



Best-response dynamics in two-person random games with correlated payoffs

Hlafo Alfie Mimun^{a,*}, Matteo Quattropani^b, Marco Scarsini^a

^a Dipartimento di Economia e Finanza, Luiss University, Viale Romania 32, Roma, 00197, RM, Italy

^b Dipartimento di Matematica “Guido Castelnuovo”, Sapienza Università di Roma, Piazzale Aldo Moro 5, Roma, 00185, RM, Italy

ARTICLE INFO

Keywords:

Pure Nash equilibrium
Random games
Potential games
Best response dynamics

ABSTRACT

We consider finite two-player normal form games with random payoffs. Player A's payoffs are i.i.d. from a uniform distribution. Given $p \in [0, 1]$, for any action profile, player B's payoff coincides with player A's payoff with probability p and is i.i.d. from the same uniform distribution with probability $1 - p$. This model interpolates the model of i.i.d. random payoff used in most of the literature and the model of random potential games. First we study the number of pure Nash equilibria in the above class of games. Then we show that, for any positive p , asymptotically in the number of available actions, best response dynamics reaches a pure Nash equilibrium with high probability.

1. Introduction

1.1. The problem

Consider the class of two-person normal-form finite games. Some properties hold for the entire class, for instance, the mixed extension of each game in the class admits a Nash equilibrium (Nash, 1950, 1951). Some properties hold generically, for instance, generically the number of Nash equilibria is finite and odd (Wilson, 1971; Harsanyi, 1973). Some properties do not hold generically and neither does their negation; for instance having a pure Nash equilibrium or not having a pure Nash equilibrium is not a generic property of finite games. Still, it may be relevant to know how likely it is for a finite game to admit a pure equilibrium. Along a similar line of investigation, how likely is a recursive procedure—such as best response dynamics—to reach a pure Nash equilibrium in finite time?

One way to formalize these questions is to assume that the game is drawn at random according to some probability measure. It is not clear what a natural probability measure is in this setting; a good part of the literature on the topic has focused on measures that make the payoffs i.i.d. with zero probability of ties. Few papers have relaxed this assumption. For instance, Rinott and Scarsini (2000) considered payoff vectors that are i.i.d. across different action profiles, but can have some positive or negative dependence within the same action profile. Amiet et al. (2021b) considered i.i.d. payoffs whose distribution may have atoms and, as a consequence, may produce ties. Durand and Gaujal (2016) studied the class of random potential games, i.e., a class of games that admit a potential having i.i.d. entries.

* Corresponding author.

E-mail address: hmimun@luiss.it (H.A. Mimun).

1.2. Our contribution

In this paper we want to study two-person games with random payoffs where the stochastic model for the payoffs parametrically interpolates the case of i.i.d. payoffs with no ties and the case of random potential games. In particular, we start with a model where all payoffs are i.i.d. according to a continuous distribution function (without loss of generality, uniform on $[0, 1]$) and we consider an i.i.d. set of coin tosses, one for each action profile. If the toss gives head, then the original payoff of the second player is made equal to the payoff of the first player; if the toss gives tail, the payoff remains unchanged. The relevant parameter is the probability p of getting heads in the coin toss. If $p = 0$, we obtain the classical model of random games with continuous i.i.d. payoffs. If $p = 1$, we get the model of common-interest random games. From the viewpoint of pure Nash equilibria (PNE) any potential game is strategically equivalent to a common-interest game. Therefore, the above class of games parametrically interpolates the case of i.i.d. payoffs with no ties and the case of random potential games. When p is small, the game is close to a game with i.i.d. payoffs; when p is large, the game is close to a potential game.

For this parametric class of games we first compute the expected number of PNE as a function of p , and then study its asymptotic behavior as the numbers of actions of the two players diverge, possibly at different speeds. It is well known (Powers, 1990) that, as the number of action increases, the asymptotic distribution of the number of PNE is a Poisson(1) distribution, for i.i.d. random payoffs. Our result shows an interesting phase transition around $p = 0$, in the sense that for every $p > 0$ the expected number of PNE diverges.

We then consider best response dynamics (BRD) for the above class of games. Durand and Gaujal (2016) considered BRD for random potential games with an arbitrary number of players and the same number of actions for each player. In this class of games a PNE is reached by a BRD in finite time. Durand and Gaujal (2016) studied the asymptotic behavior of the expectation of this random time. In our paper we first consider potential games and we compute the distribution of the time that the BRD needs to reach a PNE. Moreover we compute exactly the first two moments of this random time, when the two players have the same action set.

Amiet et al. (2021a) showed that, for games with i.i.d. continuous payoffs, when players have the same action set, as the number of actions increases, the probability that a BRD reaches a PNE goes to zero. Here we generalize the result of Amiet et al. (2021a) to the case of possibly different action sets for the two players. Moreover, we prove that for every positive p , asymptotically in the number of actions, a BRD reaches a PNE in finite time with probability arbitrarily close to 1. Again this shows a phase transition in $p = 0$ for the behavior of the BRD.

1.3. Related literature

Games with random payoffs have been studied for more than sixty years. We refer the reader to Amiet et al. (2021b); Heinrich et al. (2023) for an extensive survey of the literature on the topic. Here we mention just some recent papers and some articles that are more directly connected with the results of our paper. Powers (1990) proved that in random games with i.i.d. payoffs having a continuous distribution, as the number of actions of at least two players diverges, the asymptotic distribution of the number of PNE is Poisson(1). Stanford (1995) computed the exact nonasymptotic form of this distribution, from which the result in Powers (1990) can be obtained as a corollary. Rinott and Scarsini (2000) retained the i.i.d. assumptions for payoff vectors corresponding to different action profiles, but allowed dependence for payoffs within the same profile. They proved an interesting phase transition in terms of the payoffs' correlation: asymptotically in either the number of players or the number of actions, for negative dependence the number of PNE goes to 0, for positive dependence it diverges, and for independence it is Poisson(1), as proved by Powers (1990). Baldi et al. (1989) studied the distribution of the number of local maxima on a graph, which—by choosing a suitable graph—can be translated into the number of PNE in a random potential game.

Pei and Takahashi (2019) studied point-rationalizable strategies in two-person random games. Since the number of point-rationalizable strategies for each player is weakly larger than the number of PNE, they were interested in the typical magnitude of the difference between these two numbers. A game is dominance solvable if iterated elimination of strictly dominated strategies leads to a unique action profile, which must be a PNE. Alon et al. (2021) used recent combinatoric results to prove that the probability that a two-person random game is dominance solvable vanishes with the number of actions.

Several papers studied the behavior of various learning dynamics BRD in games with random payoffs. For instance, Galla and Farmer (2013) studied a type of reinforcement learning called experience-weighted attraction in two-person games and showed the existence of three different regimes in terms of convergence to equilibria. Sanders et al. (2018) extended their analysis to games with an arbitrary finite number of players. Pangallo et al. (2019) compared through simulation the behavior of various adaptive learning procedures in games whose payoffs are drawn at random. Heinrich et al. (2023) compared the behavior of BRD in games with random payoffs, when the order of acting players is fixed vs when it is random and they showed that, asymptotically in either the number of players or the number of strategies, the fixed-order BRD converges with vanishing probability, whereas the random-order does converge to a PNE whenever it exists. Similar results were obtained by Wiese and Heinrich (2022).

Coucheney et al. (2014); Durand and Gaujal (2016); Durand et al. (2019) focused on random potential games and measured the speed of convergence of BRD to a PNE. Amiet et al. (2021a) dealt with two-person games where the players have the same action set and payoffs are i.i.d. with a continuous distribution. They compared the behavior of best response dynamics and better response dynamics. They proved that, asymptotically in the number of actions, the first reaches a PNE only with vanishing probability, whereas the second does reach it, whenever it exists. Amiet et al. (2021b) studied a class of games with n players and two actions for each player where the payoffs are i.i.d. but their distribution may have atoms. They proved that the relevant parameter for the analysis of this class of games is the probability of ties in the payoffs, called α . They showed that, whenever this parameter is positive,

the number of PNE diverges, as $n \rightarrow \infty$ and proved a central limit theorem for this random variable. Moreover, using percolation techniques, they studied the asymptotic behavior of BRD, as a function of α , and they showed a phase transition at $\alpha = 1/2$. Johnston et al. (2023) considered the class of games whose random payoffs are i.i.d. with a continuous distribution; they showed that in almost every game in this class that has a pure Nash equilibrium, asymptotically in the number of players, best response dynamics can lead from every action profile that is not a pure Nash equilibrium to every pure Nash equilibrium.

Potential functions in games were introduced by Rosenthal (1973) and their properties were extensively studied by Monderer and Shapley (1996). Among them, existence of PNE and convergence to one of these equilibria of the most common learning procedures, including BRD.

Goemans et al. (2005) introduced the concept of sink equilibrium. Sink equilibria are strongly connected stable sets of action profiles that are never abandoned once reached by a BRD. A sink equilibrium that is not a PNE is what in this paper is called a trap.

Fabrikant et al. (2013) studied the class of weakly acyclical games, i.e., the class of games for which from every action profile, there exists some better-response improvement path that leads from that action profile to a PNE. This class includes potential games and dominance solvable games as particular cases.

The goal of our paper is to consider probability measures on spaces of finite noncooperative games that go beyond the usual assumption of i.i.d. payoffs. In particular, we define a parametric family of probability measures that interpolates random games with i.i.d. payoffs and random potential games. The interpolation is achieved locally by acting on each action profile of the game and replacing—with some fixed probability and independently across profiles—the payoff of the second player with the payoff of the first player. A different interpolation could be achieved by considering a convex combination of a game with i.i.d. payoffs and a random potential game. This was done, e.g., in Rinott and Scarsini (2000), where in each action profile the payoffs are obtained by summing a Gaussian vector with i.i.d. components and an independent Gaussian vector with identical components (which plays the role of the *random potential*). This approach is somehow comparable with the idea of decomposing the space of finite games proposed by Candogan et al. (2011). This decomposition was then used by Candogan et al. (2013) to analyze BRD in games that are close to potential games.

1.4. Organization of the paper

Section 2 introduces some basic game theoretic concepts. Section 3 defines the parametric family of distributions on the space of games and deals with the number of PNE in games with random payoffs. Section 4 studies the behavior of BRD in games with random payoffs in the above parametric class, for different values of the parameter. Section 5 contains all the proofs. Conclusions and open problems can be found in Section 6. Appendix A lists the symbols used throughout the paper. Appendix B contains two well-known results about the Beta distribution.

1.5. Notation

Given an integer n , the symbol $[n]$ indicates the set $\{1, \dots, n\}$. Given a finite set A , the symbol $|A|$ denotes its cardinality. The symbol \sqcup denotes the union of disjoint sets. We use the notation $x \wedge y := \min\{x, y\}$. The symbol \xrightarrow{p} denotes convergence in probability.

Given two nonnegative sequences h_n, g_n , we use the following common asymptotic notations:

$$h_n = o(g_n) \quad \text{if} \quad \lim_{n \rightarrow \infty} \frac{h_n}{g_n} = 0, \tag{1.1}$$

$$h_n = \mathcal{O}(g_n) \quad \text{if} \quad \limsup_{n \rightarrow \infty} \frac{h_n}{g_n} < \infty, \tag{1.2}$$

$$h_n = \Omega(g_n) \quad \text{if} \quad \liminf_{n \rightarrow \infty} \frac{h_n}{g_n} > 0, \tag{1.3}$$

$$h_n = \omega(g_n) \quad \text{if} \quad \lim_{n \rightarrow \infty} \frac{h_n}{g_n} = \infty, \tag{1.4}$$

$$h_n = \Theta(g_n) \quad \text{if} \quad h_n = \mathcal{O}(g_n) \quad \text{and} \quad h_n = \Omega(g_n). \tag{1.5}$$

2. Preliminaries

We consider two-person normal-form games where, for $i \in \{A, B\}$, player i 's action set is $[K^i] := \{1, \dots, K^i\}$ and $U^i : [K^A] \times [K^B] \rightarrow \mathbb{R}$ is player i 's payoff function. The game is defined by the payoff bimatrix

$$U := (U^A, U^B), \tag{2.1}$$

where, for $i \in \{A, B\}$,

$$U^i := (U^i(a, b))_{a \in [K^A], b \in [K^B]}. \tag{2.2}$$

A pure Nash equilibrium (PNE) of the game is a pair (a^*, b^*) of actions such that, for all $a \in [K^A], b \in [K^B]$ we have

$$U^A(a^*, b^*) \geq U^A(a, b^*) \quad \text{and} \quad U^B(a^*, b^*) \geq U^B(a^*, b). \tag{2.3}$$

As is well known, PNE are not guaranteed to exist. A class of games that admits PNE is the class of *potential games*, i.e., games for which there exists a *potential function* $\Psi : [K^A] \times [K^B] \rightarrow \mathbb{R}$ such that for all $a, a' \in [K^A]$, for all $b, b' \in [K^B]$, we have

$$U^A(a, b) - U^A(a', b) = \Psi(a, b) - \Psi(a', b), \tag{2.4a}$$

$$U^B(a, b) - U^B(a, b') = \Psi(a, b) - \Psi(a, b'). \tag{2.4b}$$

Games of *common interest*, i.e., games for which $U^A = U^B$, are a particular case of potential games. As far as PNE are concerned, every potential game is strategically equivalent to a common interest game, for instance to the game where $U^A = U^B = \Psi$. For the properties of potential games with an arbitrary number of players, we refer the reader to Monderer and Shapley (1996).

