only from very small sets of initial conditions. Thus, instead of seeking approximately optimal performance for (almost) all initial conditions, which is generally impractical, we need to develop and argue about average performance measures that couple the limit point of the learning process with the likelihood that such an outcome is reached by the given learning dynamics.

Based on the above, our focus is to study 1) under which conditions can we reliably compare the performance of standard learning dynamics, such as replicator and gradient descent, in games and 2) under which conditions can they be shown to be approximately optimal under random initialization even when the performance gap between different equilibria, i.e., the PoA, is unbounded.

**Model and Contributions.** To make progress in these directions, we study replicator dynamics (QRD) in general potential games (PGs). QRD are one of the most general classes of continuous time, multi-agent learning dynamics that include gradient descent (GD), replicator (RD) and log-barrier dynamics as special cases (Giannou et al., 2021). On the other hand, PGs are the standard model of multi-agent coordination and include congestion and identical interests games as important sub-classes (Wang & Sandholm, 2002; Panait & Luke, 2005; Carroll et al., 2019; Dafoe et al., 2020).

To analyze the average- and comparative-performance of QRD in PGs, an essential first step is to establish their convergence to equilibrium points. In our first result, we answer this question affirmatively by proving pointwise convergence of a QRD to NE in almost all finite PGs (Theorem 3.2). Our proof combines recent results from Swenson et al. (2020) standard convergence techniques in the study of PGs, e.g., Palaiopanos et al. (2017); Mertikopoulos & Sandholm (2016; 2018).

In Section 4, we turn to our main focus of equilibrium selection, we restrict our attention for analytical tractability to a class of 2-agent, 2-action (2) PGs. Despite its simple representation, this class retains all complexities of multi-agent learning that we aim to study. Specifically, as we show, there exist games in this class which have arbitrarily large PoA, i.e., for which the worst-case equilibrium can be arbitrarily worse than the socially optimal outcome, but for which Average Price of Anarchy (APoA) under GD admits a (tight) upper bound<sup>2</sup> (Theorem 4.8). Despite this positive news for GD, we also show that, even in this restricted class of games, GD is not always optimal. Specifically, GD reaches the payoff-dominant, i.e., socially optimal, equilibrium more often than RD, if and only if, this equilibrium is also risk-dominant (less risky) (Theorem 4.6). To the best of our knowledge, this provides the first comparison in the performance between two optimal no-regret dynamics and adds a new dimension to the long literature on the important interplay between risk- and payoff-dominance in game-theoretic dynamics (Kim, 1996; Schmidt et al., 2003; Sato et al., 2005; Kaisers & Tuyls, 2011; Kianercy & Galstyan, 2012; Leonardos & Piliouras, 2022; Pangallo et al., 2022).

Finally, in section 5, we present experiments in larger PGs for which risk- and payoff-dominance can be properly generalized (diagonal payoff matrices). Such games are, in fact, designed to be as hard as possible from the perspective of average case performance due to their exponentially large number of NE and unbounded PoA. However, as our experimental results rather surprisingly indicate, increasing the size of the game seems to have little to no effect on average-case or relative performance and the analogues of Theorem 4.6, i.e., comparative performance of RD and GD, and Theorem 4.8, i.e., bounded APoA (in fact typically around 1.2), continue to hold.

**Related Work** The literature on average-case performance is scarce. As exhibited by Panageas & Piliouras (2016); Zhang & Hofbauer (2015); Pangallo et al. (2022) one of the main reasons is that regions of attraction are complex geometric manifolds that become mathematically intractable even in low-dimensional settings. Average performance metrics are first introduced in Panageas & Piliouras (2016) who is the main precursor to our work. Panageas & Piliouras (2016) focus exclusively on replicator dynamics and prove pointwise convergence only for the special cases of congestion games with linear cost functions and network coordination games. Their bounds only apply in restricted instances of common interest games (Stag Hunt). Critically, prior to our work it was not known whether a rigorous comparison between RD and GD is possible in any single game setting.

<sup>2</sup>Swenson et al. (2020) shows that all NEs, almost all, i.e., in all but a closed set with Lebesgue measure zero, PGs are regular in the sense of Harsanyi (1973), i.e., they are isolated and highly robust.

Regarding the technical tools, our approach makes use of two connected, yet fundamentally different, theories. The first part, i.e. convergence, relies on the theory of Lyapunov analysis and the properties of dissipative systems, i.e., systems that lose momentum over time till they converge to a steady state (Haddad & Chellaboina, 2008). This in itself is a standard staple of analyzing learning in potential games (Kleinberg et al., 2009; Cohen et al., 2017; Anagnostides et al., 2022). By contrast, the second part, i.e. performance, relies on the existence of invariant functions a feature most often studied in conservative systems, that characterize stable and unstable areas in the state space of such systems (Palaiofanos et al., 2017; Nagarajan et al., 2020). It is important to note that this proof technique is scalable to higher dimensional settings. The key idea is to identify one or more (in more complex games) invariants of motion for QRD or other game dynamics. This is arguably a challenging task, but several recent works showcase that this is a viable proof strategy in rather general settings Balduzzi et al. (2018); Mertikopoulos et al. (2018); Piliouras & Wang (2021); Paik & Grifon (2023).

## 2 PRELIMINARIES: GAME-THEORETIC AND BEHAVIORAL MODELS

**Game-theoretic model.** A multi-agent finite potential game  $\Gamma := fN; (A_k; u_k)_{k \in 2N}; g$  denotes the interaction between a set of  $n := f1; \dots; ng$  agents. Each agent  $k \in 2N$  has a finite set of actions,  $A_k$ , with size  $|A_k|$ , and a reward function  $u_k : A \rightarrow \mathbb{R}$  where  $A := \prod_{k \in 2N} A_k$  is the set of all pure action profiles. Agents may use mixed actions or choice distributions  $x_k = (x_{k;a_k})_{a_k \in A_k} \in \Delta_{A_k}$ , where  $x_{k;a_k}$  is the probability with which agent  $k$  uses their action  $a_k \in A_k$  and  $\Delta_{A_k} := \{x_k \in \mathbb{R}^{|A_k|} \mid \sum_{a_k \in A_k} x_{k;a_k} = 1; x_{k;a_k} \geq 0\}$  is the  $(|A_k| - 1)$ -dimensional simplex. Given any mixed-action  $x_k \in \Delta_{A_k}$ , we write  $\text{supp}(x_k) := \{a_k \in A_k \mid x_{k;a_k} > 0\}$  to denote the support of action  $x_k$ , i.e., the set of all pure actions  $a_k \in A_k$  that are selected with positive probability at  $x_k$ . Using conventional notation, we also write  $s = (s_k; s_{-k}) \in A$  and  $x = (x_k; x_{-k}) \in X := \prod_{k \in 2N} \Delta_{A_k}$  to denote the joint pure and mixed action profile, where  $s_k$  and  $x_{-k}$  are the vectors of pure and mixed actions, respectively, of all agents other than  $k$ . When time is relevant, we will use the index  $t$  for all the above, e.g., we will write  $x_k(t)$  for agent  $k$ 's choice distribution at time  $t \geq 0$ . Finally, a function  $u : A \rightarrow \mathbb{R}$  is called a potential function if it satisfies  $u_k(s) - u_k(s'_k; s_{-k}) = \phi(s) - \phi(s'_k; s_{-k})$ , for all  $k \in 2N$  and all  $s; s'_k \in A$ . The agents' reward functions and the potential function extend naturally to mixed action profiles with  $u_k(x) = E_{s \sim x}[u_k(s)]$  and  $\phi(x) = E_{s \sim x}[\phi(s)]$ .

**Regular Nash and restricted equilibria.** A Nash equilibrium (NE) of  $\Gamma$  is an action profile  $x \in X$  such that  $u_k(x) \geq u_k(x_k; x_{-k})$ , for all  $k \in 2N$  and for all  $x \in X$ . By linearity of expectation, the above definition is equivalent to  $u_k(x) \geq u_k(a_k; x_{-k})$ , for all  $a_k \in A_k$ , and all  $k \in 2N$ , where  $u_k(a_k; x_{-k})$  is the reward of agent  $k$  for playing pure action  $a_k$  against mixed strategies  $x_{-k}$  by the remaining agents.  $\text{NE}(\Gamma)$  denote the set of all NE of  $\Gamma$ . A NE is called symmetric if  $x_1 = \dots = x_n$ , and is called fully mixed if  $\text{supp}(x) = \prod_{k \in 2N} \text{supp}(x_k) = A$ . A NE is called regular if it satisfies the following definition.

**Definition 2.1 (Regular Nash equilibria (Harsanyi, 1973; Swenson et al., 2020))** A Nash equilibrium,  $x \in \text{NE}(\Gamma)$ , is called regular if it is (i) quasi-strict, i.e., if for each player  $k \in 2N$ ,  $x_k$  assigns positive probability to all best responses of player  $k$  against  $x_{-k}$ , and (ii) second-order non-degenerate, i.e., if the Hessian  $H(x)$  of the potential function  $\phi$ , taken with respect to  $\text{supp}(x)$  is non-singular.

Finally, a restriction,  $\Gamma^0 := fN; (A_k^0; u_k^0)_{k \in 2N}; g$  of  $\Gamma$ , is a game where  $A_k^0 \subseteq A_k$  and  $u_k^0 : A^0 \rightarrow \mathbb{R}$  is the restriction of  $u_k$  to  $A^0 := \prod_{k \in 2N} A_k^0$  for all  $k \in 2N$ . We write  $\mathcal{R}$  to denote the set of all restrictions of  $\Gamma$ . An action-profile  $x \in X$  is called a restricted equilibrium of  $\Gamma$  if it is a Nash equilibrium of a restriction of  $\Gamma$ , cf. Mertikopoulos & Sandholm (2018). It is easy to see that all restrictions of a potential game  $\Gamma := fN; (A_k; u_k)_{k \in 2N}; g$  are also potential games, whose potential functions are restrictions of  $\phi$  to the respective subspace  $A^0$ .

**Behavioral-learning model.** The evolution of agents' choice distributions in the joint action space,  $X$ , is governed by the replicator dynamics (QRD) which are described by the parametric system of differential equations (equations of motions)  $\dot{x} = V_q(x)$ , where  $V_q : X \rightarrow \mathbb{R}^{|A|}$  is given by

$$\dot{x}_{k;a_k} = x_{k;a_k}^q \left( u_k(a_k; x_{-k}) - \frac{\sum_{a_j \in A_j} x_{k;a_j}^q u_k(a_j; x_{-k})}{\sum_{a_j \in A_j} x_{k;a_j}^q} \right); \text{ for all } k \in 2N; a_k \in A_k; \text{ (QRD)}$$

for any  $q \geq 0$ . Special cases of the above dynamics are the projection gradient (PG) dynamics, for  $q = 0$ , the (standard) replicator (RD) dynamics, for  $q = 1$ , and the log-barrier or inverse update dynamics, for  $q = 2$  (Mertikopoulos & Sandholm, 2018; Giannou et al., 2021).

### 3 POINTWISE CONVERGENCE OF QRD TO NASH EQUILIBRIA

Our results consist of two parts. In the first part, which is the subject of this section, we show convergence of QRD to Nash equilibria in a subclass of potential games we call Perfectly-Regular Potential Games (PRPG) that contains almost all finite potential games.

**Definition 3.1 (Perfectly-regular potential games)** A potential game is called regular if it has only regular Nash equilibria. A regular potential game is called perfectly-regular potential game (PRPG) if all its restrictions are regular potential games, i.e., if they only possess regular Nash equilibria.

As we show in Lemma B.2 in the Appendix, almost all finite potential games are PRPGs; this is a direct generalization of Swenson et al. (2020) who prove that almost all potential games are regular. The PRPG class contains many widely-studied subclasses of games, e.g., congestion games, or games with identical reward functions (Wang & Sandholm, 2002; Panait & Luke, 2005; Carroll et al., 2019; Dafoe et al., 2020). Examples of non-PRPG games are degenerate games with equal payoffs, e.g., a two-player game with actions  $A_1 = A_2 = \{a_1, a_2\}$  and payoffs  $u_k(a_i, a_j) = 1$  for all  $i, j \in A_k$  and  $k = 1, 2$ . The convergence result of QRD in this class is stated formally in Theorem 3.2.

