Response Dynamics, Nash Equilibria and Random Games

Ben Amiet

A thesis submitted for the degree of

Doctor of Philosophy

at Monash University in 2022

School of Mathematics
Monash University
Melbourne, Australia

2022

People tell you that wishful thinking is bad. Do not believe it, this is just one of those generally accepted errors.

George Pólya

Copyright notice

Abstract

Pure Nash equilibria have been used to model a variety of phenomena such as human behaviour in social settings, evolution of a population, how bank customers behave during various economic environments, and network usage. As such, answers to questions regarding their existence and how to locate them are highly sought after. This thesis presents a study of pure Nash equilibria and response dynamics through the use of random games; instead of a deterministic line of inquiry, we aim to answer the question of how likely existence and convergence are.

Chapter 2 addresses games in which a large number of players can pick from one of two actions. Payoffs in this game are independent and identically distributed, but the probability of ties between payoffs is nonzero. We provide results describing a central limit theorem and geometric growth of the number of pure Nash equilibria in the game as the number of players increases. Moreover, we establish a new connection between percolation and game theory to prove a phase transition in the probability of ties regarding where the pure Nash equilibria are located. Finally, we demonstrate that when the probability of ties is small enough, a best-response dynamics reaches a pure Nash equilibrium with high probability.

Chapter 3 explores games in which two players can pick from a large number of actions. We specifically consider games in which payoffs are continuously distributed, so the probability of ties is zero. In this setting, we show that the better-response dynamics is very likely to converge to a pure Nash equilibrium, whereas the best-response dynamics fails to do so, and instead gets trapped in a subset of action profiles.

Chapter 4 extends the aforementioned two-player model by making the probability of ties between payoffs nonzero. We show limiting distribution results for the number of pure Nash equilibria when the payoff distribution is either continuous at the top of its support or the maximal payoff value has positive probability of being attained. When the payoff distribution decays at a speed at least as fast as a geometric distribution, we show that the expected number of pure Nash equilibria is bounded. Moreover, specifically for geometrically-distributed payoffs, we prove the unintuitive result that the expected number of pure Nash equilibria oscillates in the number of actions, with the period increasing geometrically. Finally, we provide simulation results for games in which payoffs follow power law and Poisson distributions, and conjecture on the limiting behaviour in both of these cases.

Declaration

This thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Ben Amiet

14 October 2022

Publications during enrolment

Chapter 2 is based on joint work with Andrea Collevecchio, Marco Scarsini and Ziwen Zhong. The manuscript, titled 'Pure Nash Equilibria and Best-Response Dynamics in Random Games', has been published in Mathematics of Operations Research [3].

Chapter 3 is based on joint work with Andrea Collevecchio and Kais Hamza. The manuscript, titled 'When "Better" is better than "Best"', has been published in Operations Research Letters [2].

Both manuscripts are available on arXiv under the same name.

Acknowledgements

I would not have been able to get to where I am today without the help of my supervisors Andrea Collevecchio and Kais Hamza. Both were instrumental in my upbringing as a mathematician, and without their encouragement and guidance throughout the last seven years, I can confidently say I wouldn't have been able to produce anything near the quality of work that I present here. I learned so much from them, not just in terms of mathematics, but also work ethic and how to be a good educator, and I am deeply grateful that my PhD journey was under their guidance.

I would also like to thank my collaborators Marco Scarsini and Ziwen Zhong. To Marco—thank you for teaching me the joys of international collaboration and introducing me to the world of game theory. To Ziwen—thank you for the endless discussions we had, both when we were working on the same paper, and when we were simply discussing our own roadblocks. Many insights and inspirations were garnered from our discussions, and I just hope I was as helpful to you as you were to me.

Finally, thank you to Mum, Dad and Nicole, for supporting me on this journey, for providing a roof over my head, and for being a soundboard for presentations. And to Kaycee—your unconditional support throughout the last four years made all the difference in the world. Your efforts on your own academic journey provided so much inspiration to me, I'm not sure I would have reached this point without you there beside me.