Given a finite game, it is interesting to see whether an equilibrium can be reached iteratively by allowing players to deviate whenever they have an incentive to do so. In particular, we will consider a procedure where, starting from a fixed action profile, players in alternation choose their best response to the other player’s action. If the procedure gets stuck in an action profile, then it has reached a pure Nash equilibrium. In general, there is no guarantee that this occurs.

Assume that the payoffs of each player are all different, i.e., $U^A(a, b) \neq U^A(a', b)$ for all $a \neq a'$ and all b (and similarly for the second player). The *best response dynamics* (BRD) is a learning algorithm taking as input a two-player game (U^A, U^B) and a starting action profile (a_0, b_0) . For each $t \geq 0$ we consider the process $\text{BRD}(t)$ on $[K^A] \times [K^B]$ such that

$$\text{BRD}(0) = (a_0, b_0) \tag{2.5}$$

and, if $\text{BRD}(t) = (a', b')$, then, for t even,

$$\text{BRD}(t + 1) = (a'', b'), \tag{2.6}$$

where $a'' \in \arg \max_{a \in [K^A]} U^A(a, b') \setminus \{a'\}$, if the latter set is not empty, otherwise

$$\text{BRD}(t + 1) = \text{BRD}(t); \tag{2.7}$$

for t odd,

$$\text{BRD}(t + 1) = (a', b''), \tag{2.8}$$

where $b'' \in \arg \max_{b \in [K^B]} U^B(a', b) \setminus \{b'\}$, if the latter set is not empty, otherwise

$$\text{BRD}(t + 1) = \text{BRD}(t). \tag{2.9}$$

It is easy to see that, if, for some positive \hat{t} , we have

$$\text{BRD}(\hat{t}) = \text{BRD}(\hat{t} + 1) = (a^*, b^*), \tag{2.10}$$

then $\text{BRD}(t) = (a^*, b^*)$ for all $t \geq \hat{t}$ and (a^*, b^*) is a PNE of the game.

The algorithm stops when it visits an action profile for the second time. If this profile is the same as the one visited at the previous time, then a PNE has been reached.

Inspired by the concept of *sink equilibrium* of Goemans et al. (2005), we give a definition of trap in a way that is suitable for our two-player environment.

Definition 2.1. A *trap* is a finite set M of action profiles such that

- (a) $|M| \geq 2$,
- (b) if $\text{BRD}(t) \in M$, then $\text{BRD}(t + 1) \in M$;
- (c) for every $(a, b) \in M$, there exists t such that $\text{BRD}(t + k|M|) = (a, b)$, for every $k \in \mathbb{N}$.

Moreover, the definition of BRD implies that for every trap M we have $|M| \geq 4$ and $|M|$ even. Moreover, for every game (U^A, U^B) and every initial profile (a, b) , the BRD eventually visits a PNE or a trap in finite time, say τ .

Even if the game admits PNE, there is no guarantee that a BRD reaches one of them; it could cycle over a trap, i.e., it could start to periodically visit the same set of action profiles and never stabilize. On the other hand, if the game is a potential game, then a BRD always reaches a PNE. This is due to the fact that at every iteration of the BRD the payoff of one player increases, and so does the potential. Since the game is finite, in finite time the BRD reaches a local maximum, which is a PNE (see, e.g., Karlin and Peres, 2017, proposition 4.4.6).

The goal of this paper is to study the number of PNE and the behavior of BRD in a “typical” game. To make sense of the above sentence, we need to formalize the meaning of the term typical. The approach that we will take is stochastic. That is, we will assume the bimatrix U to be random and drawn from a distribution that will be specified later. In any game with random payoffs, the set

NE of pure Nash equilibria is a (possibly empty) random set of action profiles, i.e., a random subset of $[K^A] \times [K^B]$. Therefore, since the game is finite, the number of PNE is an integer-valued random variable. Moreover, we will be able to speak about the probability that a BRD converges (to a PNE).

3. Number of pure Nash equilibria in random games

As mentioned in the Introduction, several attempts have been made in the literature to put a probability measure on a space of games. Most of the existing papers assume all the entries of U to be i.i.d. with a continuous marginal distribution. There are some notable exceptions to the independence assumption. Rinott and Scarsini (2000) considered a setting where the payoff vectors of different action profiles have a continuous distribution and are i.i.d., but some dependence is allowed within each profile. Durand and Gaujal (2016) studied random potential games where the entries of the potential are i.i.d. with a continuous distribution.

Our stochastic model is quite general, since—in a sense that will be made precise—it interpolates the i.i.d. payoffs and the random potential.

By definition, the concept of pure Nash equilibrium is ordinal, that is, if all payoffs in a game are transformed according to a strictly increasing function, then the set of pure Nash equilibria remains the same. Assume that each entry of U has a marginal distribution that is uniform on the interval $[0, 1]$. Its distribution function will be denoted by F . The above consideration implies that this uniformity assumption is without loss of generality, i.e., any other continuous distribution would produce the same conclusions.

Start with (U^A, U^B) , where all the entries are i.i.d. with distribution F . Then, for each action profile (a, b) , with probability p set $U^B(a, b)$ to be equal to $U^A(a, b)$, independently of the other action profiles. In other words, for every pair (a, b) ,

- with probability $1 - p$, the random payoffs $U^A(a, b)$ and $U^B(a, b)$ are independent,
- with probability p , we have $U^A(a, b) = U^B(a, b)$.

The larger p , the closer the game is to a potential game. The smaller p , the closer the game is to a random game with i.i.d. payoffs. The game whose payoff bimatrix is obtained as above will be denoted by $U(p)$.

We now compute the expected number of PNE in the above-defined class of random games.

Proposition 3.1. *If W is the number of PNE in the game $U(p)$, then*

$$E[W] = p \frac{K^A K^B}{K^A + K^B - 1} + (1 - p). \tag{3.1}$$

The analysis of this class of games is quite complicated for fixed K^A, K^B . Therefore, as it is done in much of the literature, we will take an asymptotic approach, letting the number of actions grow. More formally, we will consider a sequence $(U_n)_{n \in \mathbb{N}}$ of payoff bimatrices, where the numbers of actions in game U_n are K_n^A and K_n^B , and these two integer sequences are increasing in n and diverge to ∞ . In particular, we allow the number of actions of the two players to diverge at different speeds.

The following proposition shows the asymptotic behavior of the random number of PNE where the parameter p may vary with n . We write p_n to highlight this dependence. In what follows, every asymptotic equality holds for $n \rightarrow \infty$. The proof uses a second-moment argument.

For every $n \in \mathbb{N}$, let W_n be the number of PNE in the game U_n and

$$K_n := \min(K_n^A, K_n^B). \tag{3.2}$$

Proposition 3.2. *If $p_n = \omega(1/K_n)$, then*

$$\frac{K_n^A + K_n^B}{p_n K_n^A K_n^B} W_n \xrightarrow{P} 1. \tag{3.3}$$

The following corollary deals with the case of fixed p .

Corollary 3.3. *If $p_n = p$ for all $n \in \mathbb{N}$ and $K_n^B = \omega(K_n^A)$, then*

$$\frac{W_n}{K_n^A} \xrightarrow{P} p. \tag{3.4}$$

If $K_n^B = \alpha_n K_n^A$, with $\alpha_n \rightarrow \alpha$, then

$$\frac{W_n}{K_n^A} \xrightarrow{P} \frac{\alpha}{\alpha + 1} p. \tag{3.5}$$

In particular, if $K_n^A = K_n^B$, then

$$\frac{W_n}{K_n} \xrightarrow{P} \frac{p}{2}. \tag{3.6}$$

When $p = 0$, i.e., the payoffs are i.i.d. the number of PNE converges in distribution to a Poisson with parameter 1 (see Powers, 1990). When the payoffs within the same action profile are positively correlated, Rinott and Scarsini (2000) showed that the number of PNE diverges. A similar phenomenon happens here when $p > 0$.

4. Best response dynamics

We now want to study the behavior of BRD in the class of random games introduced in Section 3. First notice that the continuity of F implies that the probability of ties in the payoffs of the same player is zero; as a consequence, once the game is realized, a BRD is almost surely deterministic. In this respect, the symbols P and E refer solely to the randomness of the payoffs, not to any randomness in the BRD. Moreover, the symmetry of our model implies that, without loss of generality, we can assume the starting position of the BRD to be any fixed profile. In the rest of the paper, without loss of generality, the starting point of any BRD will always be the profile $(1, 1)$, i.e., for every $n \in \mathbb{N}$

$$P(\text{BRD}_n(0) = (1, 1)) = 1. \tag{4.1}$$

4.1. BRD and related stopping times

As stressed above, for a given realization of the payoffs and a starting point, the BRD is a deterministic algorithm that decides its next step only on the basis of *local information*. In what follows we will exploit this fact by revealing the players' payoffs only when this information is needed to select the next position of the BRD. This whole process, in which the BRD moves on a sequentially sampled random game, can be thought of as a non-Markov stochastic process, and the time at which the BRD stops can be seen as a stopping time for such a stochastic process.

We will focus our attention on the distribution of the first time the BRD reaches a PNE. For the sake of brevity, we write NE_n for the (random) set of PNE in the game $U_n(p_n)$, and we define

$$\tau_n^{\text{NE}} := \min\{t : \text{BRD}_n(t) \in \text{NE}_n\}. \tag{4.2}$$

In words, τ_n^{NE} is the first time the process $\text{BRD}_n(t)$ visits a PNE. Notice that the first step of the BRD is somehow different from the following, because, by the definition of the model, at time 0 no player is assumed to be already in a best response. Contrarily, for any $t \geq 1$ odd (resp. even) we have that the first (resp. second) player is in a best response, and the other player's action can be changed at the following step, if it is not itself a best response. As a consequence of this fact, the forthcoming definitions that depend on the step t of the BRD, require a special treatment for the first few steps, i.e., $t = 0, 1, 2$. Clearly, this issue could be solved by assuming that at $t = 0$ the second player is in best response. The latter assumption has been made, e.g., in Amiet et al. (2021a). For the sake of generality, we prefer to avoid this assumption and rather treat the case of $t = 0, 1, 2$ separately. We will make use of the following sequence of random sets:

$$R_n(t) := \{(a, b) : \text{either } \text{BRD}_n(s) = (a', b) \text{ or } \text{BRD}_n(s) = (a, b'), \\ \text{for some } a' \in [K_n^A], b' \in [K_n^B] \text{ and } 1 \leq s \leq t\}, \quad t \geq 1, \tag{4.3}$$

$$R_n(0) := \{(a, 1) : a \in [K_n^A]\}. \tag{4.4}$$

Roughly speaking, the random set $R_n(t)$ is the set of all rows and columns that contain one element that has been visited by the BRD up to time t . More precisely, for $t \geq 1$, in order to determine where $\text{BRD}_n(t)$ is and whether it is in a PNE or not, for each action profile $(a, b) \in R_n(t)$, at least one of the payoffs $U^A(a, b)$ and $U^B(a, b)$, at some time $s \leq t$, needs to be revealed.