**Theorem 3.2 (Pointwise convergence of QRD to NE in PRPGs)**. Given any perfectly-regular potential game (PRPG), and any interior initial condition  $x(0) \in \text{int } X$ , the  $q$ -replicator dynamics, defined as in equation (1), converge pointwise to a Nash equilibrium of  $G$  for any parameter  $q \geq 0$ . Furthermore, the set  $\text{set}(\text{int } X) := \{x \in \text{int } X \mid \exists x_0 \in \text{int } X, x \in \lim_{t \rightarrow \infty} x(t) = x; x(0) = x_0\}$ , i.e., the set of all limit points of interior initial conditions, is finite.

Importantly, Theorem 3.2 states that QRD converge to some NE for almost all initial conditions in almost all potential games. A direct implication is that when reasoning about the quality of the collective learning outcome in cooperative multi-agent settings, as captured by PRPGs, we can restrict our attention to NE. This is the subject of the next section. Before proceeding with this, we provide a sketch of the proof of Theorem 3.2 (cf. Appendix B for the formal proof).

The proof of Theorem 3.2 proceeds in two steps, which utilize that PRPGs have a finite number of regular equilibria, and (ii) the probability of optimal actions near an equilibrium point is increasing in time with respect to QRD. In the first step, we prove that for any initial condition, the sequence of joint action profiles  $x(t)_{t \geq 0}$ , that is generated by QRD for any  $q \geq 0$  converges to a restricted equilibrium of a PRPG. This relies on the fact that the set of cluster (limit) points of the trajectory—also called the  $\omega$ -limit set—is a finite, and in fact, as we show, a single element set for a PRPG. In turn, this holds because a PRPG provably contains only a finite number of restricted equilibria. In the second step, we show that any such limit point has to be a NE of  $G$ . To establish this, we exclude convergence to restricted equilibria that are not NE by coupling the structure of PRPGs, which ensures that there is a finite number of (regular) restricted equilibria, with the natural QRD which guarantees that in the vicinity of a limit point, optimal actions, i.e., best responses, need to be played with increasingly higher probability. As a result, all actions in the support of the limit choice distribution of each agent must be best responses against the actions of all other agents, which implies that all points that can be reached by QRD are NE of  $G$ .

### 4 PERFORMANCE OF THE COLLECTIVE LEARNING OUTCOME

Beyond static performance metrics. We next turn our attention to the main challenge of quantifying the quality of the collective learning outcome. To do that, we need to derive appropriate performance metrics. In static regimes, we can rely on a variety of meaningful metrics, e.g., the Price of Anarchy (PoA) (Koutsoupias & Papadimitriou, 1999; Christodoulou & Koutsoupias, 2005; Roughgarden, 2015), which is defined as the ratio between the socially worst NE of the game and the socially optimal outcome (in terms of the agents' sum of rewards). However, despite the useful insights that PoA provides in general games, it is not difficult to find PRPGs in which the PoA fails to provide any meaningful information about the game. Let us consider the following example:

Example 4.1 (A simple example of unbounded performance loss) Consider a parametrized 2-PRPG with identical payoff functions  $u_{w;1}(s_1; s_2) = u_{w;2}(s_2; s_1) = A_w(s_1; s_2)$ , where matrix  $A_w \in \mathbb{R}^{2 \times 2}$  is given by  $A_w = \text{diag}(1; w)$ , with  $1 \leq w$ .<sup>3</sup> The game  $\Gamma_w$  has two pure NE, one that corresponds to  $x_1^0 = (1; 0)$  and  $x_2^0 = (1; 0)$  with social welfare  $SW(x^0) = 1 + 1 = 2$ , and one that corresponds to  $x_1^0 = (0; 1)$  and  $x_2^0 = (0; 1)$  with social welfare  $SW(x^0) = w + w = 2w$ . Since  $w$  can take any value larger than 1, the difference in performance can be arbitrary large with respect to the PoA. Specifically,  $\text{PoA}(\Gamma_w) = \frac{SW(x^0)}{SW(x^*)} = w \neq 1$  as  $w \neq 1$ .

While useful in static environments, the PoA metric does not capture the dynamic nature of multi-agent learning. In particular, it does not provide an answer to the question: how likely is it for the agents to reach a good or bad outcome given that the multi-agent system converges?

#### 4.1 REGIONS OF ATTRACTION AND AVERAGE PERFORMANCE METRICS

To answer the above question and argue about the collective performance of the game dynamics, we need to quantify the likelihood of each outcome when the initial conditions of the system are randomly sampled. A region of attraction of a given outcome formalizes this notion.

Definition 4.2 (Regions of attraction) Let  $\Gamma$  be any game and assume that its joint action profile,  $x \in X$ , is evolving according to the equations of motion  $\dot{x} = f(x)$ . Then for any  $x^* \in X$ , the set  $\text{RoA}_i(x^*) := \{x_0 \in X \mid \lim_{t \rightarrow \infty} x(t) = x^*; x(0) = x_0\}$  is called the region of attraction (RoA) of  $x^*$  with respect to the dynamics.

In other words, the RoA of a point  $x^* \in X$  is the set of all initial conditions  $x_0$  for which the dynamics asymptotically converge to  $x^*$ . Note that RoAs do not intersect. If we can determine the regions of attraction of some game dynamics, then given a certain static performance metric, e.g., the social welfare, we can define a corresponding average-performance metric that weighs-in all possible outcomes, in the sense of limit points, according to their likelihood of occurring with respect to the given dynamics. In order for this average to be meaningful, a minimum requirement, is that the dynamics converge for almost all, i.e., all but a measure zero, initial conditions. An Average Performance Metric (APM), refined from (Panageas & Piliouras, 2016), is defined as follows

Definition 4.3 (Average-performance metric) Let  $\Gamma$  be a multi-agent game and assume that its joint action profile,  $x \in X$ , is evolving according to the equations of motion  $\dot{x} = f(x)$ . Let  $X_0 \subseteq X$  be a set of initial conditions such that the set of convergence points  $\mathcal{C}(X_0)$  is finite. Then, given a performance metric  $g: X \rightarrow \mathbb{R}$ , the average-performance of the dynamics governed by  $f$  with respect to the performance metric  $g$  and the set of initial conditions  $X_0$ , is given by

$$\text{APM}_{g;X_0}(f; \mu) := \int_{x \in \mathcal{C}(X_0)} (\text{RoA}_i(x)) g(x) \mu(dx); \quad (\text{APM})$$

where  $\mu$  is a probability measure on  $X_0$ .

In other words, APM is the expected optimality of a random initialization of the dynamics  $\dot{x} = f(x)$  with respect to some metric  $g$ . For instance, if the performance metric is the social welfare, then the average-performance metric with respect to  $\mu$  measures the expected social welfare of the system for any random initialization in  $X_0$ . The average-performance metric that we are going to use in the remainder of this section is the Average Price of Anarchy (APoA) (Panageas & Piliouras, 2016). The APoA is an APM with respect to the social welfare, re-normalised such that APoA is greater than or equal to 1, with equality only if (almost) all the initial conditions converge to the socially optimal outcome of the system. Formally, given a multi-agent game, equations of motion  $\dot{x} = f(x)$  that describe the evolution of the agents' actions, and a set of initial conditions  $X_0 \subseteq X$  that consists of almost all points in  $X$ , the APoA is given by

$$\text{APoA}(f; \mu) := \frac{\max_{x \in X} SW(x)}{\text{APM}_{SW;X_0}(f; \mu)}; \quad (\text{APoA})$$

<sup>3</sup>Here,  $\text{diag}(a_1; a_2)$  denotes a matrix with  $a_1, a_2$  on the diagonal and zeros otherwise. Also, we write both  $x; y \in [0; 1]$  to denote the mixed choice distribution  $(x; 1-x); (y; 1-y)$ , of players 1 and 2, respectively.

<sup>4</sup>Recall that a probability measure on a compact space  $X$  is countably-additive function from the powerset of  $X$  to  $\mathbb{R}_+$  such that  $\mu(X) = 1$  and  $\mu(\emptyset) = 0$ .

Notice that a large APoA, like a large PoA, is a negative trait that depends on the game, but, in contrast to PoA, it also depends on the game dynamics. In other words, a low APoA is an indication of better performance in compared to some other dynamics with larger APoA.

Remark 4.4. As we mentioned above, Definition 4.3 does not ensure that APM is always a meaningful metric for the system. However, as long as one can prove that (i) the dynamics converge pointwise to some  $x \in Q(X) \cap NE(\cdot)$  for almost all initial conditions  $x_0 \in X$ , and (ii) the set of limit points,  $Q(X)$ , is finite—two conditions that are satisfied by any PRPG that evolves with respect to some QRD (cf. Theorem 3.2)—the APoA has an intuitive interpretation. Specifically, in this setup, the APoA is always bounded between the PoA and the Price of Stability (PoS) of the game, i.e., the ratio between the socially optimal outcome and the socially optimal NE (Roughgarden, 2015).

## 4.2 THE TAXONOMY OF QRD IN 2-2 PRPGs

Having established appropriate performance measures and PRPGs are an appropriate class in which standard learning dynamics, like QRD, can be compared in terms of performance, in this section, we perform a complete theoretical analysis of our two motivating questions in the subclass of symmetric 2-2 PRPGs. The main takeaway of this analysis is that if the payoff-dominant NE requires relatively low risk for the players compared to the one of the game, i.e., if it is risk-dominant, the GD performs better than RD, and vice versa, if the payoff-dominant NE of the game fails to be the risk-dominant one, RD performs better than GD. To the best of our knowledge, this is the first rigorous analysis on the relative performance between GD and RD. Concerning our second question, we establish that if the payoff- and risk-dominant equilibria coincide, the APoA of GD is upper bounded by 1. It is important to note that the class of symmetric 2-2 PRPGs is a non-trivial regime for the comparison of the two dynamics, since PoA is provably unbounded in this setup, cf. Example 4.1. Omitted materials may be found in Appendix C.

Representation of symmetric 2-2 PRPGs. We begin this section by showing that any 2-2 symmetric PRPG is equivalent to a game  $w$ , as defined in Lemma 4.5. The only non-generic games that are excluded from this reformulation are dominance-solvable games whose analysis is trivial and, therefore, outside of our scope. To proceed with the formal statement of Lemma 4.5, recall that a NE  $x^*$  of a symmetric potential game is called payoff-dominant if  $u_k(x^*) \geq u_k(x^0)$  for all  $x^0 \in NE(\cdot)$ , and it is called risk-dominant if  $x^*$  is unilaterally optimal against the uniform distribution of the rest of the agents (Harsanyi & Selten, 1988).

Lemma 4.5. Any 2-2 symmetric PRPG  $\Gamma$ , with payoff functions  $u_1(s_1; s_2) = u_2(s_2; s_1)$  can be equivalently represented by a game  $w$  with payoff functions  $u_{w;1}(s_1; s_2) = u_{w;2}(s_2; s_1) = A_{w;1; s_1; s_2}$ , where the matrix  $A_{w;1; s_1; s_2} \in \mathbb{R}^2 \times \mathbb{R}^2$  is given by  $A_{w;1; s_1; s_2} = \begin{pmatrix} 1 & 0 \\ w & w \end{pmatrix}$ , for  $s_1 = w$ . The game

$w$  has the same NE as the original game  $\Gamma$ , retains its payoff- and risk-dominance properties, and preserves the limiting behavior of any QRD in  $\Gamma$ . Accordingly, each game  $w$  has three NE, two pure at  $x = y = 0$  and  $x = y = 1$ , with social welfare  $SW(0; 0) = 2w$  and  $SW(1; 1) = 2$ , respectively, as well as one fully-mixed NE at  $x = y = \frac{1}{w+1}$ .

We are going to refer to the first pure NE of  $w$  as  $x_w$ . Note that  $x_w$  is payoff-dominant for any  $w > 0$ , and it is also risk-dominant whenever  $w > 1$ , or equivalently, whenever  $w > 0.5$ .