This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

Contents

1 Introduction

In a given day, we encounter a myriad of situations in which we consciously (or even subconsciously) compete or co-operate with other people, other entities, and our environment to obtain a favourable outcome for ourselves. There are obvious, artificial examples that can easily demonstrate this—deciding on when and how much to bet in a game of poker to maximise our winnings, or what the best choice of shot is in a table tennis rally—but there also exist very natural, everyday occurrences which also match this description. Consider, for example, choosing when to leave for a destination to minimise the time spent stuck in traffic. In each of these examples, we need to consider the desires of the other actors in the system; we cannot simply behave in our best interests and assume everything will go according to plan. If we are holding four-of-a-kind, we will almost certainly win the hand, but making too large a bet will likely dissuade other players from placing bets on the same hand, thus reducing our overall winnings. Similarly, we know that other drivers will be using the roads during peak hour, so we can try to avoid these times to minimise our own time on the road.

This analysis of other players' behaviours in the context of maximising some heuristic

(P1, P2)	Co-operation	Silence
Co-operation	(5, 5)	(1, 7)
Silence	(7, 1)	(3, 3)

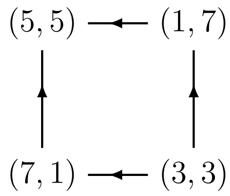

Figure 1.1: A tabular representation of the prisoners' dilemma (left) and its corresponding graphical representation (right).

function for ourselves is the cornerstone of game theory. More formally, each of the above situations is an example of a *game*, comprised of a set of *players* acting to try and maximise their individual *payoff*, a function of the actions each player chooses to make at any given turn in the game. We will entirely focus on single-turn games; the canonical example of such a game is the prisoners' dilemma. In this game, two prisoners are isolated in two interrogation rooms, and are told the following:

(i) if they co-operate with the investigation and the other prisoner does not, they will receive a sentence of 1 year, and the other prisoner will receive a sentence of 7 years;

(ii) if they both co-operate (and necessarily give conflicting statements), they will both receive a sentence of 5 years;

(iii) if neither prisoner co-operates, they will both receive a sentence of 3 years.

This is an example of a *normal-form representation* of a game: a representation of all possible action profiles and their respective payoffs for each player. From the normal-form representation, we can also generate a graphical representation of the game. Given an action profile, if a player can change their action and improve their payoff, then we direct an edge from that profile to the profile corresponding to their changed action. See Fig. 1.1 for a tabular representation of the prisoners' dilemma and its corresponding graphical representation. Importantly, it is always possible to recover the graphical representation from the normal-form representation, but the converse is not true. In this thesis, we will concern ourselves with games consisting of random payoffs. Chapter 2 will focus on games with a large number of players that can choose from one of two actions, while Chapters 3 and 4 will feature games with a large number of actions from which our two players can choose.

Consider the set of actions in which both prisoners choose to co-operate with the investigation. If either of them changes their action, their sentence will increase from

5 to 7 years, so neither prisoner is incentivised to do so. This is an example of a *pure Nash equilibrium* (PNE), an action profile such that no player can change their action to improve their payoff. Notice that in the graphical representation in Fig. 1.1, this is depicted by an action profile with no outgoing edges.

PNEs offer a means to analyse the effects of strategic interaction between players in a game. One method may involve comparing the action profile which attains the best payoff for each player to the PNEs in a game. The ratio of some meaningful function (e.g., sum, minimum) of these two payoffs for each of these profiles is known as the *price of anarchy*. In a sense, the price of anarchy offers a metric describing the impact of players if they opt to act selfishly in a game instead of working together to reach a common goal. In the prisoner's dilemma scenario described above, it is clear that the prisoners should both refuse to co-operate with the investigation, as this results in the minimal total sentence (6 years). As the PNE has a total sentence of 10 years, the price of anarchy in this case (if we choose our function to be the sum of payoffs) is 1.66. Notice that if the sentences for both prisoners co-operating with the investigators is increased, this will increase the price of anarchy in our game.