For instance, $\text{BRD}_n(t) = (a, 1)$ for some $a \in [K_n^A]$. To determine whether this profile $(a, 1)$ is a PNE or not, the payoff $U^B(a, 1)$ has to be compared with all payoffs $U^B(a, b)$ for each $b \in [K_n^B]$.

To better understand the above definition, consider the random time

$$\tau_n^R := \min\{t \geq 2 : \text{BRD}_n(t) \in R_n(t-2)\}. \tag{4.5}$$

By the definitions in (4.3) and (4.4), at time τ_n^R

- (i) either the BRD has reached an equilibrium, i.e., $\text{BRD}_n(\tau_n^R) = \text{BRD}_n(\tau_n^R - 1) \in \text{NE}_n$;
- (ii) or the BRD has reached a trap, i.e., there exists some $t \leq \tau_n^R - 3$ such that $\text{BRD}_n(\tau_n^R + 1) = \text{BRD}_n(t)$.

An example of the first steps of the BRD is given in Fig. 1. The red (respectively blue) dots are the action profiles whose payoff is compared by the row (respectively column) player. The red (respectively blue) lines help visualize the action profiles that the BRD considers when the row (respectively column) player is active in the BRD. The intersections between these lines are action profiles in which the payoffs of both the column and row player have been examined by the BRD. Such points are represented by two overlapping dots of different size: the color of the biggest (respectively smallest) one is associated to the first (respectively second) player that compares the payoff of such action profiles. The numbers indicate the positions of the BRD at the various times. The odd (respectively even) numbers are in red (respectively blue) since the associated action profiles are visited for the first time when the row (respectively column) player is active. The left side of Fig. 1 shows an instance of what occurs in case (i) in the above list, whereas the right side of the figure contains a graphical explanation of what happens in case (ii).




Fig. 1. Both figures show an instance of the first seven steps in the BRD. The figure on the left describes the case in which $\tau_n^{NE} = 6$. To compute $BRD_n(7)$, the row player visits the action profiles on the red dashed lines and finds the maximum payoff at $BRD_n(6)$; hence $BRD_n(6) \in NE_n$. In this case $R_n(5)$ consists of all the action profiles on the solid red and blue lines in the figure. Hence, $\tau_n^R = 7$ and, consequently, $\tau_n^R - 1 = \tau_n^{NE}$. The figure on the right describes the case in which the BRD discovers a trap. Since $BRD_n(6)$ and $BRD_n(3)$ are in the same column, we have $BRD_n(7) = BRD_n(3)$. In this case $\tau_n^R = 6$ because $R_n(4)$ consists of all the action profiles on the solid lines in the figure, except for the blue line passing through the action profiles labeled 5 and 6. Hence, $BRD_n(\tau_n^R + 1) = BRD_n(t)$ for $t = 3 \leq \tau_n^R - 3$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

In general, the trajectory of the BRD for all $t \geq 0$ is completely determined by the trajectory up to the random time τ_n^R . Concerning the random time τ_n^{NE} defined in (4.2), notice that

- (a) in the case described in (i), we have $\tau_n^{NE} = \tau_n^R - 1$ if $\tau_n^R > 2$, and $\tau_n^{NE} = 0$ if $\tau_n^R = 2$ (see Fig. 1 on the left);
- (b) whereas, in the case in (ii), $\tau_n^{NE} = \infty$.

By (4.3) and (4.5), we have

$$P(\tau_n^R \leq 2K_n) = 1. \tag{4.6}$$

Moreover, by (4.6) and (a), if a PNE is eventually reached, then the BRD must visit the set NE_n for the first time within $2K_n - 1$ steps. Formally,

$$P(\tau_n^{NE} \leq 2K_n - 1) = 1 - P(\tau_n^{NE} = \infty). \tag{4.7}$$

As shown in the following lemma, another relevant feature of (4.3) and (4.4) is that, for all $t \geq 1$, conditionally on the event $\{\tau_n^R > t\}$, even if the set $R_n(t)$ is random, its cardinality $|R_n(t)|$ is almost surely equal to the deterministic value r_n .

Lemma 4.1. Fix any $n \in \mathbb{N}$ and $p_n \in [0, 1]$. Call, for every $t \geq 1$,

$$r_n(t) := \left\lceil \frac{t+1}{2} \right\rceil K_n^A + \left\lfloor \frac{t+1}{2} \right\rfloor K_n^B - \left\lfloor \frac{t+1}{2} \right\rfloor \left\lceil \frac{t+1}{2} \right\rceil. \tag{4.8}$$

Then, for every $t \geq 1$,

$$P(|R_n(t)| = r_n(t) \mid \tau_n^R > t) = 1. \tag{4.9}$$

4.2. Potential games

We start studying the case of $p = 1$, i.e., the case of potential games. A potential game cannot have traps. Therefore, thanks to (a),

$$\{\tau_n^R = t\} = \{\tau_n^{NE} = t - 1\}, \quad \text{for } t > 2, \tag{4.10}$$

and, together with (4.7) and (4.6), we recover the well-known fact

$$P(\tau_n^{NE} < \infty) = 1, \quad n \geq 1. \tag{4.11}$$

Define

$$q_{t,n} := P(\tau_n^{\text{NE}} = t \mid \tau_n^{\text{NE}} \geq t) \quad \text{for all } t \in \{0, \dots, 2K_n - 1\}. \tag{4.12}$$

In words, $q_{t,n}$ represents the conditional probability that a PNE is reached at time t , given that it was not reached before. Notice that if $t = 2K_n - 1$, then $q_{t,n} = 1$. We start with a non-asymptotic result, which, for every $n \in \mathbb{N}$, provides the value of $q_{t,n}$ for all $t \leq 2K_n - 1$.

Lemma 4.2. Fix $n \in \mathbb{N}$, let $p_n = 1$ and recall the definition of $r_n(t)$ in (4.8). Then,

$$q_{0,n} = \frac{1}{K_n^A + K_n^B - 1}, \tag{4.13}$$

$$q_{1,n} = \frac{K_n^A - 1}{K_n^A + K_n^B - 2}, \tag{4.14}$$

and, for $t \geq 2$,

$$q_{t,n} = \frac{r_n(t-1)}{r_n(t)}. \tag{4.15}$$

As an immediate consequence of the previous lemma, we obtain the exact form of the distribution of the random time τ_n^{NE} .

Theorem 4.3. With the definitions of Lemma 4.2, we have

$$P(\tau_n^{\text{NE}} > t) = \prod_{j=0}^t (1 - q_{j,n}). \tag{4.16}$$

The following proposition provides the asymptotic expectation and variance of τ_n^{NE} when the two players have the same action set.

Proposition 4.4. If $p_n = 1$ and $K_n^A = K_n^B$ for every $n \in \mathbb{N}$. Then

$$\lim_{n \rightarrow \infty} E[\tau_n^{\text{NE}}] = e - 1 \tag{4.17}$$

$$\lim_{n \rightarrow \infty} \text{Var}[\tau_n^{\text{NE}}] \approx 0.767. \tag{4.18}$$

We point out that our expression for $E[\tau_n^{\text{NE}}]$ coincides with the one in Durand and Gaujal (2016, theorem 4), but the proof techniques exploited in Theorem 4.3 are completely different from the ones used by Durand and Gaujal (2016); moreover, our analysis considers the whole distribution of τ_n^{NE} and not just its expectation.

4.3. Games with i.i.d. payoffs

In a recent paper Amiet et al. (2021a, theorem 2.3) showed that, when $p_n = 0$ and $K_n^A = K_n^B$, with high probability as $n \rightarrow \infty$, the BRD does not converge to the set NE_n . We start by generalizing their result to the general setting $K_n^A \neq K_n^B$.

Theorem 4.5. Let $p_n = 0$ for all $n \in \mathbb{N}$. Then

$$\lim_{n \rightarrow \infty} P(\tau_n^{\text{NE}} < \infty) = 0. \tag{4.19}$$

Even though the result in Theorem 4.5 can be achieved by naturally adapting the proof of Amiet et al. (2021a, theorem 2.3) to the rectangular case, for completeness we present a detailed proof in Section 5. It is worth stressing that, contrarily to our setting, Amiet et al. (2021a) assume that the BRD starts at an action profile in which the second player is already at best-response. The fact that we dispense with this assumption makes the quantities appearing in our proof slightly different from those in Amiet et al. (2021a).

4.4. The general case

The main purpose of this paper is to complement the negative result in Theorem 4.5 by showing that a tiny bit of (positive) correlation in the players payoffs, namely $p_n > 0$, is enough to dramatically change the picture and make it similar to the case of a potential game, i.e., $p_n = 1$.

More precisely, the following result shows that, if p_n is not too small compared to K_n , then the probability of the event $\{\tau_n^{\text{NE}} = \infty\}$ vanishes as n goes to infinity.

Theorem 4.6. Fix a positive sequence p_n . If

$$\lim_{n \rightarrow \infty} \frac{\log(p_n)}{\log(K_n)} = 0, \tag{4.20}$$

then

$$\lim_{n \rightarrow \infty} P(\tau_n^{\text{NE}} = \infty) = 0. \tag{4.21}$$

Notice that the latter result is qualitative in nature: it tells us that the BRD will eventually converge to a Nash equilibrium rather than to a trap, but does not provide any bound on the rate of convergence beyond the trivial one presented in (4.7). The following result—which implies Theorem 4.6—provides a much better bound on the time of convergence. Indeed, it states that, under the condition in (4.20), the time of convergence of the BRD to an equilibrium can be upper bounded, with high probability, by any function diverging exponentially faster than p_n^{-1} .

Proposition 4.7. Fix a positive sequence p_n satisfying (4.20). Then, for any sequence ℓ_n such that

$$\lim_{n \rightarrow \infty} \ell_n = \infty \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{\log(p_n)}{\log(\ell_n)} = 0, \tag{4.22}$$

we have

$$\lim_{n \rightarrow \infty} P(\tau_n^{\text{NE}} > \ell_n) = 0. \tag{4.23}$$

Notice that Theorem 4.6 follows from Proposition 4.7 by choosing $\ell_n = 2K_n - 1$. In the particular case when $p_n = p > 0$ for all $n \in \mathbb{N}$, Proposition 4.7 states that, with high probability, a PNE is reached by the BRD in at most ℓ_n steps, no matter how slowly the sequence ℓ_n diverges.

Remark 4.8. Notice that, whenever $p_n < 1$, the random variable τ_n^{NE} takes value $+\infty$ with positive probability. Therefore, even though (4.21) holds true, the random variable τ_n^{NE} cannot have a finite expectation.

5. Proofs

Proofs of Section 3

Proof of Proposition 3.1. By linearity of expectation, we have

$$E[W] = K^A K^B P((1, 1) \in \text{NE}). \tag{5.1}$$

Conditioning on the event $\{U^A(1, 1) = U^B(1, 1)\}$ and using the law of total probability, we get

$$P((1, 1) \in \text{NE}) = p \frac{1}{K^A + K^B - 1} + (1 - p) \frac{1}{K^A K^B}. \tag{5.2}$$

To see that (5.2) holds, consider the following: when $U^A(1, 1) = U^B(1, 1)$, the profile $(1, 1)$ is a PNE if and only if $U^A(1, 1)$ is larger than or equal to all payoffs $U^A(a, 1)$ and all payoffs $U^B(1, b)$ for all $a \in [K^A]$ and all $b \in [K^B]$. By symmetry, this happens with probability $1/(K^A + K^B - 1)$. On the other hand, when $U^A(1, 1) \neq U^B(1, 1)$, the profile $(1, 1)$ is a PNE if and only if $U^A(1, 1)$ is larger than or equal to all payoffs $U^A(a, 1)$ for all $a \in [K^A]$ and $U^B(1, 1)$ is larger than or equal to all payoffs $U^B(1, b)$ for all $b \in [K^B]$. By independence of the payoffs and by symmetry, this happens with probability $1/K^A K^B$. Plugging (5.2) into (5.1), we get the result. \square

Proof of Proposition 3.2. By Proposition 3.1 we get

$$E[W_n] = p_n \frac{K_n^A K_n^B}{K_n^A + K_n^B - 1} + (1 - p_n) = (1 + o(1)) p_n \frac{K_n^A K_n^B}{K_n^A + K_n^B}, \tag{5.3}$$

where, in the asymptotic equality we used the fact that $p_n K_n \rightarrow \infty$. This shows that $E[W_n] \rightarrow \infty$.