Replicator dynamics (RD) versus gradient descent (GD). The first result of this section is that whenever the risk- and payoff-dominant equilibria of  $w$  coincide, i.e.,  $w > 0.5$ , then the gradient descent (GD) dynamics, i.e., the replicator dynamics, perform better (or equally in the generic case  $w = 0.5$ ) on average than the standard replicator dynamics (RD) with respect to the social welfare of their outcomes, i.e., they yield a smaller APoA. In any other instance of these games, i.e., for  $w < 0.5$ , the RD perform better than GD with respect to the same metric.

Theorem 4.6 (Comparative Performance of RD and GD in 2-2 PRPGs). Consider an arbitrary 2-2 symmetric PRPG which, without any loss of generality, can be represented as an instance  $w$  and let  $V_0; V_1$  denote the equations of motion of the replicator, i.e., gradient descent dynamics, and 1-replicator, i.e., replicator dynamics, respectively, cf. equation QRD. Then, it holds that

$$\text{APM}_{\text{SW}; \text{int } X}(V_0; w) \leq \text{APM}_{\text{SW}; \text{int } X}(V_1; w); \quad (1)$$

if and only if the payoff-dominant equilibrium is also risk-dominant. Equality obtains only when  $w = 0.5$ , i.e., when  $w = 1$ .

Figure 1: The invariant function,  $\varphi_q(x; y)$ , for all  $x, y \in [0; 1]^2$  (cf. Lemma 4.7) in the game  $w_i$  for  $w = 2$ ,  $\alpha = 0$ , and  $q = 0$  (gradient descent),  $q = 1$  (standard replicator),  $q = 2$  (log-barrier), and  $q = 20$ . The invariant function becomes steep at the boundary as  $q$  increases, taking both arbitrarily large negative (dark) and positive (light) values in the vicinity of the NE. The contour lines at zero, highlighted as solid blue and dashed red lines on the surfaces, correspond to the stable  $x = 0$ , and unstable  $x = 1$ , manifolds, respectively.

Theorem 4.6 demonstrates that the riskiness of the socially optimal equilibrium is decisive in characterising the relative performance of GD and RD. In particular, as stated in Theorem 4.6, the expected social welfare is optimized by GD whenever risk and payoff-dominant equilibria coincide and is optimized by RD when risk and payoff-dominant equilibria differ. From a practical perspective, Theorem 4.6 provides a concrete recommendation on the optimal behavior of the agents (GD vs RD) based solely on the properties of the underlying game. However, it also suggests that even in low-dimensional  $2 \times 2$  potential games, there is not a uniform recommendation, and the optimal behavior largely depends on the features of the underlying game. As it turns out, in this case, the decisive feature is the riskiness of the payoff-dominant equilibrium (see also Appendix C).

**Invariant functions.** Technically, the proof of Theorem 4.6 (see Appendix C) proceeds with a first order analysis of the manifolds that separate the regions of attractions of the two pure equilibria for the different dynamics. This approach leverages the existence of motion-invariant functions (Nagarajan et al., 2020), i.e., quantities that remain constant along the trajectories of the learning dynamics. The rationale is that if we can identify such a function, then, by finding its value at the unique mixed equilibrium of the game, we can determine all initial conditions that asymptotically converge to it: these will be all points at the same level set of the invariant function. The manifold, i.e., the geometric locus of all the points that converge to the equilibrium, formally stable manifold  $\mathcal{M}^s$ , is the one that separates the regions of attractions of the two NEs of the game. Because of this property, we may also refer to the stable manifold of the MNE as the separatrix (Panageas & Piliouras, 2016). Note that, since the dynamics are also backward-invariant (Panageas & Piliouras, 2016; Mertikopoulos & Sandholm, 2018), their level-set will also contain a set of initial conditions that converge to it when moving backward in time. These points constitute the unstable manifold  $\mathcal{M}^u$ . In Lemma 4.7, we identify such an invariant for all QRD.

**Lemma 4.7 (Invariant functions of QRD in  $2 \times 2$  symmetric PRGs).** Given a  $2 \times 2$  symmetric PRG  $w_i$ , whose agents evolve with respect to the replicator dynamics QRD, the separable function  $\varphi_q : (0; 1)^2 \rightarrow \mathbb{R}$ , with  $\varphi_q(x; y) := \varphi_q(x) - \varphi_q(y)$ , and  $\varphi_q : (0; 1) \rightarrow \mathbb{R}$  given by

$$\varphi_q(x) = \begin{cases} \frac{x^{2-q} + (1-x)^{2-q}}{2-q} + \frac{1-x^{1-q}}{1-q} - \frac{(1-x)^{1-q}}{1-q}; & q \in [1; 2); \\ \ln(x) + (1-x)\ln(1-x); & q = 1; \\ \ln(x) + \ln(1-x) + \frac{1}{x} + \frac{1}{1-x}; & q = 2; \end{cases} \quad (2)$$

remains constant along any trajectory  $(x(t); y(t))_{t \geq 0}$  of the system. The function  $\varphi_q(x)$  is continuous with respect to the parameter  $q$  at both,  $q = 1$  and  $q = 2$ , since  $\lim_{q \downarrow 1} \varphi_q(x) = \varphi_1(x)$  and  $\lim_{q \downarrow 2} \varphi_q(x) = \varphi_2(x)$  for all  $x \in (0; 1)$ .

**Monotonicity of RoAs for all QRD.** A rigorous extension of Theorem 4.6 to all QRD is prevented by the difficulty to obtain an analytical form of the separatrix for all  $q \in [1; 2)$ . However, by visualizing the RoAs for all QRD in this range, we obtain empirical evidence that RoA is monotonically increasing in  $q$ .

In Figure 1, we visualize the invariant function  $\varphi_q(x; y)$ , for  $x, y \in (0; 1)^2$  defined in Lemma 4.7 for various values of  $q \in [0; 20]$ . From the panels of Figure 1, it is also evident that  $\varphi_q(x; y)$  acts as a handy tool to visualize the regions of attraction of the two NEs of the game. Namely, at the unique

mixed NE, i.e., at  $x = y = \frac{1}{2}$ , the invariant function,  $q$ , is equal to 0. The same holds for any point  $(x; y) \in (0; 1)^2$  with  $x = y$ . Thus, we can factorize  $q(x; y)$  as  $q(x; y) = q_{\text{Stable}}(x; y) (x - y)$  where  $q_{\text{Stable}}(x; y) = 0$  is precisely the geometric locus of all points  $(x; y) \in (0; 1)^2$  such that  $\lim_{t \rightarrow \infty} x(t) = y$ , and  $x = y$  is the geometric locus of all points such that  $\lim_{t \rightarrow \infty} x(t) = x$ . These two manifolds constitute the stable and unstable manifolds of the replicator dynamics.

If we stack the stable manifolds (solid blue lines in Figure 1), it becomes evident that the region of attraction of the payoff-dominant and risk-dominant equilibrium grows as  $q$  decreases to 0. This is depicted in Figure 2 for all values of  $q \in [0; 10]$ . In this case, the unstable manifolds are always equal to the diagonal  $x = y$  and thus, omitted. The progression of the surface of the stable manifolds continuous for larger values of  $q$ . Analogous plots, but with the results reversed as predicted by Theorem 4.6, can be generated for instances of  $w_1$ , in which the risk-dominant equilibrium is different from the payoff-dominant one, as well as, for 2 generic PRPGs. In general, putting together Theorem 4.6 and the current visualizations, we have both theoretical and empirical evidence that the region of attraction of the payoff-dominant equilibrium in  $w_1$  is decreasing (increasing) in  $q$  for  $q > 0$  whenever this equilibrium is (is not) risk-dominant.

Figure 2: Stable manifold (separatrix) for all different values of  $q \in [0; 10]$  (from blue to brown) in the  $w_1$  game for  $w = 2$  and  $\beta = 0$ . The manifolds for  $q = 0$ ,  $q = 1$ , and  $q = 2$  are shown in shades of black for reference (cf. Figure 5). The region of attraction of the payoff-dominant equilibrium (bottom-left corner) shrinks as  $q$  increases.

Application: APoA in 2x2 PRPGs. To showcase the practical importance of Theorem 4.6 and the invariant function approach, we conclude this section with a concrete evaluation of the APoA average-performance measure in the class of 2 symmetric PRPGs. For this result, we focus on symmetric 2x2 PRPGs such that payoff and risk-dominant equilibria coincide; one can prove tight bounds on APoA as shown in Theorem 4.8.

Theorem 4.8. The APoA of GD dynamics in all 2x2 symmetric PRPGs,  $w_1$ , in which the payoff- and risk-dominant NE coincide is bounded by  $\frac{1}{2}$ , i.e.,  $\text{APoA}(V_0; w_1) \leq \frac{1}{2}$ , for all  $w_1$  with  $\beta > \frac{1}{2} w$ . This bound is tight.

Figure 3: APoA of a 2x2 symmetric PRPG for the gradient descent dynamics and various values of  $\beta$  and  $w$ . The APoA is upper bounded by (dark to light values) as shown in Theorem 4.8.

The proof of Theorem 4.8 proceeds by first order analysis of the function depicted in Figure 3 which, in turn, depends on the invariant function of the gradient descent dynamic, see Appendix C. Note that the bound of Theorem 4.8 also holds for  $\beta < \frac{1}{2} w$ , but in this case, there exists no risk-dominant equilibrium. Moreover, one way to check that this bound is tight, is to set  $\beta = \frac{1}{2} w + \epsilon$ , for a small  $\epsilon > 0$  and let  $w$  increase (cf. Figure 3). In combination Theorem 4.6 and Theorem 4.8 imply that the APoA of the RD (QRD with  $q = 1$ ), is not upper bounded by  $\frac{1}{2}$  whenever  $\beta < 0.5$ , i.e., whenever the risk- and payoff-dominant equilibria are different. However, for the case  $\beta > 0.5$ , the separatrices for  $\beta = 0$  as visualized in Figure 2, (empirically) imply that similar bounds hold for all values of  $\beta$ .

<sup>5</sup>In Figure 2, we visualize the stable manifolds which have the largest region of attraction, i.e., the lowest APoA. The case  $\beta < 0.5$ , in which the manifolds are mirrored on the  $x = 1 - x$  diagonal, is in Appendix C.

m	APoA std (Max APoA)			
	RD		GD	
3	1:16	:05(1:23)	1:12	:04(1:19)
5	1:19	:06(1:30)	1:14	:04(1:22)
10	1:22	:05(1:38)	1:13	:03(1:20)
20	1:21	:04(1:28)	1:11	:02(1:22)

Figure 4 & Table 1: Numerical results for the APoA metric of replicator dynamics (RD) and gradient descent (GD) in diagonal PRPGs of dimensions  $m = 2; \dots; 20$ . The two panels show the APoA of each dynamic together with its standard deviation (left), their difference (solid line enveloped by the standard deviation shaded region) and the difference of the maximum APoAs of the dynamics (right) observed over 100 sampled games at each dimension. The GD dynamics throughout outperform the RD dynamics suggesting that Theorem 4.6, with the notions of pure- and risk-dominance properly re-named, may still hold for larger dimensional PRPGs (for the reverse implication in Theorem 4.6 see Section 5). Table 1 provides detailed statistics for  $m = 3, 5, 10$  and  $20$ .

### 5 HIGHER DIMENSIONAL NUMERICAL EXPERIMENTS

In this section, we present experimental evidence that our approach can be extrapolated on larger setups, i.e., beyond 2-games. Notice that, the extrapolation needs to be done with care since the central notion of risk dominance may not be well-defined in larger game setups. We choose to treat this, by restricting our focus to the class of two-player diagonal PRPGs. Such games can be described by a common payoff matrix with the only non-zero elements of which lie on its diagonal, i.e., it is of the form,  $U := \text{diag}(u_1; u_2; \dots; u_m)$ , where w.l.o.g, we assume  $0 < u_1 < u_2 < \dots < u_m$ . Notice that a diagonal PRPG defined as above has the following important traits:

1. Profiles  $(1; 1; \dots; m; m)$  are the only pure NE and, since the game is PRPG the only stable NE.
2. The pure NE can be ordered in terms of payoffs as  $u_1 < \dots < u_m$ . Consequently the NE  $(i; i)$  dominates all NE  $(j; j)$ ,  $j = 1; \dots; i - 1$ .
3. The difference in payoffs from any pure NE  $(i; i)$  to some non-diagonal profile  $(i; j)$ , or  $(j; i)$ , is  $u_i$  for any  $j \notin i$ . If we define risk to be the negative deviation payoffs Harsanyi & Selten (1988), then the risk of NE  $(i; i)$  is  $u_i^2$  for all  $i = 1; \dots; m$ , and therefore the pure NE can be ordered in terms of risk with NE  $(i; i)$  be less risky than  $(j; j)$  for all  $j = 1; \dots; i - 1$ .

Experimental setup. We run experiments in random agent, symmetric diagonal PRPGs of dimensions  $m = 2; 3; \dots; 20$  (size of each agent’s action space). In each game, the payoffs  $u_1; \dots; u_m$  are selected (pseudo-)randomly and satisfy the following properties: (i) the lowest diagonal payoff,  $u_1$ , is at least as large as some predefined positive constant (set equal to 2 for the experiments), (ii) the highest payoff,  $u_m$ , is equal to the dimension  $m$ , of the game, i.e.  $u_m = 2; 3; \dots$  and  $20$  respectively, and (iii)  $u_2; \dots; u_{m-1}$  are in ascending order strictly between  $u_1$  and  $u_m$  with randomly selected distances between them. For each dimension, we sample 100 random games and run the gradient descent and standard replicator dynamics from random initial conditions till convergence.

The outputs of the simulations of the above experiments are summarized in Figure 4 and Table 1. The outputs provide indications that (i) the gradient descent dynamics (continue to) outperform the replicator dynamics in all diagonal games in terms of average performance, and that (ii) the APoA is lower-bounded by  $\frac{1}{2}$ , the theoretical bound that we presented in the case of 2 games. As Table 1 suggests, the APoA is, in fact, considerably lower than all sampled games.

### 6 CONCLUSION

We study the class of replicator dynamics (RD), and showed that all RD converge pointwise to Nash equilibria in perfectly-regular potential games, a class of games that encompasses almost all potential games, i.e., the standard models of multi-agent coordination. From the perspective of equilibrium selection and quality, however, convergence provides little information. Even if two dynamics converge and even if they have joint optimal no-regret guarantees, they may still exhibit very different equilibrium selection properties, which, in turn, determine their practical performance. Our analysis provides the first formal comparative analysis of different optimization-driven dynamics by combining diverse techniques and establishes an intriguing direction for future work.

## ACKNOWLEDGMENTS

This research was supported in part by the National Research Foundation, Singapore and DSO National Laboratories under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), grant PIESGP-AI-2020-01, AME Programmatic Fund (Grant No.A20H6b0151) from A\*STAR and Provost's Chair Professorship grant RGEPPV2101. Ioannis Panageas would like to acknowledge startup grant from UCI. Part of this work was conducted while Ioannis was visiting Archimedes Research Unit. This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

## REFERENCES

- I. Anagnostides, I. Panageas, G. Farina, and T. Sandholm. On last-iterate convergence beyond zero-sum games. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato (eds.), Proceedings of the 39th International Conference on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pp. 536–581. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/anagnostides22a.html>
- J. P. Bailey and G. Piliouras. Fast and Furious Learning in Zero-Sum Games: Vanishing Regret with Non-Vanishing Step Sizes. Proceedings of the 33rd International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 2019. Curran Associates Inc.
- David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In Jennifer Dy and Andreas Krause (eds.), Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, pp. 354–363. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/balduzzi18a.html>
- N. Bard, J. N. Foerster, S. Chandar, N. Burch, M. Lanctot, H. F. Song, E. Parisotto, V. Dumoulin, S. Moitra, E. Hughes, I. Dunning, S. Mourad, H. Larochelle, M. G. Bellemare, and M. Bowling. The Hanabi challenge: A new frontier for AI research. Artificial Intelligence, 280:103216, 2020. ISSN 0004-3702. doi: 10.1016/j.artint.2019.103216.
- Y. Cai, A. Oikonomou, and W. Zheng. Tight Last-Iterate Convergence of the Extragradient Method for Constrained Monotone Variational Inequalities. CoRR, abs/2204.09228, 2022. URL <https://doi.org/10.48550/arXiv.2204.09228>
- M. Carroll, R. Shah, Mark K. Ho, T. L. Griffiths, S. A. Seshia, P. Abbeel, and A. Dragan. On the Utility of Learning about Humans for Human-AI Coordination. Proceedings of the 33rd International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 2019. Curran Associates Inc.
- G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC '05, pp. 67–73, New York, NY, USA, 2005. Association for Computing Machinery. doi: 10.1145/1060590.1060600.
- J. Cohen, A. Héliou, and P. Mertikopoulos. Learning with Bandit Feedback in Potential Games. In Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 6372–6381, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- A. Dafoe, E. Hughes, Y. Bachrach, T. Collins, K. R. McKee, J. Z. Leibo, K. Larson, and T. Graepel. Open Problems in Cooperative AI, 2020. URL <https://arxiv.org/abs/2012.08630>
- A. Dafoe, Y. Bachrach, G. Hadfield, E. Horvitz, K. Larson, and T. Graepel. Cooperative AI: machines must learn to find common ground. Nature, 7857:33–36, 2021. doi: 10.1038/d41586-021-01170-0.
- C. Daskalakis and I. Panageas. Last-Iterate Convergence: Zero-Sum Games and Constrained Min-Max Optimization. In Avrim Blum (ed.), 10th Innovations in Theoretical Computer Science Conference (ITCS 2019) volume 124 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 27:1–27:18, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-095-8. doi: 10.4230/LIPIcs.ITCS.2019.27. URL <http://drops.dagstuhl.de/opus/volltexte/2018/10120>

- Y. N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, pp. 2933–2941, Cambridge, MA, USA, 2014. MIT Press.
- L. Flokas, E. V. Vlatakis-Gkaragkounis, T. Lianas, P. Mertikopoulos, and G. Piliouras. No-Regret Learning and Mixed Nash Equilibria: They Do Not Mix. *Proceedings of the 34th International Conference on Neural Information Processing Systems NIPS'20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping From Saddle Points — Online Stochastic Gradient for Tensor Decomposition. In P. Grünwald, E. Hazan, and S. Kale (eds.), *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pp. 797–842, Paris, France, 03–06 Jul 2015. PMLR. URL: <https://proceedings.mlr.press/v40/Ge15.html>.
- A. Giannou, E.-V. Vlatakis-Gkaragkounis, and P. Mertikopoulos. On the Rate of Convergence of Regularized Learning in Games: From Bandits and Uncertainty to Optimism and Beyond. Part of *Advances in Neural Information Processing Systems 34 pre-proceedings (NeurIPS 2021)*, 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/bf40f0ab4e5e63171dd16036913ae828-Abstract.html>.
- Wassim M Haddad and VijaySekhar Chellaboina. *Nonlinear dynamical systems and control: a Lyapunov-based approach*. Princeton university press, 2008.
- J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2(1):235–250, Dec 1973. ISSN 1432-1270. doi: 10.1007/BF01737572. URL <https://doi.org/10.1007/BF01737572>.
- J. C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. Cambridge: MIT Press., 1988.
- J. Honda. Equilibrium selection for symmetric coordination games with an application to the minimum-effort game. *CiteSeerX*, 2012.
- M. Kaisers and K. Tuyls. FAQ-Learning in Matrix Games: Demonstrating Convergence near Nash Equilibria, and Bifurcation of Attractors in the Battle of Sexes. *Proceedings of the 13th AAAI Conference on Interactive Theory and Game Theory AAIWS'11-13*, pp. 36–42. AAAI Press, 2011.
- A. Kianercy and A. Galstyan. Dynamics of Boltzmann learning in two-player two-action games. *Phys. Rev. E*, 85:041145, Apr 2012. doi: 10.1103/PhysRevE.85.041145.
- Y. Kim. Equilibrium Selection in n-Person Coordination Games. *Games and Economic Behavior*, 15(2):203–227, 1996. doi: 10.1006/game.1996.0066.
- R. Kleinberg, G. Piliouras, and E. Tardos. Multiplicative Updates Outperform Generic No-Regret Learning in Congestion Games: Extended Abstract. *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, pp. 533–542, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585062. doi: 10.1145/1536414.1536487.
- E. Koutsoupias and C. Papadimitriou. Worst-Case Equilibria. In Christoph Meinel and Sophie Tison (eds.), *STACS 99*, pp. 404–413, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg. doi: 10.1007/3-540-49116-3\_38.
- J. D. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. I. Jordan, and B. Recht. First-order methods almost always avoid strict saddle points. *Math. Program.*, 176(1-2):311–337, 2019.
- S. Leonardos and G. Piliouras. Exploration-exploitation in multi-agent learning: Catastrophe theory meets game theory. *Artificial Intelligence*, 304:103653, 2022. ISSN 0004-3702. doi: 10.1016/j.artint.2021.103653.
- P. Mertikopoulos and W. H. Sandholm. Learning in Games via Reinforcement and Regularization. *Mathematics of Operations Research*, 41(4):1297–1324, 2016. doi: 10.1287/moor.2016.0778.

- P. Mertikopoulos and W. H. Sandholm. Riemannian game dynamics. *Journal of Economic Theory* 177:315–364, 2018. ISSN 0022-0531. doi: 10.1016/j.jet.2018.06.002.
- P. Mertikopoulos, C. Papadimitriou, and G. Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the 2018 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2703–2717, 2018. doi: 10.1137/1.9781611975031.172.
- P. Mertikopoulos, B. Lecouat, H. Zenati, C.-S. Foo, V. Chandrasekhar, and G. Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019 OpenReview.net*, 2019. URL <https://openreview.net/forum?id=Bkg8jjC9KQ>.
- S. G. Nagarajan, D. Balduzzi, and G. Piliouras. From Chaos to Order: Symmetry and Conservation Laws in Game Dynamics. In Hal Daumé III and Aarti Singh (eds), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7186–7196. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/nagarajan20a.html>.
- Joshua Paik and Christopher Griffin. Completely integrable replicator dynamics associated to competitive networks. *Phys. Rev. E* 107:L052202, May 2023. doi: 10.1103/PhysRevE.107.L052202. URL <https://link.aps.org/doi/10.1103/PhysRevE.107.L052202>.
- G. Palaiopanos, I. Panageas, and G. Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*, pp. 5874–5884, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- I. Panageas and G. Piliouras. Average case performance of replicator dynamics in potential games via computing regions of attraction. *Proceedings of the 2016 ACM Conference on Economics and Computation (EC '16)*, pp. 703–720, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450339360. doi: 10.1145/2940716.2940784.
- L. Panait and S. Luke. Cooperative Multi-Agent Learning: The State of the Art. *Autonomous Agents and Multi-Agent Systems* 11(3):387–434, Nov 2005. doi: 10.1007/s10458-005-2631-2.
- M. Pangallo, J. B.T. Sanders, T. Galla, and J. D. Farmer. Towards a taxonomy of learning dynamics in  $2 \times 2$  games. *Games and Economic Behavior* 132:1–21, 2022. ISSN 0899-8256. doi: 10.1016/j.geb.2021.11.015.
- Georgios Piliouras and Xiao Wang. Constants of motion: The antidote to chaos in optimization and game dynamics, 2021.
- T. Roughgarden. Intrinsic Robustness of the Price of Anarchy. *ACM*, 62(5), nov 2015. ISSN 0004-5411. doi: 10.1145/2806883.
- Y. Sato, E. Akiyama, and J. P. Crutchfield. Stability and diversity in collective adaptation. *Physica D: Nonlinear Phenomena* 210(1):21–57, 2005.
- D. Schmidt, R. Shupp, J. M. Walker, and E. Ostrom. Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior* 42(2): 281–299, 2003. doi: 10.1016/S0899-8256(02)00552-3.
- B. Swenson, R. Murray, and S. Kar. Regular potential games. *Games and Economic Behavior* 124: 432–453, 2020. ISSN 0899-8256. doi: 10.1016/j.geb.2020.09.005.
- E. van Damme. *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin, Heidelberg, 1987. ISBN 0387171010.
- X. Wang and T. Sandholm. Reinforcement Learning to Play an Optimal Nash Equilibrium in Team Markov Games. In *Proceedings of the 15th International Conference on Neural Information Processing Systems (NIPS'02)*, pp. 1603–1610, Cambridge, MA, USA, 2002. MIT Press.
- B. Zhang and J. Hofbauer. Equilibrium selection via replicator dynamics in 2 coordination games. *International Journal of Game Theory* 44(2):433–448, May 2015. ISSN 1432-1270. doi: 10.1007/s00182-014-0437-7.