We now show that W_n concentrates around its expectation. To this end, we use an upper bound on the second moment of W_n . Notice that

$$\begin{aligned} E[W_n^2] &= E\left[\left(\sum_{(a,b) \in [K_n^A] \times [K_n^B]} \mathbb{1}_{(a,b) \in NE_n}\right)^2\right] \\ &= \sum_{(a,b) \in [K_n^A] \times [K_n^B]} \sum_{(a',b') \in [K_n^A] \times [K_n^B]} P((a,b), (a',b') \in NE_n). \end{aligned} \tag{5.4}$$

From (5.4) it follows immediately that the computation of the second moment amounts to studying probabilities of the form

$$P((a,b), (a',b') \in NE_n), \quad a, a' \in [K_n^A], b, b' \in [K_n^B].$$

We argue that there are only three relevant cases:

- (a) $a = a'$ and $b = b'$;
- (b) $a = a'$ and $b \neq b'$ or vice versa;
- (c) $a \neq a'$ and $b \neq b'$;

Case (a) is trivial:

$$P((a,b), (a',b') \in NE_n) = P((a,b) \in NE_n). \tag{5.5}$$

The continuity of F implies that in Case (b) we have

$$P((a,b), (a',b') \in NE_n) = 0. \tag{5.6}$$

To analyze Case (c), we remark that, if $a \neq a'$ and $b \neq b'$, then the event $\{(a',b') \in NE_n\}$ depends on the event $\{(a,b) \in NE_n\}$ only through the payoffs $U^A(a',b)$ and $U^B(a,b')$. Therefore

$$\begin{aligned} P((a,b), (a',b') \in NE_n) &= P((a,b) \in NE_n)P((a',b') \in NE_n \mid (a,b) \in NE_n) \\ &\leq P((a,b) \in NE_n)P\left(U^A(a',b') = \max_{a'' \in [K_n^A] \setminus \{a\}} U^A(a'',b'), \right. \\ &\quad \left. U^B(a',b') = \max_{b'' \in [K_n^B] \setminus \{b\}} U^B(a',b'')\right) \\ &= P((a,b) \in NE_n) \left(p_n \frac{1}{K_n^A + K_n^B - 3} + (1 - p_n) \frac{1}{(K_n^A - 1)(K_n^B - 1)} \right) \\ &= (1 + o(1))(P((a,b) \in NE_n))^2. \end{aligned} \tag{5.7}$$

The inequality in (5.7) stems from the fact that the event $(a',b') \in NE_n$ coincides with the event

$$\left\{ U^A(a',b') = \max_{a'' \in [K_n^A]} U^A(a'',b'), U^B(a',b') = \max_{b'' \in [K_n^B]} U^B(a',b'') \right\},$$

which in turn implies the event

$$\left\{ U^A(a',b') = \max_{a'' \in [K_n^A] \setminus \{a\}} U^A(a'',b'), U^B(a',b') = \max_{b'' \in [K_n^B] \setminus \{b\}} U^B(a',b'') \right\}.$$

The independence between this latter event and $\{(a,b) \in NE_n\}$ proves the inequality. The second equality in (5.7), follows from the fact that, when $U^A(a',b') = U^B(a',b')$, which happens with probability p_n , the event

$$\left\{ U^A(a',b') = \max_{a'' \in [K_n^A] \setminus \{a\}} U^A(a'',b'), U^B(a',b') = \max_{b'' \in [K_n^B] \setminus \{b\}} U^B(a',b'') \right\} \tag{5.8}$$

has the same probability as the event of picking the maximum among $K_n^A + K_n^B - 3$ equally probable objects; on the other hand, when $U^A(a',b') \neq U^B(a',b')$, the event in (5.8) has the same probability of picking independently the maximum of $K_n^A - 1$ equally probable objects and the maximum of $K_n^B - 1$ equally probable objects. The last equality stems from (5.2).

In conclusion, plugging (5.5)–(5.7) into (5.4), we obtain

$$\begin{aligned} E[W_n^2] &= \sum_{(a,b) \in [K_n^A] \times [K_n^B]} P((a,b) \in NE_n) \\ &\quad + (1 + o(1)) \sum_{(a,b) \in [K_n^A] \times [K_n^B]} \sum_{\substack{(a',b') \in [K_n^A] \times [K_n^B] \\ a \neq a', b \neq b'}} P((a,b) \in NE_n))^2 \\ &< E[W_n] + (1 + o(1))(E[W_n])^2 \\ &= (1 + o(1))(E[W_n])^2, \end{aligned} \tag{5.9}$$

where the last equality stems from the fact that $E[W_n] \rightarrow \infty$, which implies $E[W_n] = o((E[W_n])^2)$.

By Chebyshev’s inequality, we have that, for every $\varepsilon > 0$,

$$P(|W_n - E[W_n]| \geq \varepsilon E[W_n]) \leq \frac{E[W_n^2] - (E[W_n])^2}{\varepsilon^2 (E[W_n])^2} = o(1), \tag{5.10}$$

where the asymptotic upper bound follows from (5.9). \square

Proofs of Section 4.1

Proof of Lemma 4.1. We prove (4.9) by induction. For $t = 1$ we have $|R_n(1)| = K_n^A + K_n^B - 1$, which is trivially true. Assume now that (4.9) holds up to $t - 1 < \tau_n^R$. Notice that, conditioning on $\{\tau_n^R > t\}$, $BRD_n(t)$ cannot visit a row or column visited at time $1, 2, \dots, t - 1$. Since player A plays first, by the inductive hypothesis and thanks to the conditioning, for $0 < t < \tau_n^R$, we have almost surely,

$$|R_n(t)| = \begin{cases} r_n(t - 1) + K_n^B - \lfloor \frac{t+1}{2} \rfloor, & \text{if } t \text{ is odd;} \\ r_n(t - 1) + K_n^A - \lfloor \frac{t+1}{2} \rfloor, & \text{if } t \text{ is even.} \end{cases} \tag{5.11}$$

Since

$$r_n(t - 1) = \lfloor \frac{t}{2} \rfloor K_n^A + \lfloor \frac{t}{2} \rfloor K_n^B - \lfloor \frac{t}{2} \rfloor \lfloor \frac{t}{2} \rfloor, \tag{5.12}$$

we can rewrite the right hand side of (5.11) as (4.8). Hence, (4.9) holds. \square

Proofs of Section 4.2

Proof of Lemma 4.2. Since $p_n = 1$, we have $U_n^A = U_n^B$ almost surely. Therefore, to simplify the notation, we write

$$U_n(a, b) := U_n^A(a, b) = U_n^B(a, b). \tag{5.13}$$

Moreover, in a potential game $\{\tau_n^R = t\} = \{\tau_n^{NE} = t - 1\}$, for $t > 2$, as in (4.10). We start with $t = 0$. We have

$$\begin{aligned} q_{0,n} &= P(\tau_n^{NE} = 0) \\ &= P\left(U_n(1, 1) = \max \left\{ \max_{a \in [K_n^A]} U_n(a, 1), \max_{b \in [K_n^B]} U_n(1, b) \right\} \right) \\ &= \frac{1}{K_n^A + K_n^B - 1}. \end{aligned} \tag{5.14}$$

Let now $t = 1$. We have

$$P(\tau_n^{NE} \geq 1) = P(\tau_n^{NE} \geq 0) - P(\tau_n^{NE} = 0) = 1 - \frac{1}{K_n^A + K_n^B - 1} = \frac{K_n^A + K_n^B - 2}{K_n^A + K_n^B - 1}. \tag{5.15}$$

Moreover, if we define the event

$$A_a := \{\text{player A's best response to action 1 is } a\}, \tag{5.16}$$

we have

$$\begin{aligned}
 P(\tau_n^{\text{NE}} = 1) &= \sum_{a \in [K_n^A]} P(\tau_n^{\text{NE}} = 1 \mid A_a) P(A_a) \\
 &= \sum_{a=2}^{K_n^A} P(\text{Beta}(K_n^A, 1) \geq \text{Beta}(K_n^B - 1, 1)) \frac{1}{K_n^A} \\
 &= \sum_{a=2}^{K_n^A} \frac{K_n^A}{K_n^A + K_n^B - 1} \cdot \frac{1}{K_n^A} \\
 &= \frac{K_n^A - 1}{K_n^A + K_n^B - 1},
 \end{aligned}
 \tag{5.17}$$

where, with an abuse of notation, we have identified a random variable with its distribution. Conditionally on A_a , we have $\tau_n^{\text{NE}} = 1$ if the payoff in $(a, 1)$ is the largest among all payoffs in the same row. To get the result we have applied Proposition B.1 about the maximum of uniform independent random variables and Proposition B.2 about the probability that a $\text{Beta}(a, 1)$ is larger than an independent $\text{Beta}(b, 1)$. This result can be applied because the payoffs, and consequently the two Beta random variables, are independent. Therefore, combining (5.15) and (5.17), we obtain

$$\begin{aligned}
 q_{1,n} &= P(\tau_n^{\text{NE}} = 1 \mid \tau_n^{\text{NE}} \geq 1) \\
 &= \frac{P(\tau_n^{\text{NE}} = 1)}{P(\tau_n^{\text{NE}} \geq 1)} \\
 &= \frac{K_n^A + K_n^B - 1}{K_n^A + K_n^B - 2} \cdot \frac{K_n^A - 1}{K_n^A + K_n^B - 1} \\
 &= \frac{K_n^A - 1}{K_n^A + K_n^B - 2}.
 \end{aligned}
 \tag{5.18}$$

Let now $t \geq 2$. Call Π_t the set of sequences $\pi = (\pi_0, \pi_1, \dots, \pi_t)$ such that

- $\pi_i = (a_i, b_i) \in [K_n^A] \times [K_n^B]$ for $i \geq 0$,
- $\pi_0 = (1, 1)$,
- if i is odd, then $b_i = b_{i-1}$, whereas, if $i > 0$ is even, then $a_i = a_{i-1}$,
- there is no pair of distinct odd indices i, j such that $b_i = b_j$,
- there is no pair of distinct even indices i, j such that $a_i = a_j$.