# SUPPLEMENTARY MATERIAL FOR BEATING PRICE OF ANARCHY AND GRADIENT DESCENT WITHOUT REGRET IN POTENTIAL GAMES

## A EXTENSIONS AND CHALLENGES

A rigorous average-case analysis in arbitrary games is mathematically and computationally challenging, since closed-form solutions for the boundaries of RoAs not always exist, and since computing RoAs is generally NP-Hard. Nevertheless, it is immediate to extend the empirical analysis by approximating the APM of interest through counts of frequencies for each limiting equilibrium for multiple random initializations. However, the invariant function approach offers a scalable alternative that is promising to work in more complicated settings. Finally, non-uniform distributions over initial conditions that keep the expectation in Definition 4.3 well-defined can be used to study games in which the dynamics converge to limit cycles instead of NE.

## B MISSING PROOFS AND MATERIALS: SECTION 3

To complete the steps in the sketch of the proof of Theorem 3.2, we start by showing that the class of perfectly-regular potential games (PRPG) is well-defined, since restrictions of potential games are also potential games.

Lemma B.1. A restriction of a potential game is also a potential game.

Proof of Lemma B.1 Let  $\gamma^0$  a restriction of a potential game and potential function:  $A \rightarrow \mathbb{R}$ . Take  $\gamma^0: A^0 \rightarrow \mathbb{R}$  to be the restriction of  $\gamma$  to  $A^0 \subseteq A$ . Then for all  $k \in N$ , and  $s^0 \in A^0$  we have that:

$$u_k^0(s) - u_k^0(s_k^0; s_{-k}) = u_k(s) - u_k(s_k^0; s_{-k}) = \nabla_k \gamma(s) - \nabla_k \gamma(s_k^0; s_{-k}) = \nabla_k \gamma^0(s) - \nabla_k \gamma^0(s_k^0; s_{-k});$$

where the second equality follows from the definition of potential games (cf. section 2). Hence,  $\gamma^0$  is a potential function, and therefore  $\gamma^0$  a potential game.  $\square$

Using the recent results of Swenson et al. (2020), it is not difficult to show that perfectly-regular potential games are generic, and have a finite number of restricted equilibria. These are the statements of Lemma B.2 and Lemma B.3, respectively.

Lemma B.2. Almost all finite potential games are perfectly-regular.

Proof of Lemma B.2 Let  $\text{RPG}$  and  $\text{PRPG}$  denote the sets of regular potential, and perfectly-regular potential games, respectively. Let also  $\gamma$  a random finite potential game. Since  $\gamma$  is finite, there exist  $2^m$  distinct restrictions of  $\gamma$ , where  $m := \sum_{k \in N} |A_k|$ . Then, by Lemma B.1, we have that any restriction  $\gamma^0$  of  $\gamma$  is also a random finite potential game, and therefore  $\Pr(\gamma^0 \in \text{RPG}) = 1$ , with respect to the Lebesgue measure Swenson et al. (2020). It follows that:

$$\begin{aligned} \Pr(\gamma \in \text{PRPG}) &= 1 - \Pr(\gamma \notin \text{PRPG}) \\ &= 1 - \Pr\left(\bigcup_{\gamma^0} \gamma^0 \notin \text{RPG}\right) \\ &= 1 - \int_{\mathcal{R}} \Pr(\gamma^0 \notin \text{RPG}) \, d\mu; \end{aligned}$$

where the last equality follows from the fact that  $\mathcal{R}$  is finite.  $\square$

Lemma B.3. Every perfectly-regular finite potential game has a finite number of restricted equilibria.

Proof of Lemma B.3 Let  $\Gamma$  be a perfectly-regular finite potential game, and let  $\Gamma^0$  be one of its restrictions. By the definition of a perfectly-regular potential game, we have that  $\Gamma^0$  is a regular potential game. Furthermore, since  $\Gamma$  is finite and  $A_k^0 = A_k$  for all  $k \in N$ , it follows that  $\Gamma^0$  is finite. But then  $\Gamma^0$  is a finite regular potential game and as such it has a finite number of Nash equilibria, i.e.,  $|\text{NE}(\Gamma^0)| < \infty$  Swenson et al. (2020). Finally, since each restricted equilibrium is a Nash equilibrium of a restriction of  $\Gamma$ , it follows that there exist at most:

$$\sum_{\Gamma^0 \in \mathcal{R}} |\text{NE}(\Gamma^0)| \leq 2^m \max_{\Gamma^0 \in \mathcal{R}} |\text{NE}(\Gamma^0)| < \infty$$

restricted equilibria of  $\Gamma$ . Therefore, the number of restricted equilibria of  $\Gamma$  is finite.  $\square$

We now consider the replicator dynamics of a finite potential game given by the dynamical system of equations in equation QRD. Our goal for the remainder of this section is to prove that for any interior initial condition  $x(0) \in \text{int } X$  the dynamics in equation QRD converge pointwise to a Nash equilibrium of a given perfectly-regular finite potential game.

The proof proceeds in two parts: First, we prove that the dynamics converge to a restricted equilibrium of the game for any initial condition. Second, we prove that for any interior initial condition, the dynamics are bound to deviate from any rest-point that is not a Nash equilibrium, and therefore, they have to converge to a Nash equilibrium.

Let us begin by proving the first of the two claims. For this, we will use the notion of the  $\omega$ -limit set of a sequence  $(x(t))_{t \geq 0} \subset X$  that is generated by the QRD, which is defined as

$$\omega(x(t)) := \bigcap_{t \geq 0} \text{cl} \{x(t^0) \mid t^0 > t\}$$

where  $\text{cl } S$  denotes the closure of a set  $S$ .

Lemma B.4. Given a perfectly-regular finite potential game, every  $\omega$ -limit set, with respect to the replicator dynamics, is a singleton  $\{x^*\}$ , where  $x^* \in X$  is a rest-point of the dynamics. Specifically,  $x^*$  is a Nash equilibrium, if  $q = 0$ , or a restricted equilibrium, if  $q > 0$ . Furthermore, the set  $Q(X) := \{x \in X \mid \exists x_0 \in X \text{ s.t. } \lim_{t \rightarrow \infty} x(t) = x; x(0) = x_0\}$ , i.e., the set of all limit points, is finite.

Proof of Lemma B.4 Let  $\Gamma$  be a perfectly-regular finite potential game. Since  $\Gamma$  is a potential game, by Proposition 6.4 of Mertikopoulos & Sandholm (2018), we have that every  $\omega$ -limit set consists entirely of rest-points of the dynamics. In particular, these are Nash equilibria if  $q = 0$ , or restricted equilibria of  $\Gamma$ , if  $q > 0$ . However, since  $\Gamma$  is perfectly-regular—it suffices for it to be regular for the case  $q = 0$ —it follows by Lemma B.3 that every  $\omega$ -limit set is a finite set. Consider now, the  $\omega$ -limit set of an orbit  $(x(t))_{t \geq 0}$  of the dynamics for some arbitrary initial condition  $x(0) = x_0$ . Since  $x(t)$  is continuous, the  $\omega$ -limit set  $\omega(x(t))$  is the decreasing intersection of compact, connected sets and, therefore, it is connected. Since the  $\omega$ -limit set is finite, the above implies that, in fact it has to be a singleton  $\{x^*\}$ , where  $x^*$  is a Nash equilibrium if  $q = 0$  Mertikopoulos & Sandholm (2018), or a restricted equilibrium if  $q > 0$ , respectively. Finally, from the above, we have that the set of all limit points  $Q(X)$  is a subset of the restricted equilibria of  $\Gamma$  and, therefore, since  $\Gamma$  is a perfectly-regular finite potential game, we have, by Lemma B.3, that the set of restricted equilibria of  $\Gamma$  and, consequently,  $Q(X)$  are finite.  $\square$

To prove Theorem 3.2, it remains to exclude convergence to restricted equilibria that are not Nash equilibria of the original game,  $\Gamma$ . To establish that, we will show that as  $q \rightarrow 0$  in the QRD approach a limit point  $x_k$ , the probability  $x_{k,a_k}$  of non-optimally performing actions must go to zero. Thus, all actions  $a_k$  in  $\text{supp } x_k$  must be a best response against  $x_k$  for all agents  $k \in N$  which implies that  $x_k$  is a NE of  $\Gamma$ .

Proof of Theorem 3.2 If  $q = 0$ , the statement follows directly from Lemma B.4. So we only need to consider the replicator dynamics for  $q > 0$ . Let  $\Gamma$  be a perfectly-regular finite potential game, and let  $(x(t))_{t \geq 0}$  be a trajectory of the replicator dynamics with initial condition  $x(0) = x_0 \in \text{int } X$ .

<sup>1</sup>Recall that the interior of the set  $X$ ,  $\text{int } X$  is the set of all joint choice distributions  $x \in X$  with full support, i.e.,  $x_{k,a_k} > 0$  for all  $a_k \in A_k$  and for all  $k \in N$ ; all points of  $X$  that are not in the interior, are called boundary points

Since  $\epsilon > 0$ , we know that the support of  $\alpha(t)$  remains constant for all  $t \in \mathbb{R}$  (Mertikopoulos & Sandholm (2018)). Thus, since  $\alpha(t) \in \text{int } X$ , it follows that  $x(t)$  remains in the interior of  $X$  for all  $t \geq 0$ , i.e.,  $x_{k;a_k}(t) > 0$  for all  $k \in N$  and for all  $a_k \in A_k$ . Furthermore, by Lemma B.4, we have that the limit  $\lim_{t \rightarrow \infty} x(t)$  exists and is a rest-point of the dynamics, say

Assume, now, that  $x$  is not a Nash equilibrium. This implies that there exists a player  $i$  and a pure action  $a_i$  of  $i$  such that  $u_i(a_i; x) > u_i(x)$ . Moreover, since  $x$  is, by construction, a rest-point,  $\dot{x}_{i;a_j} = 0$  for all actions  $a_j$ . Observe that  $\dot{x}_{i;a_i} = 0$ . To see this, let us assume the absurd  $\dot{x}_{i;a_i} > 0$ . Since  $\dot{x}_{i;a_j} = 0$  for all  $a_j$ , by equation QRD, we have

$$u_i(a_j; x) = \frac{\prod_{a_k \in A_k} x_{i;a_k}^q u_i(a_k; x)}{\prod_{a_k \in A_k} x_{i;a_k}^q}; \text{ for all } a_j \text{ whenever } x_{i;a_j} > 0. \quad (3)$$

Furthermore, since  $u_i(a_i; x) > u_i(x)$ , we also have

$$\begin{aligned} u_i(a_i; x) &> u_i(x) \\ &= \int_X x_{i;a_j} u_i(a_j; x) \\ &= \int_{a_j \in A_i} x_{i;a_j} \frac{\prod_{a_k \in A_k} x_{i;a_k}^q u_i(a_k; x)}{\prod_{a_k \in A_k} x_{i;a_k}^q} \text{ by equation 3} \\ &= \frac{\prod_{a_k \in A_i} x_{i;a_k}^q u_i(a_k; x)}{\prod_{a_k \in A_i} x_{i;a_k}^q}. \end{aligned} \quad (4)$$

Clearly, the above reveals a contradiction, and therefore, it must be the case that  $\dot{x}_{i;a_i} = 0$ .

Fix  $\epsilon > 0$ , and consider the set:

$$B := \left\{ x \in X \mid u_i(a_i; x) > \frac{\prod_{a_j \in A_j} x_{i;a_j}^q u_i(a_j; x)}{\prod_{a_j \in A_j} x_{i;a_j}^q} + \epsilon \right\}$$

By continuity,  $B$  is open and by equation 4, it contains  $x$ —given  $\epsilon$  small enough. Since  $x(t)$  converges to  $x^*$  as  $t \rightarrow \infty$ , there exists a time  $t_0 \geq 0$  such that  $x(t) \in B$  for all  $t > t_0$ . Therefore, for each  $t > t_0$  we have that:

$$\dot{x}_{i;a_i}(t) = \left( x_{i;a_i}(t) \right)^q u_i(a_i; x) - \frac{\prod_{a_j \in A_j} (x_{i;a_j}(t))^q u_i(a_j; x)}{\prod_{a_j \in A_j} (x_{i;a_j}(t))^q} > \left( x_{i;a_i}(t) \right)^q > 0;$$

where the last inequality follows because  $x(t) \in \text{int } X$  for all  $t \geq 0$ . Finally, by integrating with respect to time, we have that for all  $t > t_0$ :

$$x_{i;a_i}(t) = \int_{t_0=t}^t \dot{x}_{i;a_i}(t) dt + x_{i;a_i}(t_0) > x_{i;a_i}(t_0) > 0;$$

Therefore, by the continuity of  $\alpha(t)$ , we have that  $\alpha_{i;a_i} = \lim_{t \rightarrow \infty} x_{i;a_i}(t) > 0$ , which is a contradiction to our assumption that  $\alpha_{i;a_i} = 0$ , which is a direct consequence of not being a NE of the game; thus,  $x^*$  has to be a Nash equilibrium of  $\Gamma$ .  $\square$

## C MISSING PROOFS AND MATERIALS: SECTION 4

For this part, it is useful to use the generic notation used by Pangallo et al. (2022), who provide a taxonomy of 2-player games. Consider an arbitrary 2-symmetric PRPG with payoff functions  $u_1(s_1; s_2) = u_2(s_2; s_1) = B_{s_1; s_2}$ , where the payoff matrix  $B \in \mathbb{R}^{2 \times 2}$  is given by:

$$B = \begin{pmatrix} a & c \\ b & d \end{pmatrix}; \quad (5)$$

and set without any loss of generality  $a > b$ —that may always be done by possibly re-indexing the agents' actions. If  $a > d$ , then the game is dominance-solvable (cf. Table 1 of Pangallo et al. (2022)) and as such the dynamics of the game are trivial. Hence, we may narrow our scope

to payoff matrices that satisfy  $c > a + b$ . Also, PRPG games have a finite number of Nash equilibria, and so, we may exclude games where the above inequality is strict. All in all, we are going to assume without any loss of the generality that the following conditions hold:

$$d > a > b \quad \text{and} \quad d > c: \tag{6}$$

Proof of Lemma 4.5 Recall that in the notation of Pangallo et al. (2022), an arbitrary symmetric PRPG can be written with payoff functions  $u_1(s_1; s_2) = u_2(s_2; s_1) = B_{s_1; s_2}$ , where the payoff

matrix  $B \in \mathbb{R}^2 \times \mathbb{R}^2$  is given by  $B = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$ , where we assumed without loss of generality that

$d > a > b$  and  $d > c$  (cf. equation equation 6. Let also  $x = (x; 1 - x)$  and  $y = (y; 1 - y)$  with  $x; y \in [0; 1]$  denote the choice distributions of the two agents, adopting the common notation for the statespace of game dynamics in 2 games. Thus, by slightly abusing notation, the choice distributions can be conveniently represented by single variable  $x$  and  $y$  for agents 1 and 2, respectively. The 3 Nash equilibria of such games are  $x = y = 1$  and  $x = y = 0$ , with payoffs  $a$  and  $d$ , respectively, for both players, and  $x = y = \frac{d - c}{a + b + d - c}$ , with payoff  $(x - y)^2 a + x(1 - x)(b + c) + (1 - x)^2 d$  for both players. Recall that, by definition, the Nash equilibrium  $x = y = 0$  is always payoff-dominant—due to the possible re-indexing of the actions—and risk-dominant if  $d - c > a - b$ .

To prove the claim of the Lemma, we begin by presenting the equations of motion of the replicator dynamics as functions of  $x$  and  $y$ . For the first agent that is:

$$\dot{x} = x^q u_1(a_1; y) - \frac{x^q u_1(a_1; y) + (1 - x)^q u_1(a_2; y)}{x^q + (1 - x)^q} \tag{7}$$

$$\begin{aligned} &= \frac{x^q(1 - x)^q}{x^q + (1 - x)^q} (u_1(a_1; y) - u_1(a_2; y)) \\ &= \frac{x^q(1 - x)^q}{x^q + (1 - x)^q} (ay + c(1 - y) - by - d(1 - y)) \end{aligned} \tag{8}$$

$$\begin{aligned} &= \frac{x^q(1 - x)^q}{x^q + (1 - x)^q} [(a - b + d - c)y - (d - c)] \\ &= \frac{x^q(1 - x)^q}{x^q + (1 - x)^q} (y - y^*); \end{aligned} \tag{9}$$

where  $y^* := \frac{d - c}{a - b + d - c}$ . Similarly, we may derive the equation of motions for the second agent as  $\dot{y} = \frac{y^q(1 - y)^q}{y^q + (1 - y)^q} (x - x^*)$ . Here,  $(x^*; y^*)$  is the mixed Nash equilibrium of the game (see above), and it holds that  $y^* = x^*$ . Thus, apart from the variables  $x$  and  $y$  and the hyperparameter  $q$  which is exogenously given, the replicator dynamics depend on the payoffs of the game only through  $\frac{d - c}{a - b}$ . It follows that any transformation that preserves the value of  $\frac{d - c}{a - b}$  and scales by a constant may only scale  $x$  and  $y$  by the same constant; that is, may only affect the convergence rate of the dynamics, but not their limiting behavior. Starting from an arbitrary payoff matrix  $B$  as given in equation 5 let us assume, without any loss of the generality, that the conditions in equation 6 apply. Notice that, since by the aforementioned assumptions we have  $d - c > a - b$  we may set some  $2 \in \mathbb{R}$  such that  $0 < a + 2 < d - c$ . Accordingly, we consider the following sequence of transformations: (T1) Add  $2$  to the first column, (T2) subtract  $2$  from the second column, and (T3) divide by  $a + 2$ . These lead to:

$$B = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \xrightarrow{(T1)} \begin{pmatrix} a + 2 & c \\ b + 2 & d \end{pmatrix} \xrightarrow{(T2)} \begin{pmatrix} a + 2 & 0 \\ b + 2 & d - c \end{pmatrix} \xrightarrow{(T3)} \begin{pmatrix} 1 & 0 \\ \frac{b + 2}{a + 2} & \frac{d - c}{a + 2} \end{pmatrix} =: A_{\frac{d - c}{a + 2}, \frac{b + 2}{a + 2}};$$

Notice that  $A_{\frac{d - c}{a + 2}, \frac{b + 2}{a + 2}}$  is the payoff matrix of a parametric game  $\tilde{g}$ , where  $\tilde{a} := \frac{d - c}{a + 2}$  and  $\tilde{b} := \frac{b + 2}{a + 2}$ .

Observe, that (T1), (T2) and (T3) leave  $\frac{d - c}{a - b}$  unaltered and only scale by a constant  $\frac{1}{a + 2}$ ; that is, the limiting behavior of the replicator dynamics is preserved. Furthermore, the risk-dominance of the equilibrium points is preserved, because:

$$d - c > a - b \quad \text{if and only if} \quad \frac{d - c}{a + 2} > 1 - \frac{b + 2}{a + 2}$$

Finally, the payoff-dominance of the Nash equilibrium  $x = y = 0$  is also preserved because, by the definition of  $\tilde{g}$ , we have that  $\frac{d - c}{a + 2} > 1 - \frac{b + 2}{a + 2}$ .  $\square$

Next, we are going to construct the invariant functions of  $\dot{x}$  with respect to the  $q$ -replicator dynamics, which are given by Lemma 4.7 that we restate below for completeness. Recall that  $\bar{x}$  is defined to be the equilibrium point of the game  $G$ ; that is,  $x = y = \bar{x}$ .

**Lemma 4.7 (Invariant functions of QRD in 2 × 2 symmetric PRPGs).** Given a 2 × 2 symmetric PRPG  $G_w$ , whose agents evolve with respect to the replicator dynamics QRD, the separable function  $\phi_q : (0; 1)^2 \rightarrow \mathbb{R}$ , with  $\phi_q(x; y) := \phi_q(x) + \phi_q(y)$ , and  $\phi_q : (0; 1) \rightarrow \mathbb{R}$  given by

$$\phi_q(x) = \begin{cases} \frac{x^{2-q} + (1-x)^{2-q}}{2-q} + \frac{1}{1-q} \frac{x^{1-q} - (1-x)^{1-q}}{1-x}; & q \notin \{1; 2\}; \\ \ln(x) + (1-x)\ln(1-x); & q = 1; \\ \ln(x) + \ln(1-x) + \frac{1}{x} + \frac{1}{1-x}; & q = 2; \end{cases} \quad (2)$$

remains constant along any trajectory  $(x(t); y(t))_{t \geq 0}$  of the system. The function  $\phi_q(x)$  is continuous with respect to the parameter  $q$  at both  $q = 1$  and  $q = 2$ , since  $\lim_{q \rightarrow 1} \phi_q(x) = \phi_1(x)$  and  $\lim_{q \rightarrow 2} \phi_q(x) = \phi_2(x)$  for all  $x \in (0; 1)$ .

**Proof of Lemma 4.7.** To prove the statement, we will show that the time derivative of the function  $\phi_q(x(t); y(t))$  is equal to zero. Let us begin by constructing the derivative  $\dot{\phi}_q$ . For  $q \notin \{1; 2\}$  we have that:

$$\dot{\phi}_q(x) = \frac{x}{x^q} - \frac{1-x}{(1-x)^q} - \frac{1}{x^q} + \frac{1}{(1-x)^q} = \frac{x}{x^q} + \frac{x}{(1-x)^q} = \frac{(x - (1-x)^q + x^q)}{x^q(1-x)^q}.$$

Similarly, for  $q = 1$  we have that:

$$\dot{\phi}_1(x) = \frac{1}{x} - \frac{1}{1-x} = \frac{x}{x(1-x)};$$

and, for  $q = 2$ , we have that:

$$\dot{\phi}_2(x) = \frac{1}{x} - \frac{1}{1-x} - \frac{1}{x^2} + \frac{1}{(1-x)^2} = \frac{(x - (1-x)^2 + x^2 - (1-x)^2)}{x^2(1-x)^2}.$$

That is the derivative of  $\phi_q(x)$  has the general form:

$$\dot{\phi}_q(x) = \frac{(x - (1-x)^q + x^q)}{x^q(1-x)^q},$$

for all  $q \geq 0$ , where  $2 \leq q < 1q$ . Notice, that the choice for  $\phi_q$  is purely stylistic because the invariance of a function is not affected by scalar transformations. Using equation 7 from the proof of Lemma 4.5 we have that:

$$\begin{aligned} -\dot{\phi}_q(x; y) &= \frac{\partial \phi_q(x; y)}{\partial x} \dot{x} - \frac{\partial \phi_q(x; y)}{\partial y} \dot{y} \\ &= \dot{\phi}_q(x) \dot{x} - \dot{\phi}_q(y) \dot{y} \\ &= [(\dot{x} - \dot{y}) - (y - x) \dot{y} + (x - y) \dot{x}]: \quad \square \end{aligned}$$

To proceed with the proof of Theorem 4.6, we first need to provide the formal definition of the stable and unstable manifolds of the mixed NE  $\bar{x}$ .

**Definition C.1 (Stable and unstable manifolds of  $\bar{x}$  under QRD).** Let  $\phi_q : [0; 1]^2 \rightarrow \mathbb{R}$  with  $\phi_q(x; y) = \phi_{q; \text{Stable}}(x; y) - \phi_q(x) + \phi_q(y)$  for all  $x; y \in [0; 1]$  denote the invariant function of the replicator dynamics for the 2 × 2 symmetric PRPG  $G_w$ . The unstable manifold of the mixed NE  $\bar{x}; y = (\bar{x}; \bar{y})$  under the  $q$ -replicator dynamics is the curve  $M_{\text{Unstable}} := \{f(x; y) \in (0; 1)^2 \mid x = \bar{x}; y \geq \bar{y}\}$ ; that is, the set of points for which  $\lim_{t \rightarrow 1} x(t) = \bar{x}$  and  $\lim_{t \rightarrow 1} y(t) = \bar{y}$ . Analogously, the stable manifold of the mixed NE is the curve  $M_{\text{Stable}} := \{f(x; y) \in (0; 1)^2 \mid \phi_{q; \text{Stable}}(x; y) = 0; y \leq \bar{y}\}$ ; that is, the set of points for which  $\lim_{t \rightarrow 1} x(t) = \bar{x}$  and  $\lim_{t \rightarrow 1} y(t) = \bar{y}$ .