Notice that the set Π_t coincides with all possible trajectories of length t of BRD_n satisfying the event $\{\tau_n^{\text{NE}} \geq t\}$. Define the event

$$G_t^\pi = \{\text{BRD}_n(s) = \pi_s \text{ for } 0 \leq s \leq t\}.
 \tag{5.19}$$

Notice that

$$q_{t,n} := P(\tau_n^{\text{NE}} = t \mid \tau_n^{\text{NE}} \geq t) = \sum_{\pi \in \Pi_t} P(\tau_n^{\text{NE}} = t \mid G_t^\pi, \tau_n^{\text{NE}} \geq t) P(G_t^\pi \mid \tau_n^{\text{NE}} \geq t).
 \tag{5.20}$$

The conditional probability $P(\tau_n^{\text{NE}} = t \mid G_t^\pi, \tau_n^{\text{NE}} \geq t)$ equals the probability that the maximum of $r_n(t - 1)$ i.i.d. uniform random variables is bigger than the maximum of $K_n^* - \lfloor \frac{t+1}{2} \rfloor$ uniform random variables, where

$$K_n^* = \begin{cases} K_n^A & \text{if } t \text{ is even,} \\ K_n^B & \text{if } t \text{ is odd.} \end{cases}
 \tag{5.21}$$

The conditioning event $G_t^\pi \cap \{\tau_n^{\text{NE}} \geq t\}$ determines the position of $\text{BRD}_n(t)$ and the fact that the payoffs associated to $r_n(t - 1)$ action profiles up to time t have been computed by the BRD. The payoff associated to $\text{BRD}_n(t)$ is the maximum of $r_n(t - 1)$ i.i.d. uniform random variables. The probability that the action profile $\text{BRD}_n(t)$ is a PNE is the probability that its payoff is larger than all the previously uncomputed payoffs associated to action profiles in its same column (if t is even) or in its same row (if t is odd). This requires a comparison with $K_n^* - \lfloor \frac{t+1}{2} \rfloor$ uniformly distributed payoffs. So

$$\begin{aligned}
 P(\tau_n^{\text{NE}} = t \mid G_t^\pi, \tau_n^{\text{NE}} \geq t) &= P\left(\text{Beta}(r_n(t - 1), 1) > \text{Beta}\left(K_n^* - \left\lfloor \frac{t+1}{2} \right\rfloor, 1\right)\right) \\
 &= \frac{r_n(t - 1)}{r_n(t - 1) + K_n^* - \left\lfloor \frac{t+1}{2} \right\rfloor}
 \end{aligned}$$

$$= \frac{r_n(t-1)}{r_n(t)},$$

where the last identity is due to (5.11). Using (5.20) we get

$$\begin{aligned} P(\tau_n^{NE} = t \mid \tau_n^{NE} \geq t) &= \sum_{\pi \in \Pi_t} P(\tau_n^{NE} = t \mid G_t^\pi, \tau_n^{NE} \geq t) P(G_t^\pi \mid \tau_n^{NE} \geq t) \\ &= \frac{r_n(t-1)}{r_n(t)} \sum_{\pi \in \Pi_t} P(G_t^\pi \mid \tau_n^{NE} \geq t). \end{aligned} \tag{5.22}$$

Note that $\{\tau_n^{NE} \geq t\} = \sqcup_{\pi \in \Pi_t} G_t^\pi$. Hence

$$\sum_{\pi \in \Pi_t} P(G_t^\pi \mid \tau_n^{NE} \geq t) = P(\tau_n^{NE} \geq t \mid \tau_n^{NE} \geq t) = 1. \tag{5.23}$$

Therefore, (5.22) becomes

$$P(\tau_n^{NE} = t \mid \tau_n^{NE} \geq t) = \frac{r_n(t-1)}{r_n(t)}. \quad \square$$

Proof of Theorem 4.3. Note that

$$P(\tau_n^{NE} > t) = P(\tau_n^{NE} > t \mid \tau_n^{NE} > t-1) P(\tau_n^{NE} > t-1) = (1 - q_{t,n}) P(\tau_n^{NE} > t-1). \tag{5.24}$$

Hence, by iteration, we get

$$P(\tau_n^{NE} > t) = \prod_{j=1}^t (1 - q_{j,n}) \cdot P(\tau_n^{NE} > 0) = \prod_{j=0}^t (1 - q_{j,n}). \quad \square$$

Proof of Proposition 4.4. We first compute $E[\tau_n^{NE}]$. By Theorem 4.3 we have

$$E[\tau_n^{NE}] = \sum_{t=0}^{2K_n-2} P(\tau_n^{NE} > t) = \sum_{t=0}^{2K_n-2} \prod_{j=0}^t (1 - q_{j,n}). \tag{5.25}$$

We split the first sum in (5.25) into two parts: $t \in \{0, \dots, \ell_n - 1\}$ and $t \in \{\ell_n, \dots, 2K_n - 2\}$, where $\ell_n = \lfloor (\log K_n)^2 \rfloor$. We start by showing that

$$\lim_{n \rightarrow \infty} \sum_{t=\ell_n}^{2K_n-2} \prod_{j=0}^t (1 - q_{j,n}) = 0, \tag{5.26}$$

for which it suffices to show that for $t \geq \ell_n$

$$P(\tau_n^{NE} > t) = \prod_{j=0}^t (1 - q_{j,n}) = o(K_n^{-1}). \tag{5.27}$$

Notice that the sequence $q_{t,n}$, defined in (4.12), is increasing in t . Hence, for $t \geq \ell_n$, we have

$$P(\tau_n^{NE} > t) \leq (1 - q_{0,n})(1 - q_{1,n})^{\ell_n} \leq \frac{1}{2^{\ell_n}} = o(K_n^{-1}), \tag{5.28}$$

where in the last asymptotic equality we used $q_{1,n} = 1/2$ and $\ell_n = \omega(\log(K_n))$.

We are left to show that

$$\lim_{n \rightarrow \infty} \sum_{t=0}^{\ell_n} \prod_{j=0}^t (1 - q_{j,n}) = e - 1. \tag{5.29}$$

Notice that

$$1 - q_{0,n} = \frac{2K_n - 2}{2K_n - 1} = 1 - \mathcal{O}(K_n^{-1}) \quad \text{and} \quad 1 - q_{1,n} = \frac{1}{2}. \tag{5.30}$$

Moreover, since

$$r_n(t) - r_n(t-1) = K_n - \left\lceil \frac{t}{2} \right\rceil, \tag{5.31}$$

we have, for all $t \leq \ell_n$,

$$\begin{aligned}
 1 - q_{t,n} &= \frac{r_n(t) - r_n(t-1)}{r_n(t)} \\
 &= \frac{K_n - \left\lfloor \frac{t}{2} \right\rfloor}{(t+1)K_n - \left\lfloor \frac{t+1}{2} \right\rfloor \left\lfloor \frac{t+1}{2} \right\rfloor} \\
 &= \frac{1}{t+1} \cdot \frac{1 - \Theta(t/K_n)}{1 - \Theta(t/K_n)} \\
 &= \frac{1}{t+1} \cdot \left(1 + \mathcal{O}\left(\frac{\ell_n}{K_n}\right) \right).
 \end{aligned} \tag{5.32}$$

Hence, by (5.30) and (5.32), for all $t \leq \ell_n$,

$$\begin{aligned}
 \prod_{j=0}^t (1 - q_{j,n}) &= (1 + o(1)) \cdot \frac{1}{2} \cdot \prod_{j=2}^t \left[\frac{1}{j+1} \cdot \left(1 + \mathcal{O}\left(\frac{\ell_n}{K_n}\right) \right) \right] \\
 &= \frac{1}{(t+1)!} \left(1 + \mathcal{O}\left(\frac{\ell_n^2}{K_n}\right) \right) \\
 &= (1 + o(1)) \frac{1}{(t+1)!},
 \end{aligned} \tag{5.33}$$

where in the last two steps we used that, by definition, $\ell_n = o(\sqrt{K_n})$. Hence

$$\left(1 + \mathcal{O}\left(\frac{\ell_n}{K_n}\right) \right)^{t-1} \leq \left(1 + \mathcal{O}\left(\frac{\ell_n}{K_n}\right) \right)^{\ell_n} \leq \left(1 + \mathcal{O}\left(\frac{\ell_n^2}{K_n}\right) \right) = (1 + o(1)), \quad 2 \leq t \leq \ell_n. \tag{5.34}$$

Notice that (5.33) implies

$$\sum_{t=0}^{\ell_n} \prod_{j=0}^t (1 - q_{j,n}) = (1 + o(1)) \sum_{t=0}^{\ell_n} \frac{1}{(t+1)!} = (1 + o(1))(e - 1), \tag{5.35}$$

and (5.29) follows by taking the limit as $n \rightarrow \infty$. At this point (4.17) follows from (5.26) and (5.29).

We now prove (4.18). Call H_n the distribution function of τ_n^{NE} . By (5.33), we have, for $t \leq \ell_n$,

$$H_n(t) = 1 - \mathbb{P}(\tau_n^{\text{NE}} > t) = (1 + o(1)) \cdot \left(1 - \frac{1}{(t+1)!} \right). \tag{5.36}$$

In what follows, we will use the following lemma, whose proof can be found, for instance, in Ogryczak and Ruszczyński (1999, corollary 3).

Lemma 5.1. *Let X be a nonnegative random variable with finite expectation μ , finite variance, and distribution function G . Define the function $\Phi(x) := \int_0^x G(t) dt$. Then*

$$\frac{1}{2} \text{Var}[X] = \int_0^{+\infty} (\Phi(t) - [t - \mu]_+) dt, \tag{5.37}$$

where $[t - \mu]_+ = \max\{t - \mu, 0\}$.

We now go back to the proof of (4.18). Since the random variable τ_n^{NE} is bounded, we can apply Lemma 5.1 to get

$$\text{Var}[\tau_n^{\text{NE}}] = 2 \int_0^\mu \Phi_n(t) dt + 2 \int_\mu^{+\infty} [\Phi_n(t) - (t - \mu)] dt, \tag{5.38}$$

where, thanks to (5.36), for all $t \leq \ell_n$,

$$\Phi_n(t) = \int_0^t H_n(s) ds \sim \int_0^t \left(1 - \frac{1}{[s+1]!} \right) ds. \tag{5.39}$$

By explicit numerical integration, using (5.39) and the fact that $\mu := \mathbb{E}[\tau_n^{\text{NE}}] \sim e - 1 \leq \ell_n$ for all n sufficiently large, we get

$$2 \int_0^\mu \Phi_n(t) dt \sim 2 \int_0^\mu dt \int_0^t \left(1 - \frac{1}{\lfloor s+1 \rfloor!}\right) ds \approx 0.258. \tag{5.40}$$

Moreover, for all $t \geq 0$,

$$\begin{aligned} \Phi_n(t) - (t - \mu) &= \int_0^t (H_n(s) - 1) ds + \int_0^{+\infty} (1 - H_n(s)) ds \\ &= \int_t^{+\infty} (1 - H_n(s)) ds \\ &= \int_t^{\ell_n \vee t} (1 - H_n(s)) ds + \int_{\ell_n \vee t}^{+\infty} (1 - H_n(s)) ds. \end{aligned} \tag{5.41}$$

Since the sequence $q_{t,n}$, defined in (4.12), is increasing in t and $q_{1,n} = 1/2$, by (5.27) we get, for all $s \geq 0$

$$1 - H_n(s) = P(\tau_n^{NE} > s) \leq \frac{1}{2^s}. \tag{5.42}$$

Hence, for all $t \geq 0$,

$$\int_{\ell_n \vee t}^{+\infty} (1 - H_n(s)) ds \leq \frac{1}{\log 2} 2^{-(\ell_n \vee t)}, \tag{5.43}$$

which goes to zero when $n \rightarrow \infty$. Therefore, by (5.36), (5.41), and (5.43), we conclude that

$$\Phi_n(t) - (t - \mu) = \begin{cases} (1 + o(1)) \int_t^{\ell_n} \frac{1}{\lfloor s+1 \rfloor!} ds + \mathcal{O}(2^{-\ell_n}) & \text{if } t \leq \ell_n, \\ \mathcal{O}(2^{-t}) & \text{if } t > \ell_n. \end{cases} \tag{5.44}$$

It is now convenient to split the second integral in (5.38) as follows

$$2 \int_\mu^{+\infty} [\Phi_n(t) - (t - \mu)] dt = 2 \int_\mu^{\ell_n} [\Phi_n(t) - (t - \mu)] dt + 2 \int_{\ell_n}^{+\infty} [\Phi_n(t) - (t - \mu)] dt. \tag{5.45}$$

At this point, using (5.44), we can bound the second integral on the right hand side of (5.45) as follows

$$\int_{\ell_n}^{\infty} [\Phi_n(t) - (t - \mu)] dt = \int_{\ell_n}^{+\infty} \mathcal{O}(2^{-t}) dt = \mathcal{O}(2^{-\ell_n}). \tag{5.46}$$

On the other hand, the first integral on the right hand side of (5.45) can be bounded by

$$\begin{aligned} \int_\mu^{\ell_n} [\Phi_n(t) - (t - \mu)] dt &= (1 + o(1)) \int_\mu^{\ell_n} \int_t^{\ell_n} \frac{1}{\lfloor s+1 \rfloor!} ds dt + \mathcal{O}(\ell_n 2^{-\ell_n}) \\ &= (1 + o(1)) \int_\mu^{\ell_n} \int_t^{\ell_n} \frac{1}{\lfloor s+1 \rfloor!} ds dt \\ &= (1 + o(1)) \int_\mu^{\infty} \int_t^{\infty} \frac{1}{\lfloor s+1 \rfloor!} ds dt. \end{aligned} \tag{5.47}$$

In conclusion, by (5.45), (5.46), and (5.47) and numerical integration, we have

$$2 \int_\mu^{+\infty} [\Phi_n(t) - (t - \mu)] dt \sim 2 \int_\mu^{+\infty} dt \int_t^{+\infty} \frac{1}{\lfloor s+1 \rfloor!} ds \approx 0.509. \tag{5.48}$$