We are now ready to prove Theorem 4.6.

Proof of Theorem 4.6. To establish the claim, we will need to have at hand an explicit form of the stable manifold  $M_{\text{Stable}}$  of the mixed NE of  $w_i$  with respect to the replicator dynamics, and  $\mathcal{D}$ . In that regard, we are now going to construct that manifold. Recall that from equation 2, the invariant function of any QRD in  $w_i$  is given by  $\phi_0(x; y) := \phi_0(x) - \phi_0(y) = 0$ , where:

$$\phi_0(x) = \frac{x^2 + (1-x)^2}{2} + 1-x = x^2 - 2x + 1$$

Therefore,  $\phi_0(x; y) = x^2 - 2x + 1 - (y^2 - 2y + 1) = (x-y)(x+y-2) = 0$  and the stable manifold (cf. subsection 4.2) satisfies  $\phi_{\text{q;Stable}}(x; y) = x + y - 2 = 0$ . In other words, the stable manifold is the line segment:

$$M_{\text{Stable}} = \{(x; y) \in (0; 1)^2 \mid y = 2 - x\} \tag{10}$$

By equation 10 and Definition C.1, the stable and unstable manifolds of gradient descent are given by the lines  $y = 2 - x$  and  $y = x$ , respectively, where  $\alpha = \frac{w}{w+1}$ , while the stable and unstable manifolds of the standard replicator dynamics are given as solutions to:

$$\phi_1(x; y) = \phi_1(x) - \phi_1(y) = \ln \frac{x}{y} + (1-x) \ln \frac{1-x}{1-y} = 0;$$

where the line  $y = x$  corresponds to a solution. We are going to prove that the remaining solution of the previous equation, although it cannot be expressed explicitly, satisfies  $x < \frac{1}{2}, y > \frac{1}{2}$  if  $x < \frac{1}{2}$  (with equality in both cases only if  $x = \frac{1}{2}$ ), and  $y < \frac{1}{2}, x > \frac{1}{2}$  otherwise; hence, the statements of Theorem 4.6 follow naturally.

It is not difficult to show that  $\phi_1'(x) = \frac{x}{x(1-x)}$  (cf. proof of Lemma 4.7), and  $\phi_1''(x) = \frac{x^2 - 2x + 1}{x^2(1-x)^2} < 0$ . Therefore,  $\phi_1$  is a strictly concave function with maximum at  $x = \frac{1}{2}$ . To proceed, it will be useful to define the implicit function  $\gamma : (0; 1) \rightarrow (0; 1)$  such that  $\phi_1(\gamma) = \phi_1(x)$ , and  $x < \frac{1}{2} \iff \gamma > \frac{1}{2}$ , and  $x > \frac{1}{2} \iff \gamma < \frac{1}{2}$ . By applying the Intermediate Value Theorem on  $\phi_1$  in the intervals  $(0; \frac{1}{2})$ , and  $(\frac{1}{2}; 1)$ , we can verify that  $\gamma$  is a well-defined bijective function. Note that  $y = \gamma(x)$  has to correspond to the remaining solution of  $\phi_1(x; y) = 0$ .

Without any loss of generality, let us consider the case of  $x < \frac{1}{2}$ . Since  $\phi_1$  is strictly decreasing in  $(0; \frac{1}{2})$  ( $\phi_1'' < 0$ ), we have that  $\phi_1'(x) > 0$  for all  $x \in (0; \frac{1}{2})$ , and  $\phi_1'(x) < 0$  for all  $x \in (\frac{1}{2}; 1)$ . We begin by proving that for all  $x \in (0; \frac{1}{2})$  it holds that  $|\phi_1'(x)| < |\phi_1'(\gamma(x))|$ , i.e.,  $\phi_1'(x) < \phi_1'(\gamma(x))$ . Specially, we have the following series of equivalences:

$$\phi_1'(x) < \phi_1'(\gamma(x)) \iff \frac{x}{x(1-x)} < \frac{x}{\gamma(x)(1-\gamma(x))}$$

Since  $x < \frac{1}{2}$  and  $x \in (0; \frac{1}{2})$ , the latter is equivalent to  $x^2(1-\gamma(x)) < 0$  or equivalently to  $\gamma(x) > 1-x$  which holds by assumption. Next, by taking advantage of the above, we can prove that, for all  $x \in (0; \frac{1}{2})$ , it holds  $\phi_1(x) > \phi_1(\gamma(x))$ ; that is:

$$\begin{aligned} \phi_1(x) &= \int_x^0 \phi_1'(t) dt = \int_x^0 \phi_1'(t) dt + \int_0^x \phi_1'(t) dt \\ &> \int_0^0 \phi_1'(t) dt + \int_0^x \phi_1'(t) dt = \phi_1(x) \end{aligned}$$

Let us, now, consider some  $x \in (0; \frac{1}{2})$ . There are three possibilities: 1.  $x < \frac{1}{2}$ ; 2.  $x \in (\frac{1}{2}; 1)$ ; and 3.  $x = \frac{1}{2}$ . If  $x < \frac{1}{2}$ , then, since  $\gamma(x) \in (\frac{1}{2}; 1)$ , it follows, trivially, that  $x + \gamma(x) < 2$ .

If  $x \in (\frac{1}{2}; 1)$ , let us set  $\gamma = 2 - x$ . Then, since  $\phi_1$  is strictly convex, we have that  $\phi_1(x) = \phi_1(\gamma) > \phi_1(\frac{x+\gamma}{2})$ . Consequently, by the Intermediate Value Theorem, there must exist some  $y \in (\frac{x+\gamma}{2}; \frac{1}{2})$  such that  $\phi_1(x) = \phi_1(y)$ . Clearly  $y < \frac{1}{2}$ , and, therefore, since  $\phi_1(x) = \phi_1(y)$ , it must be the case that  $\gamma(x) = y$ . Furthermore,  $x + \gamma(x) = x + y < x + \frac{x+\gamma}{2} = 2$ , and, therefore, our claim holds in that case as well.

Finally, if  $x = \frac{1}{2}$ , let us set  $\gamma = x$ . Since  $\phi_1' > 0$ , we have that  $\gamma \in (0; \frac{1}{2})$ , and therefore,  $\phi_1(x) < \phi_1(\gamma) = \phi_1(x)$ . Furthermore, by the monotonicity of  $\phi_1$  in the intervals  $(0; \frac{1}{2})$ ,  $(\frac{1}{2}; 1)$  and  $(x; 1)$ , we also have, respectively, that

$$\phi_1(y) < \phi_1(\gamma) < \phi_1(x) \text{ for all } y \in (0; \frac{1}{2}) \tag{11}$$

$$\phi_1(y) < \phi_1(x) \text{ for all } y \in (\frac{1}{2}; x) \tag{12}$$

$$\phi_1(y) > \phi_1(x) \text{ for all } y \in (x; 1) \tag{13}$$

In other words, it has to be the case that  $y(x) < 2 - x$ , and, therefore,  $x + y(x) < 2$ .

It follows that, if  $\alpha > 0.5$ ,  $y(x) < 2 - x$ . We remark that the case of  $\alpha = \frac{1}{2}$  follows identical arguments, while the case  $\alpha < \frac{1}{2}$  is trivial.  $\square$

**Remark C.2 (Technical intuition of Theorem 4.6)** It is important to provide the mathematical (technical) interpretation of Theorem 4.6. From equation (10), we observe that the speed of the projection dynamics ( $q = 0$ ), does not depend on the current state, since the term is eliminated from the right hand side. As a result, the separatrix is a straight line, see e.g., panel 1 in Figure 1. As  $q$  increases for values between 0 and 1 for which our analysis applies, the dynamics leave the boundary must faster, cf. panels 2 to 4 in Figure 1, while being fixed between the three points, i.e., the corners  $(0; 1)$ ;  $(1; 0)$  and the mixed NE. This implies a certain curvature for the separatrix. On the other hand, the center of the simplex, i.e.,  $(0.5; 0.5)$  point, belongs to the region of attraction of the risk-dominant equilibrium under all dynamics. Combining these two, we get that the replicator dynamics create a convex separatrix when risk- and payoff-dominant equilibria differ, and a concave separatrix when risk- and payoff-dominant equilibria coincide, cf. Appendix D. This provides the picture that we have, with the straight-line separatrix of the gradient descent (projection) dynamics being below the curved separatrix of the replicator when risk- and payoff-dominant equilibria differ and above when risk- and payoff-dominant equilibria coincide. The above explanation drives our current proof of Theorem 4.6. Deriving a game-theoretic interpretation (along with the technical one above) is an interesting open question.

Finally, a direct (technical) implication from the proof of Theorem 4.6 is provided in Lemma C.3 which may be of independent interest. Recall from the proof of Theorem 4.6 that the solutions to  $\phi_1(x; y) = 0$  are the functions  $y = x$  and  $y = (0; 1) - (0; 1)x$  such that  $y(x) = 1 - x$ , and for all  $x \in (0; 1) \cap \text{int } G$  it holds that  $y(x) \in G$ , and  $\phi_1(y(x)) = \phi_1(x)$ .

**Lemma C.3 (Curvature of the stable manifold of RD).** Consider the  $\phi$ -replicator dynamics (RD) in the parametric game  $w_\alpha$ . The stable manifold  $M^{\text{Stable}}$  of RD in  $w_\alpha$  is given by the curve  $y = y(x)$ . If the payoff-dominant equilibrium  $x = y = 0$ , is also risk-dominant, then  $y$  is strictly concave. Conversely, if the non-payoff dominant equilibrium  $x = y = 1$ , is risk-dominant, then  $y$  is strictly convex. Otherwise,  $y(x) = 1 - x$ .

**Proof of Lemma C.3** By differentiating both sides of the implicit function  $\phi_1(y(x)) = \phi_1(x)$  with respect to  $x$ , we get that  $\phi_1'(y(x))y'(x) = \phi_1'(x)$ , i.e.,  $y'(x) = \frac{\phi_1'(x)}{\phi_1'(y(x))}$ . Notice that, since  $\phi_1$  is bijective,  $\phi_1'$  is monotonic ( $\phi_1'' < 0$ ), and  $\phi_1'(0) = \phi_1'(1) = 0$ , the above equality is well-defined for all  $x \in (0; 1) \cap \text{int } G$ . Hence, we have that:

$$\begin{aligned} y''(x) &= \frac{\phi_1''(x)[\phi_1'(y)]^2 - \phi_1'(x)\phi_1''(y)y'}{[\phi_1'(y)]^3} = \frac{\phi_1''(x)[\phi_1'(y)]^2 - [\phi_1'(x)]^2 \phi_1''(y)}{[\phi_1'(y)]^3} \\ &= \frac{y^3(1-y)^3(1-x)(x-y)(y+x-2)}{(1-y)^3x^2(1-x)^2y^2(1-y)^2}, \end{aligned}$$

where the dependency of  $y$  to  $x$  is implied for compactness. Hence, we have  $y''(x) < 0$  if and only if  $\frac{(x-y)(y+x-2)}{y} < 0$ . However, by the Intermediate Value Theorem (IVT) applied on  $\phi_1$  in  $(0; \frac{1}{2})$ , and  $(\frac{1}{2}; 1)$ , and the definition of  $y$ , it follows, trivially, that  $x > y$  if, and only if,  $x > \frac{1}{2}$ , with equality in both inequalities only if  $x = \frac{1}{2}$ . Therefore,  $\frac{x-y}{y} > 0$ ;  $\forall x \in (0; 1) \cap \text{int } G$ ; hence  $y''(x) < 0$  if, and only if,  $y + x - 2 < 0$ , which by the proof of Theorem 4.6 is equivalent to  $x > \frac{1}{2}$ . This concludes the proof for the first statement. The second statement follows in a similar manner by requesting  $y''(x) > 0$ , while the last statement is trivial.  $\square$