The combination of (5.38), (5.40), and (5.48) gives

$$\text{Var}[\tau_n^{NE}] \approx 0.767. \quad \square$$

Proofs of Section 4.3

Proof of Theorem 4.5. First observe that either $\tau_n^{NE} \leq 2K_n - 1$ or $\tau_n^{NE} = \infty$. Call

$$C_n(t) := \{\text{BRD}_n(t) \in \text{NE}_n\}, \text{ and } D_n(t) := C_n(t)^c \cap \{\text{BRD}_n(t) \notin R_n(t-2)\}, \tag{5.49}$$

where $R_n(t)$ is defined as in (4.8). Therefore, the statement in (4.19) is equivalent to

$$\lim_{n \rightarrow \infty} P(C_n(2K_n - 1)) = 0. \tag{5.50}$$

Notice that

$$P(C_n(0)) = \frac{1}{K_n^A K_n^B} \text{ and } P(C_n(1)) = P(C_n(0)) + P(C_n(1) | C_n(0)^c)P(C_n(0)^c). \tag{5.51}$$

We have

$$\begin{aligned} P(C_n(1) | C_n(0)^c) &= \sum_{a=1}^{K_n^A} P(C_n(1) | C_n(0)^c \cap \text{BRD}_n(1) = (a, 1))P(\text{BRD}_n(1) = (a, 1) | C_n(0)^c) \\ &= \sum_{a=2}^{K_n^A} P(C_n(1) | C_n(0)^c \cap \text{BRD}_n(1) = (a, 1))P(\text{BRD}_n(1) = (a, 1) | C_n(0)^c) \\ &= \sum_{a=2}^{K_n^A} \frac{1}{K_n^B} P(\text{BRD}_n(1) = (a, 1) | C_n(0)^c) \\ &= \frac{1}{K_n^B} (1 - P(\text{BRD}_n(1) = (1, 1) | C_n(0)^c)) \\ &= \frac{1}{K_n^B} \left(1 - \frac{P(C_n(0)^c | \text{BRD}_n(1) = (1, 1))P(\text{BRD}_n(1) = (1, 1))}{P(C_n(0)^c)} \right) \\ &= \frac{1}{K_n^B} \left(1 - \frac{\frac{K_n^B - 1}{K_n^B} \frac{1}{K_n^A}}{1 - \frac{1}{K_n^A K_n^B}} \right) = \frac{K_n^A - 1}{K_n^A K_n^B - 1}. \end{aligned} \tag{5.52}$$

This implies

$$P(C_n(1)) = \frac{1}{K_n^A K_n^B} + \frac{K_n^A - 1}{K_n^A K_n^B - 1} \left(1 - \frac{1}{K_n^A K_n^B} \right) = \frac{1}{K_n^B}. \tag{5.53}$$

Notice that

$$C_n(t) \cap C_n^c(t-1) = \{\text{BRD}_n(t-1) \neq \text{BRD}_n(t) = \text{BRD}_n(t+1)\} \subseteq \{\text{BRD}_n(t) \notin R_n(t-2)\}. \tag{5.54}$$

Hence, by definition of $D_n(t-1)$, we get

$$C_n(t) \cap C_n^c(t-1) = C_n(t) \cap D_n(t-1). \tag{5.55}$$

Moreover, since

$$C_n(t-1) \subseteq C_n(t), \tag{5.56}$$

by (5.55) and (5.56) we conclude that, for $t \geq 1$,

$$P(C_n(t)) = P(C_n(t-1)) + P(C_n(t) | D_n(t-1))P(D_n(t-1)). \tag{5.57}$$

By explicit computation (see Fig. 2), we get

$$P(C_n(t) | D_n(t-1)) = \begin{cases} \frac{1}{K_n^A} \frac{K_n^B - \lfloor \frac{t}{2} \rfloor}{K_n^B - 1} & \text{for } t \text{ even,} \\ \frac{1}{K_n^B} \frac{K_n^A - \lfloor \frac{t}{2} \rfloor}{K_n^A - 1} & \text{for } t \text{ odd.} \end{cases} \tag{5.58}$$

On the other hand, since

$$D_n(t) \subseteq D_n(t-1), \tag{5.59}$$

we have




Fig. 2. The figure on the left shows an instance of the first five steps of the BRD, whereas the one on the right considers an additional step. Given the position of $BRD(t)$ for $t = 1, \dots, 5$ in the figure on the left, $D_n(5)$ coincides with the event that the payoff of player B at 5 is not the maximum of its row (the blue dashed line). Conditioning on $D_n(5)$, the probability of $D_n(6)$ is the product between the probability that BRD(6) is not at the action profiles M_1 and M_2 , that is $(K_n^B - 3)/(K_n^B - 1)$, and, given this, the probability that the payoff of player A at 6 is not the maximum of its column (the red dashed line), i.e., $(K_n^A - 1)/K_n^A$. This explains (5.61) when $t = 6$. Similarly, conditioning on $D_n(5)$, $C_n(6)$ is the intersection between two events: the first one is that the position of BRD(6) does not coincide with M_1 and M_2 , which has probability $(K_n^B - 3)/(K_n^B - 1)$; the second one is that the payoff of player A at 6 is the maximum of its column (the red dashed line), which, conditioning on the first event, has probability $1/K_n^A$. This justifies (5.58) when $t = 6$.

$$P(D_n(t)) = P(D_n(t) \cap D_n(t-1)) = P(D_n(t-1))P(D_n(t) | D_n(t-1)). \tag{5.60}$$

The latter conditional probability can be explicitly computed (see Fig. 2), obtaining

$$P(D_n(t) | D_n(t-1)) = \begin{cases} \frac{K_n^A - 1}{K_n^A} \frac{K_n^B - \lfloor \frac{t}{2} \rfloor}{K_n^B - 1} & \text{for } t \text{ even,} \\ \frac{K_n^B - 1}{K_n^B} \frac{K_n^A - \lfloor \frac{t}{2} \rfloor}{K_n^A - 1} & \text{for } t \text{ odd.} \end{cases} \tag{5.61}$$

By iterating (5.60) and (5.61), we deduce that, for all $t \leq 2K_n - 1$ odd,

$$\begin{aligned} P(D_n(t)) &= P(D_n(t-2)) \frac{K_n^A - \lfloor \frac{t}{2} \rfloor}{K_n^A} \frac{K_n^B - \lfloor \frac{t}{2} \rfloor}{K_n^B} \\ &= \frac{K_n^B - 1}{K_n^B} \prod_{j=1}^{\lfloor \frac{t}{2} \rfloor} \frac{K_n^B - j}{K_n^B} \frac{K_n^A - j}{K_n^A} \\ &\leq \prod_{j=1}^{\lfloor \frac{t}{2} \rfloor} \left(1 - \frac{j}{K_n}\right) \leq \exp\left(-\sum_{j=1}^{\lfloor \frac{t}{2} \rfloor} \frac{j}{K_n}\right) \leq \exp\left(-\frac{(t-1)^2}{4K_n}\right), \end{aligned} \tag{5.62}$$

where we have used the inequality $1 - x \leq e^{-x}$.

Moreover, by (5.61), $P(D_n(t))$ is decreasing in t and, by (5.58), $P(C_n(t) | D_n(t-1)) \leq 1/K_n$ for all t . Hence, iterating (5.57), we can write

$$\begin{aligned} P(C_n(t)) &= P(C_n(t-2)) + P(C_n(t-1) | D_n(t-2))P(D_n(t-2)) \\ &\quad + P(C_n(t) | D_n(t-1))P(D_n(t-1)) \\ &\leq P(C_n(t-2)) + \frac{2}{K_n}P(D_n(t-2)). \end{aligned} \tag{5.63}$$

By applying the estimate in (5.62) to $P(D_n(t-2))$, we obtain, for t odd,

$$P(C_n(t)) \leq P(C_n(t-2)) + \frac{2}{K_n} \exp\left(-\frac{(t-3)^2}{4K_n}\right), \tag{5.64}$$

whose iteration leads to

$$P(C_n(t)) \leq P(C_n(1)) + \frac{2}{K_n} \sum_{\substack{j=1 \\ j \text{ odd}}}^{t-3} \exp\left(-\frac{j^2}{4K_n}\right) \leq \frac{1}{K_n} + \frac{2}{K_n} \sum_{\substack{j=1 \\ j \text{ odd}}}^{t-3} \exp\left(-\frac{j^2}{4K_n}\right), \tag{5.65}$$

where in the last inequality we have used (5.53).

Taking $t = 2K_n - 1$ and using the estimate $\sum_{j=0}^H \exp\{-j^2/H\} = \mathcal{O}(\sqrt{H})$ for $H \rightarrow \infty$, we conclude that

$$P(C_n(2K_n - 1)) \leq \mathcal{O}\left(\frac{1}{\sqrt{K_n}}\right), \tag{5.66}$$

which implies (4.19). \square

Proofs of Section 4.4

We now prove the following result, and then show that Proposition 4.7 immediately follows from it.

Proposition 5.2. Fix a positive sequence p_n . Then, for every sequence ℓ_n such that

$$\lim_{n \rightarrow \infty} \ell_n = \infty, \quad \lim_{n \rightarrow \infty} \frac{\log(p_n)}{\log(\ell_n)} = 0, \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{\ell_n}{\sqrt{K_n^A \wedge K_n^B}} = 0, \tag{5.67}$$

we have

$$\lim_{n \rightarrow \infty} P(\tau_n^{\text{NE}} < \ell_n) = 1. \tag{5.68}$$

The proof of Proposition 5.2 relies on the following lemma.

Lemma 5.3. Given a sequence $\{X_i\}_{i \in \mathbb{N}_+}$ of i.i.d. random variables having a uniform distribution on $[0, 1]$, consider the event

$$L = \{X_1 < \max\{X_1, \dots, X_k\}\}. \tag{5.69}$$

Then for all $x \in [0, 1]$ we have

$$P(X_1 \leq x \mid L) \geq x = P(X_1 \leq x). \tag{5.70}$$

Proof. The conditional distribution of X_1 , given L^c , is Beta($k, 1$), i.e.,

$$P(X_1 \leq x \mid L^c) = x^k \leq x. \tag{5.71}$$

Then

$$\begin{aligned} P(X_1 \leq x \mid L) &= \frac{P(X_1 \leq x) - P(X_1 \leq x \mid L^c)P(L^c)}{P(L)} \\ &= \frac{x - x^k P(L^c)}{1 - P(L^c)} \\ &\geq \frac{x - x P(L^c)}{1 - P(L^c)} = x, \end{aligned} \tag{5.72}$$

where the last inequality is due to (5.71). \square

Proof of Proposition 5.2. Define the events

$$E_n(t) := \{\text{BRD}_n(t) \in R_n(t-2)\}, \quad \text{for } t \geq 2, \tag{5.73}$$

$$Z_n(t) := \{\text{BRD}_n(t) \neq \text{BRD}_n(t-1)\}, \quad \text{for } t \geq 1. \tag{5.74}$$

In words, $E_n(t)$ represents the event that at time t the process $\text{BRD}_n(t)$ visits a previously visited row or column, that is, $E_n(t) = \{\tau_n^R \leq t\}$, whereas $Z_n(t)$ represents the event that $\text{BRD}_n(t-1)$ is not a PNE. Therefore, the sequence $E_n(t)$ is increasing in t , i.e.,

$$E_n(t-1) \subseteq E_n(t), \tag{5.75}$$

whereas the sequence $Z_n(t)$ is decreasing in t , i.e.,

$$Z_n(t - 1) \supseteq Z_n(t). \tag{5.76}$$

Then

$$\begin{aligned} P(E_n(t) \cap Z_n(t)) &= P(E_n(t-1) \cap E_n(t) \cap Z_n(t)) + P(E_n(t-1)^c \cap E_n(t) \cap Z_n(t)) \\ &= P(E_n(t-1) \cap Z_n(t)) + P(E_n(t-1)^c \cap E_n(t) \cap Z_n(t)) \\ &\leq P(E_n(t-1) \cap Z_n(t-1)) + P(E_n(t-1)^c \cap E_n(t) \cap Z_n(t)) \\ &\leq P(E_n(t-1) \cap Z_n(t-1)) + P(E_n(t) | E_n(t-1)^c \cap Z_n(t)), \end{aligned} \tag{5.77}$$

where the first equality is just the law of total probabilities, the second derives from (5.75), the first inequality is a consequence of (5.76), and the last stems from the definition of conditional probability. Moreover, we claim that, for $t \geq 3$, we have

$$P(E_n(t) | E_n(t-1)^c \cap Z_n(t)) \leq \begin{cases} \frac{\lfloor \frac{t}{2} \rfloor - 1}{K_n^B - 1} & \text{if } t \text{ is even,} \\ \frac{\lfloor \frac{t}{2} \rfloor - 1}{K_n^A - 1} & \text{if } t \text{ is odd.} \end{cases} \tag{5.78}$$

The conditioning event on the l.h.s. of (5.78) represents the fact that $\text{BRD}_n(t-1)$ is neither a PNE nor an element of $R_n(t-3)$. Therefore (5.78) provides a bound for the conditional probability that $\text{BRD}_n(t)$ is an element of $R_n(t-2)$. To see why the bound holds, start considering the case $p_n = 0$, where the inequality in (5.78) holds as an equality. This is due to the fact that all payoffs are i.i.d.. On the other extreme, when $p_n = 1$, the left hand side equals zero, since potential games do not admit traps. In the intermediate case when $p_n \in (0, 1)$, the payoffs in the row (column) of interest are not i.i.d.. Consider an A payoff in a previously visited row; with probability $1 - p_n$ it is uniformly distributed on $[0, 1]$ and with probability p_n it has the law of a uniform random variable, conditioned on not being the largest payoff in its row. A similar argument holds for B payoffs, replacing row with column. By Lemma 5.3, the distribution of a B payoff on a previously visited row (A payoff on a previously visited column) is stochastically dominated by a uniform distribution on $[0, 1]$. This proves the inequality in (5.78).

Iterating (5.77) we obtain

$$\begin{aligned} P(E_n(t) \cap Z_n(t)) &\leq P(E_n(2) \cap Z_n(2)) + \sum_{i=3}^t \frac{\lfloor \frac{i}{2} \rfloor - 1}{K_n^A \wedge K_n^B - 1} \\ &\leq \sum_{i=1}^t \frac{i}{K_n^A \wedge K_n^B - 1} \leq \frac{t^2}{K_n^A \wedge K_n^B - 1}, \end{aligned} \tag{5.79}$$

where for the last bound we used the fact that $P(E_n(2) \cap Z_n(2)) = 0$ (since $\{E_n(2)\} = \{Z_n(2)^c\}$) and the fact $\sum_{i=1}^t i = t(t+1)/2$. Call τ_n^{cycle} the stopping time

$$\tau_n^{\text{cycle}} := \inf \{t \geq 4 : E_n(t) \cap Z_n(t) \text{ holds}\}, \tag{5.80}$$

that is, τ_n^{cycle} is the first time that the BRD re-visits an element of a trap. Hence, (5.79) says that for all $p_n \in [0, 1]$, if $\ell_n = o(\sqrt{K_n^A \wedge K_n^B})$, then, as $n \rightarrow \infty$,

$$P(\tau_n^{\text{cycle}} \leq \ell_n) \rightarrow 0. \tag{5.81}$$

By definition of τ_n^{cycle} , the stopping time τ_n^R defined in (4.5) can be rewritten as

$$\tau_n^R := \begin{cases} \tau_n^{\text{cycle}} \wedge (\tau_n^{\text{NE}} + 1), & \text{if } (1, 1) \notin \text{NE}_n, \\ 2, & \text{if } (1, 1) \in \text{NE}_n. \end{cases} \tag{5.82}$$

Notice that, for every sequence $\ell_n = o(\sqrt{K_n^A \wedge K_n^B})$, if we show that

$$P(\tau_n^R \leq \ell_n) \rightarrow 1, \tag{5.83}$$

then from (5.81) and (5.82) it follows that

$$P(\tau_n^{\text{NE}} \leq \ell_n) \rightarrow 1. \tag{5.84}$$

Let

$$S_n := \{(a, b) \in [K_n^A] \times [K_n^B] : U_n^A(a, b) = U_n^B(a, b)\}, \tag{5.85}$$




Fig. 3. The left figure shows an instance of the first four steps of the BRD. The right figure shows how the dynamics proceeds after time 4. The numbered action profiles lying on the dashed lines, i.e., 5, 6, and 7, give the same payoff to the row and column player. Note that in this case the event $J_n^{s_n, \ell_n}$ occurs with $\ell_n = 7$ and $s_n = 3$. Indeed, there exists $t \leq \ell_n - s_n$ (in this case $t = 4$) such that the BRD visits only action profiles in S_n for s_n consecutive steps. As a consequence, the payoff of the row player at the action profile 7 is the maximum of the payoffs of the row player in the action profiles lying on the red dashed lines and of the payoffs of the column player in the action profiles lying on the blue dashed lines. Such payoffs are all i.i.d. $\text{Unif}([0, 1])$ except for the payoffs associated to the action profiles $4, V_1, \dots, V_5$, for which we have additional information. Since the number of such exceptional action profiles is at most $(\ell_n/2)^2$ and the total number of action profiles lying on the dashed lines is $r_n(s_n)$, we have that the payoff of the row player at the action profile 7 is the maximum of at least u i.i.d. $\text{Unif}([0, 1])$ random variables, where u is defined as in (5.88).

be the set of action profiles that give the same payoff to the two players. Fix now a sequence $\ell_n = o(\sqrt{K_n^A \wedge K_n^B})$ and define, for every integer $s \in \{1, \dots, \ell_n - 1\}$,

$$J_n^{s, \ell_n} := \{ \exists t \leq \ell_n - s \text{ s.t. } \text{BRD}_n(t + s') \in S_n, \quad \forall s' \in \{0, \dots, s\} \}. \tag{5.86}$$

The event $J_n^{s_n, \ell_n}$ occurs if there exists an interval of s consecutive steps before ℓ_n in which the BRD visits only elements in S_n .

For every sequence $(s_n)_{n \in \mathbb{N}}$ such that $s_n \leq \ell_n$ for every n , we have

$$\mathbb{P}(\tau_n^R > \ell_n) = \mathbb{P}(\{ \tau_n^R > \ell_n \} \cap J_n^{s_n, \ell_n}) + \mathbb{P}(\{ \tau_n^R > \ell_n \} \cap \{ J_n^{s_n, \ell_n} \}^c). \tag{5.87}$$

First we show that the first term on the r.h.s. of (5.87) goes to zero as $s_n \rightarrow \infty$. To this end, it is enough to show that, under the event $J_n^{s_n, \ell_n}$, there exists some $t \leq \ell_n - s_n$ such that the best-responding player’s payoff at time $t + s_n$ is stochastically larger than a $\text{Beta}(u, 1)$ random variable, with

$$u := r_n(s_n) - \frac{\ell_n^2}{4}. \tag{5.88}$$

Notice that, for $n \in \mathbb{N}$ large enough,

$$u \geq \frac{s_n}{2}(K_n^A + K_n^B) - \frac{\ell_n^2}{2} \geq \frac{s_n}{3}(K_n^A + K_n^B), \tag{5.89}$$

where in the first inequality we used (4.8) and the fact that $s_n \leq \ell_n$; the second inequality holds for all sufficiently large n by our choice of ℓ_n . Indeed, after the interval of s_n consecutive steps in which the BRD visits only elements of S_n , the BRD visits an action profile (a, b) such that

- $(a, b) \in S_n$, that is, both players receive the same payoff;
- this common payoff is the largest of a set of random variables, among which at least u are i.i.d. $\text{Unif}([0, 1])$ (see Fig. 3 for more details);

hence, the stochastic domination follows.

Therefore, by Proposition B.1, the probability that $\text{BRD}_n(t + s_n) = \text{BRD}_n(t + s_n + 1)$, that is, $\text{BRD}_n(t + s_n)$ is a PNE, is bounded from below by the probability that a $\text{Beta}(u, 1)$ random variable is larger than the maximum of $K_n^A \vee K_n^B - 1$ i.i.d. random variables with a uniform distribution on $[0, 1]$. Hence, by Proposition B.2, we get

$$P\left(\{\tau_n^R > \ell_n\} \cap J_n^{s_n, \ell_n}\right) \leq 1 - \frac{\left\lfloor \frac{s_n}{3} \right\rfloor (K_n^A + K_n^B)}{\left\lfloor \frac{s_n}{3} \right\rfloor (K_n^A + K_n^B) + (K_n^A \vee K_n^B - 1)} \leq \frac{3}{s_n + 3}, \tag{5.90}$$

which goes to zero as $s_n \rightarrow \infty$. We now show that, under the assumption in (5.67), it is possible to find a sequence $(s_n)_n$ such that

$$\lim_{n \rightarrow \infty} s_n = \infty, \quad \lim_{n \rightarrow \infty} \frac{s_n}{\ell_n} = 0, \tag{5.91}$$

and

$$\lim_{n \rightarrow \infty} P\left(\{\tau_n^R > \ell_n\} \cap \left\{J_n^{s_n, \ell_n}\right\}^c\right) \rightarrow 0. \tag{5.92}$$

Notice that under the event $\{\tau_n^R > \ell_n\}$, the probability of the event $\left\{J_n^{s_n, \ell_n}\right\}^c$ can be upper bounded by the probability that, splitting the interval $\{0, \dots, \ell_n\}$ into subintervals of length s_n , none of them is such that the BRD visits only S_n in that subinterval. Therefore,

$$P\left(\{\tau_n^R > \ell_n\} \cap \left\{J_n^{s_n, \ell_n}\right\}^c\right) \leq P\left(\text{Binomial}\left(\left\lfloor \frac{\ell_n}{s_n} \right\rfloor, p_n^{s_n}\right) = 0\right) = (1 - p_n^{s_n})^{\lfloor \ell_n/s_n \rfloor}. \tag{5.93}$$

If $\ell_n = o(\sqrt{K_n^A \wedge K_n^B})$, then the term on the r.h.s. of (5.93) goes to zero whenever

$$\lim_{n \rightarrow \infty} \frac{\ell_n}{s_n} p_n^{s_n} = \infty. \tag{5.94}$$

A necessary and sufficient condition for (5.94) is

$$\lim_{n \rightarrow \infty} \log(\ell_n) - \log(s_n) - s_n \log(p_n^{-1}) = \infty, \tag{5.95}$$

or, equivalently,

$$\lim_{n \rightarrow \infty} \log(\ell_n) \left[1 - \frac{\log(s_n)}{\log(\ell_n)} - \frac{s_n \log(p_n^{-1})}{\log(\ell_n)}\right] = \infty. \tag{5.96}$$

Thanks to (5.67), we can choose, e.g.,

$$s_n = \log(\ell_n) \wedge \log\left(\frac{\log(\ell_n)}{\log(p_n^{-1})}\right) \rightarrow \infty. \tag{5.97}$$

Since, for every diverging positive sequence $(a_n)_{n \geq 0}$, it holds that $\log(a_n)/a_n \rightarrow 0$, by the definition in (5.97) we deduce that

$$\frac{\log(s_n)}{\log(\ell_n)} \leq \frac{\log(\log \ell_n)}{\log(\ell_n)} \rightarrow 0 \quad \text{as } n \rightarrow \infty, \tag{5.98}$$

and

$$\frac{s_n \log(p_n^{-1})}{\log(\ell_n)} \leq \frac{\log\left(\frac{\log(\ell_n)}{\log(p_n^{-1})}\right)}{\log(\ell_n)} \rightarrow 0 \quad \text{as } n \rightarrow \infty. \tag{5.99}$$

Coupling (5.98) and (5.99), we immediately validate (5.96). \square

Proof of Proposition 4.7. Assume that (4.20) is satisfied. Choose any sequence ℓ_n such that

$$\lim_{n \rightarrow \infty} \frac{\log(p_n)}{\log(\ell_n)} = 0. \tag{5.100}$$

There are two cases:

- If

$$\lim_{n \rightarrow \infty} \frac{\ell_n}{\sqrt{K_n^A \wedge K_n^B}} = 0,$$

then (4.23) follows by Proposition 5.2.