**Proof of Theorem 4.8** Let  $w_\alpha$  be a 2x2 symmetric PRPG where the payoff-dominant equilibrium,  $x = y = 0$ , is also risk-dominant, i.e.,  $\alpha > \frac{1}{2}$ , or equivalently  $\alpha > 0.5$ , where  $x = y = \frac{1}{2}$  is the mixed NE of  $w_\alpha$ . Recall that, by equation 10, the stable manifold of the mixed NE of  $w_\alpha$ , with respect to GD is the line segment  $y = 2 - x$ , for  $x \in (\max\{0, 2 - 1g; \min\{2, 1g\})$ . Since  $\alpha > 0.5$ , we have that  $2 - 1 > 0$  and  $2 < 1$ ; therefore, the extreme points of are  $(0; 2 - 1)$  and  $(2 - 1; 0)$ . That implies that the RoA of  $(1; 1)$  is the triangle with extreme points  $(1; 1)$ ,  $(2 - 1; 1)$ , and  $(1; 2 - 1)$ . Since that is a right triangle, with both its base and its height equal

Figure 5: The stable manifolds,  $q; \text{Stable}(x; y) = 0$ , (solid blue lines) for the same values of  $w$  and the same instance of  $w$ , as in Figure 1, in which the payoff- and risk-dominant NE is at the bottom left corner. For all  $q$ , the separatrix goes through the mixed NE at the intersection of the (dashed red) and (dashed black) coordinates. All panels also include the unstable manifold defined by  $q = 0$  (dashed blue line). The region of attraction of the payoff-dominant NE is larger for all values of  $q$ ; however, this is because this NE is also risk-dominant, cf. Theorem 4.6.

to  $2(1 - \frac{1}{w+1})$ , the Lebesgue measure  $\text{RoA}(0; 0)$  is  $(\text{RoA}(1; 1)) = 2(1 - \frac{1}{w+1})^2$ ; subsequently,  $(\text{RoA}(0; 0)) = 1 - (\text{RoA}(1; 1)) = 1 - 2(1 - \frac{1}{w+1})^2$ . We may calculate the APoA of GD in  $w$ , as a function of  $w$  and  $q$ . Specifically, when  $q > 0.5$ , i.e.,  $q < 1 - \frac{1}{w}$ , we have that:

$$\begin{aligned} \text{APoA}(w; q) &:= \text{APoA}(\text{GD}; w; q) = \frac{\max_{x, y \in [0; 1]} \text{SW}(x; y)}{\text{APM}_{\text{SW}; [0; 1]^2}(\text{GD}; w; q)} \\ &= \frac{\text{SW}(0; 0)}{\text{SW}(0; 0) (\text{RoA}(0; 0)) + \text{SW}(1; 1) (\text{RoA}(1; 1))} \\ &= \frac{w}{w (\text{RoA}(0; 0)) + (\text{RoA}(1; 1))} \\ &= \frac{w}{w [1 - 2(1 - \frac{1}{w+1})^2] + 2(1 - \frac{1}{w+1})^2} = \frac{w(w+1)^2}{w(w+1)^2 - 2(w-1)(1 - \frac{1}{w+1})^2} \end{aligned}$$

where we used that  $1 - \frac{1}{w+1} = \frac{w}{w+1}$ . We may, now, perform a first-order analysis of  $\text{APoA}(w; q)$ ; that is, for all  $q < 1 - \frac{1}{w}$ , we have that:

$$\frac{\partial \text{APoA}(w; q)}{\partial q} = \frac{4w^2(w-1)(w+1)(1 - \frac{1}{w+1})}{[w(w+1)^2 - 2(w-1)(1 - \frac{1}{w+1})^2]^2} > 0$$

From the above, it follows that  $\text{APoA}(w; q) < \text{APoA}(w; 1 - \frac{1}{w})$ ; that is:

$$\text{APoA}(w; q) < \frac{4w^3}{4w^3 - 2(w-1)w^2} = \frac{2w^3}{w^3 + w^2} < 2;$$

where the last inequality follows by letting  $w \rightarrow \infty$ . Notice that this bound is tight. It is not difficult to see that if  $q < 0.5$ ,  $\text{APoA}(w; q)$  is unbounded.  $\square$

## D VISUALIZATIONS: INVARIANT FUNCTIONS AND SEPARATING MANIFOLDS

In this part, we provide systematic, and essentially exhaustive, visualizations of the stable manifolds (separatrices) in the  $w$ -class.

Same payoff- and risk-dominant NE in  $w$ . As mentioned in the main part, the separatrix for different values of  $q$  can be obtained by plotting the level set of the invariant functions in Figure 1. These are depicted in Figure 5. As a sanity check, we also see from Figure 5 that the region of attraction of the payoff-dominant equilibrium  $q = 0$  (GD dynamics) is larger than the region of attraction for  $q = 1$  (RD).

Stable manifolds for all  $q \geq 0$ . In a similar vein to Figure 2, we next depict the separatrices, stacked for all values of  $q \in [0; 10]$ , for different parameterizations of the  $w$ -class (Figure 6. In all panels of Figure 6, parameters are equal to 2. We obtain qualitatively equivalent plots for any  $w > 1$  and  $q$  small enough. The main takeaways from the (essentially exhaustive) visualizations in the panels of Figure 6 are that (i) the region of attraction of the risk-dominant equilibrium is


Figure 6: Stable manifolds (separatrices) for all different values of  $q$  of  $[0; 10]$  (from blue to brown) in different parameterizations of the  $w_i$  game for  $w = 2$  and varying  $\alpha$ . In all panels, the manifolds for  $\alpha = 0$ ,  $q = 1$ , and  $q = 2$  are shown in shades of black for reference. The region of attraction of the payoff-dominant equilibrium (bottom-left corner) shrinks and increases when this equilibrium is also risk-dominant ( $\alpha > 0.5$ ) and increases with  $q$  when this equilibrium is not risk-dominant ( $\alpha < 0.5$ ). In fact, as  $\alpha$  decreases, the payoff-dominant becomes increasingly more “risky” and its region of attraction becomes arbitrarily small.

larger for all  $q > 0$  regardless of whether this equilibrium is payoff-dominant or not, (ii) the region of attraction of the payoff-dominant equilibrium may become arbitrarily small as this equilibrium becomes arbitrarily risky. In particular, observation (ii), suggests that in this case, it is hopeless to bound any static or average performance measure. This became more transparent with the analysis in the previous Section of the Appendix (cf. Theorem 4.8 in the main paper). We conclude this part with some visualizations of the stable and unstable manifolds in a non-symmetric PRPG.


non-symmetric PRPGs. Consider the identical-interest PRPG  $ID_{w_i}$ , with identical payoff functions  $u_{w_i; 1}(s_1; s_2) = u_{w_i; 2}(s_1; s_2) = A_{w_i; 1; s_2}$ , where the payoff matrix  $A_{w_i} \in \mathbb{R}^{2 \times 2}$  is given by  $A_{w_i} = \begin{pmatrix} 1 & 0 \\ w & w \end{pmatrix}$ ,  $1 < w$ . The game  $ID_{w_i}$  has the same pure NE as the games  $w_i$ , namely  $x = y = 0$ , with payoff  $w$ , and  $x = y = 1$ , with payoffs  $1$  for both players, but now the mixed NE is not symmetric and it is given by  $x = (w - 1) / (w + 1)$  and  $y = w / (w + 1)$ .

In Figure 7, we visualize the stable and unstable manifolds for all values of  $q$  of  $[0; 10]$  in an instance of  $ID_{w_i}$  with  $w = 2$  and  $\alpha = 2$ . In this case, the separating (stable) manifolds do not increase (decrease) monotonically with  $q$  as it is evident from the overlapping (equally) colored regions. Thus, it requires a different approach to estimate whether the size of the regions of attraction of the payoff-dominant equilibrium follow a certain monotonicity pattern, which again may change depending on whether this equilibrium is also risk-dominant or not. In the context of the current paper, Figure 7 highlights that (i) the geometry of the regions of attraction is highly complex under different algorithms (parametrizations QRD) even for low-dimensional, identical interest games, and (ii) given this complexity, the findings in the class  $w_i$  become even more surprising and intriguing.

Different payoff- and risk-dominant NE in  $w_i$ . The main differences in the class  $w_i$  occur between games in which the payoff- and risk-dominant equilibria coincide and games in which they differ. Recall that Figure 1 shows the invariant function in an instance, where  $w = 2$  and  $\alpha = 0$ ,



**Figure 8:** The invariant function,  $q(X; y)$ , for all  $x, y \in [0; 1]^2$  in the game  $w_i$  for  $w = 2, 4$ , and the same values of  $q$  as in Figure 1:  $q = 0$  (gradient descent dynamics),  $q = 1$  (standard replicator dynamics),  $q = 2$  (log-barrier dynamics), and  $q = 20$ . The invariant function again becomes very steep at the boundary as  $q$  increases, taking both arbitrarily large negative (**dark**) and positive (**light**) values in the vicinity of the NE.



**Figure 9:** The stable manifolds,  $q; \text{Stable}(X; y) = 0$ , (solid blue lines) for the same values of  $q$  and instance of  $w_i$  as in Figure 8, in which the payoff-dominant NE is at the bottom-left corner and the risk-dominant NE is at the upper-right corner. For all  $q$ , the separatrix goes through the mixed NE at the intersection of the  $x$  (dashed red) and  $y$  (dashed black) coordinates. All panels also include the unstable manifold defined by  $x = y = 0$  (dashed blue line). The region of attraction of the payoff-dominant NE is now smaller for all  $q$ , because this NE is not risk-dominant, cf. Theorem 4.6.

i.e., in which the payoff- and risk- dominant equilibria coincide. In Figure 8, we provide an instance in which the payoff- and risk-dominant equilibria differ. Similar to Figure 5, Figure 9 depicts the separating manifolds (stable manifolds or separatrices) of the regions of attractions of the two pure NE. These are precisely the zero-level sets of the invariant functions shown in Figure 8. As we can see, in this case, the region of attraction of the payoff-dominant equilibrium,  $w$  is smaller than the region of the, now, risk-dominant equilibrium,  $1$ . Intuitively, when a NE becomes risk-dominated, its region of attraction shrinks, even if this NE is payoff-dominant. This is because, for a mixed choice of distributions, the risk-dominant NE yields a higher utility and is that more “attractive” for the dynamics. This trade-off between high reward at a certain state (e.g.,  $w; w$ ) and high risk if that state is not reached (e.g.,  $; 0$ , with  $< 0$ ), also explains why socially optimal, but otherwise risky, outcomes, e.g., the adoption of revolutionizing technology or a social norm that challenges the status quo, are never reached in real-life situations.

**On risk-dominance in higher dimensions** In diagonal games (D-PRPGs), it is straightforward to make an equilibrium more *risky*. This is achieved simply by replacing the zero entries in the corresponding line of matrix  $U$  by some negative number, e.g.,

$$U_{\text{risky};m} = \begin{pmatrix} u_1 & 0 & \dots & 0 \\ 0 & u_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -10 & -10 & \dots & u_m \end{pmatrix} \quad A := \text{diag}(u_1; u_2; \dots; u_m; \text{risk}_n = -10)$$

In this case, the payoff-dominant equilibrium with payoff  $u_m$  to each player becomes more risky, since a failure to coordinate on it results to a negative payoff of  $-10$  for the agent who selected the corresponding action. Proceeding in a similar fashion, one may replace the zero entries with an (arbitrarily large) negative element in all lines of the matrix except for the first one. In analogy to the  $U_{\text{risky};m}$  notation, we will denote such games by  $U_{\text{risky}} := \text{diag}(u_1; u_2; \dots; u_m; \text{risk} = -r)$ , where  $r > 0$  is the *risk constant* (equal to  $-10$  in the example above). In this way, all equilibria become more risky except for the payoff-dominated one, i.e., the equilibrium with payoffs  $u_1$  to each agent which corresponds to the action profile in which every agent selects their first action.<sup>2</sup> Concerning

<sup>2</sup>An alternative interesting approach to generalize the notion of risk-dominance in arbitrary games is via pairwise comparisons of actions as proposed by Honda (2012).