• If instead

$$\limsup_{n \rightarrow \infty} \frac{\ell_n}{\sqrt{K_n^A \wedge K_n^B}} > 0, \tag{5.101}$$

then we can define the $\ell'_n = \ell_n \wedge [K_n^A \wedge K_n^B]^{1/3}$. Notice that

- $\lim_{n \rightarrow \infty} \ell'_n = \infty$;
- by the fact that $\log(\ell'_n) \geq \frac{1}{3} \log(K_n^A \wedge K_n^B) \wedge \log(\ell_n)$, combined with (4.20) and (5.100), we deduce that

$$\lim_{n \rightarrow \infty} \frac{\log(p_n)}{\log(\ell'_n)} = 0;$$

- moreover, by the definition of ℓ'_n we have

$$\lim_{n \rightarrow \infty} \frac{\ell'_n}{\sqrt{K_n^A \wedge K_n^B}} = 0.$$

Hence, thanks to Proposition 5.2 we get

$$\lim_{n \rightarrow \infty} P(\tau_n^{NE} < \ell'_n) = 1. \tag{5.102}$$

To conclude the proof, it is enough to see that for every $n \in \mathbb{N}$

$$P(\tau_n^{NE} < \ell_n) \geq P(\tau_n^{NE} < \ell'_n). \quad \square$$

6. Conclusions and open problems

We have considered a model of two-person games with random payoffs that parametrically interpolates potential games and games with i.i.d. payoffs. The interpolation acts locally on each payoff profile. We have studied both the asymptotic behavior of the random number of pure Nash equilibria of the game and the asymptotic behavior of best response dynamics, as the number of actions for each player diverges. The type of model that we chose requires combinatorial tools for its analysis.

We see this paper as a first attempt to provide a parametric model for random games where the payoffs are not independent, but have some structure that depends on a locally acting parameter. Several extensions and variations of this model are conceivable and will be the object of our future research. For instance:

- (i) It would be interesting to have a clearer view of the phase transition taking place at $p = 0$. In particular, it would be important to investigate the existence of a sequence $p_n \rightarrow 0$ such that the probability that a BRD does not lead to a PNE converges to a value smaller than 1.
- (ii) Games with more than two players could be studied.
- (iii) With more than two players, different types of deviator rules in BRD could be considered, e.g., round-robin, random order, etc.
- (iv) The behavior of better-response dynamics could be studied and compared to best response dynamics, along the lines of Amiet et al. (2021a).
- (v) When we deal with the number of pure Nash equilibria, we studied a form of Law of Large Numbers. The existence of a Central Limit Theorem could be explored.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) did not use any AI or AI assisted technologies.

Acknowledgments

The authors deeply thank three reviewers for their extremely careful reading of the manuscript and their insightful suggestions. Hlafo Alfie Mimun and Marco Scarsini are members of GNAMPA-INdAM.

Funding

Hlafo Alfie Mimun and Marco Scarsini’s research was supported by the GNAMPA project CUP_E53C22001930001, the MIUR PRIN PNRR grant P2022XT8C8, and the MIUR PRIN grant 2022EKNE5K. Matteo Quattropani thanks the German Research Foundation (project number 444084038, priority program SPP2265) for financial support.

Appendix A. List of symbols

| | |
|-------------------------|--|
| A_a | {player A’s best response to action 1 is a }, defined in (5.16) |
| BRD | best response dynamics |
| $C_n(t)$ | $\{\text{BRD}_n(t) \in \text{NE}_n\}$, defined in (5.49) |
| $D_n(t)$ | $\{\text{BRD}_n(t) \notin \text{NE}_n, \text{BRD}_n(t) \notin R_n(t-2)\}$, defined in (5.49) |
| $E_n(t)$ | $\{\text{BRD}(t) \in R_n(t-2)\}$, defined in (5.73) |
| F | uniform distribution function on $[0, 1]$ |
| G_t^π | $\{\text{BRD}_n(s) = \pi_s \text{ for } 0 \leq s \leq t\}$, defined in (5.19) |
| H_n | distribution function of τ_n^{NE} |
| J_n^{s, ℓ_n} | $\{\exists t \leq \ell_n - s \text{ s.t. } \text{BRD}_n(t + s') \in S_n, \forall s' \in \{0, \dots, s\}\}$, defined in (5.86) |
| K_n^A | number of player A’s actions in the game U_n |
| K_n^B | number of player B’s actions in the game U_n |
| K_n | $\min(K_n^A, K_n^B)$, defined in (3.2) |
| $[K^A]$ | action set of player A |
| $[K^B]$ | action set of player B |
| L | $\{X_1 < \max\{X_1, \dots, X_k\}\}$, defined in (5.69) |
| M | trap, defined in Definition 2.1 |
| NE | set of pure Nash equilibria |
| p_n | probability that $U^A(a, b) = U^B(a, b)$ in the game U_n |
| $q_{t,n}$ | $P(\tau_n^{\text{NE}} = t \mid \tau_n^{\text{NE}} \geq t)$, defined in (4.12) |
| $r_n(t)$ | $\left\lfloor \frac{t+1}{2} \right\rfloor K_n^A + \left\lfloor \frac{t+1}{2} \right\rfloor K_n^B - \left\lfloor \frac{t+1}{2} \right\rfloor \left\lfloor \frac{t+1}{2} \right\rfloor$, defined in (4.8) |
| $R_n(t)$ | defined in (4.3) |
| S_n | $\{(a, b) \in [K_n^A] \times [K_n^B] : U_n^A(a, b) = U_n^B(a, b)\}$, defined in (5.85) |
| t | (discrete) time |
| U^A | player A’s payoff function |
| U^B | player B’s payoff function |
| W_n | number of pure Nash equilibria in U_n |
| $Z_n(t)$ | $\{\text{BRD}_n(t) \neq \text{BRD}_n(t-1)\}$, defined in (5.74) |
| Π_t | set of possible paths for BRD $_n$ up to time t |
| τ_n^{NE} | first time the BRD visits a PNE, defined in (4.2) |
| τ_n^{cycle} | first time the BRD re-visits an element of a trap, defined in (5.80) |
| τ_n^R | $\min\{t \geq 2 : \text{BRD}_n(t) \in R_n(t-2)\}$, defined in (4.5) |
| $\Phi_n(t)$ | $\int_0^t H_n(s) ds$, defined in (5.39) |
| Ψ | potential function, defined in (2.4) |

Appendix B. Beta distribution

We report two well-known results about Beta distributions. For the sake of completeness, we add their simple proofs.

Proposition B.1. *Let X_1, \dots, X_k be i.i.d. random variables having a uniform distribution on $[0, 1]$ and let $M_k := \max_{i \in \{1, \dots, k\}} X_i$. Then M_k has distribution $\text{Beta}(k, 1)$.*

Proof. For any $t \in [0, 1]$ we have

$$P(M_k \leq t) = [P(X_1 \leq t)]^k = t^k,$$

i.e., a $\text{Beta}(k, 1)$ distribution function. \square

Proposition B.2. *Let X and Y be independent random variables with distributions $\text{Beta}(a, 1)$ and $\text{Beta}(b, 1)$, respectively. Then*

$$P(X > Y) = \frac{a}{a+b}. \tag{B.1}$$

Proof. We have

$$P(X > Y) = \int_0^1 \left(\int_0^t at^{a-1} \cdot bs^{b-1} ds \right) dt = \int_0^1 at^{a-1} \cdot t^b dt = \frac{a}{a+b}. \quad \square$$

References

- Alon, N., Rudov, K., Yariv, L., 2021. Dominance solvability in random games. Technical report arXiv:2105.10743.
- Amiet, B., Collevocchio, A., Hamza, K., 2021a. When “better” is better than “best”. *Oper. Res. Lett.* 49 (2), 260–264.
- Amiet, B., Collevocchio, A., Scarsini, M., Zhong, Z., 2021b. Pure Nash equilibria and best-response dynamics in random games. *Math. Oper. Res.* 46 (4), 1552–1572.
- Baldi, P., Rinott, Y., Stein, C., 1989. A normal approximation for the number of local maxima of a random function on a graph. In: *Probability, Statistics, and Mathematics*. Academic Press, Inc., pp. 59–81.
- Candogan, O., Menache, I., Ozdaglar, A., Parrilo, P.A., 2011. Flows and decompositions of games: harmonic and potential games. *Math. Oper. Res.* 36 (3), 474–503.
- Candogan, O., Ozdaglar, A., Parrilo, P.A., 2013. Dynamics in near-potential games. *Games Econ. Behav.* 82, 66–90.
- Coucheny, P., Durand, S., Gaujal, B., Touati, C., 2014. General revision protocols in best response algorithms for potential games. In: *Network Games, Control and Optimization (NetGCoop)*. IEEE Explore, Trento, Italy.
- Durand, S., Garin, F., Gaujal, B., 2019. Distributed best response dynamics with high playing rates in potential games. *Perform. Eval.* 129, 40–59.
- Durand, S., Gaujal, B., 2016. Complexity and optimality of the best response algorithm in random potential games. In: *Algorithmic Game Theory*. In: *Lecture Notes in Comput. Sci.*, vol. 9928. Springer, Berlin, pp. 40–51.
- Fabrikant, A., Jaggard, A.D., Schapira, M., 2013. On the structure of weakly acyclic games. *Theory Comput. Syst.* 53 (1), 107–122.
- Galla, T., Farmer, J.D., 2013. Complex dynamics in learning complicated games. *Proc. Natl. Acad. Sci. USA* 110 (4), 1232–1236.
- Goemans, M., Mirrokni, V., Vetta, A., 2005. Sink equilibria and convergence. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pp. 142–151.
- Harsanyi, J.C., 1973. Oddness of the number of equilibrium points: a new proof. *Int. J. Game Theory* 2, 235–250.
- Heinrich, T., Jang, Y., Mungo, L., Pangallo, M., Scott, A., Tarbush, B., Wiese, S., 2023. Best-response dynamics, playing sequences, and convergence to equilibrium in random games. *Int. J. Game Theory* 52 (3), 703–735.
- Johnston, T., Savery, M., Scott, A., Tarbush, B., 2023. Game connectivity and adaptive dynamics. Technical report arXiv:2309.10609.
- Karlin, A.R., Peres, Y., 2017. *Game Theory, Alive*. American Mathematical Society, Providence, RI.
- Monderer, D., Shapley, L.S., 1996. Potential games. *Games Econ. Behav.* 14 (1), 124–143.
- Nash, J., 1951. Non-cooperative games. *Ann. Math.* (2) 54, 286–295.
- Nash Jr., J.F., 1950. Equilibrium points in n -person games. *Proc. Natl. Acad. Sci. USA* 36, 48–49.
- Ogryczak, W., Ruszczyński, A., 1999. From stochastic dominance to mean-risk models: semideviations as risk measures. *Eur. J. Oper. Res.* 116 (1), 33–50.
- Pangallo, M., Heinrich, T., Farmer, J.D., 2019. Best reply structure and equilibrium convergence in generic games. *Sci. Adv.* 5 (2), eaat1328.
- Pei, T., Takahashi, S., 2019. Rationalizable strategies in random games. *Games Econ. Behav.* 118, 110–125.
- Powers, I.Y., 1990. Limiting distributions of the number of pure strategy Nash equilibria in N -person games. *Int. J. Game Theory* 19 (3), 277–286.
- Rinott, Y., Scarsini, M., 2000. On the number of pure strategy Nash equilibria in random games. *Games Econ. Behav.* 33 (2), 274–293.
- Rosenthal, R.W., 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* 2, 65–67.
- Sanders, J.B.T., Farmer, J.D., Galla, T., 2018. The prevalence of chaotic dynamics in games with many players. *Sci. Rep.* 8 (1), 4902.
- Stanford, W., 1995. A note on the probability of k pure Nash equilibria in matrix games. *Games Econ. Behav.* 9 (2), 238–246.
- Wiese, S.C., Heinrich, T., 2022. The frequency of convergent games under best-response dynamics. *Dyn. Games Appl.* 12 (2), 689–700.
- Wilson, R., 1971. Computing equilibria of N -person games. *SIAM J. Appl. Math.* 21, 80–87.