Social settings such as the prisoners' dilemma are only one example of a use case for PNEs. Similar to predicting actions in people, a PNE can be used to model the evolution of a population [30]. In this setting, the payoffs instead represent the probability of reproductive success. If the population reaches a PNE then the population is "evolutionarily stable"; that is, if mutations were to arise, the population would still tend to revert back to the PNE. Moreover, multiple PNEs in a system could suggest that a population evolves towards different evolutionarily stable states depending on the initial composition of the population.

In the context of network design, the payoff function can instead be thought of as the cost of any connections to which a given user has access [4]. Each user acts selfishly to achieve the best connection possible in order to get the most out of the network; if the network is at an NE, no users will want to deviate, and the network can be considered stable. Furthermore, we can compare the cost of the PNE to the cost of an optimal state which is instead centrally enforced, and obtain a measure of the cost of stability.

Because of the significance of PNEs both as one of the cornerstone objects in game theory and due to their interpretations in other areas of study, it is important to be able to locate them efficiently. To this end, we make use of iterative procedures that converge to PNEs, notably *better-response dynamics* (bRD) and *best-response dynamics* (BRD). From a given action profile, bRD allows a single player at random to change their action, and if the resultant action profile yields a higher payoff for that player, bRD moves to that action profile. The behaviour of BRD is almost identical, save for the fact that the randomly chosen player must chose the action which yields the highest payoff for them, given all other players do not deviate. Once these processes reach an action profile from which no player is willing to deviate, they have converged to a PNE.

This raises an important question: when do these processes to reach a PNE? It is completely possible that a game does not even admit PNEs in the first place, so of course in this setting both processes fail to converge. However, even in the presence of PNEs, we could see these processes encounter a set of action profiles upon which they cycle indefinitely, which we refer to as a *trap*. In this thesis, we will explore this question through the use of random games: if we pick a game at random that satisfies our assumptions, how likely are PNEs to exist? And if they do, with what probability do the response dynamics locate them?

1.1 Related Work

The approach we take in this thesis, in which we analyse the likelihood with which PNEs exist and convergence of iterative procedures on a randomly drawn game, has a storied past. Goldman [25] showed that, in the context of zero-sum games, the probability with which a PNE exists tends to 0 as the number of strategies grows. Goldberg, Goldman, and Newman [24] generalised this model to general two-player games with continuous payoffs, and showed that the probability of having at least one PNE converges to $1-e^{-1}$. Dresher [15] then generalised this convergence result further to any number of finite players.

Since then, random games have been used to address a plethora of questions on the existence of PNEs. Powers [38] showed that, as the number of actions in a game with at least two players tends to infinity, the distribution of the number of PNEs converges to a Poisson random variable with mean 1. The assumption of i.i.d. payoffs was relaxed by Rinott and Scarsini [40], in which payoff independence was maintained between different action profiles, but payoffs corresponding to the same action profile could be correlated. They recovered the result from Powers [38] in the case of independence, but also showed negative correlation forced the number of PNEs to zero, while positive correlation resulted in divergence. Stanford [42] instead derived the distribution that describes the number of PNEs in random games, and consequently obtained the Powers result. Stanford also proved results regarding the asymptotic distribution of the number of PNEs in symmetric games [45], games with vector payoffs [43], and games of common interest [44]. Kultti, Salonen, and Vartiainen [29] take a slightly different approach and instead assign a distribution to the number of best responses to each action profile, and attain asymptotic results on this model. Takahashi [46] conditioned on the game having nondecreasing best-response functions, and found that this greatly increased the number of PNEs compared to the unconditional expected number of PNEs. Further analysis on properties of PNEs when they exist such as Pareto efficiency (i.e., the sum of payoffs is maximised) was investigated by Cohen [10]. Point rationalisable strategies—a concept introduced by Bernheim [6]—in two-person games were studied by Pei and Takahashi [37], and were shown to be asymptotically uncorrelated with PNEs.

As established, questions about the existence of PNEs are not the only point of interest in games; convergence of iterative procedures like BRD is also a very active area of research, especially considering how computationally complex the act of locating a PNE is, as shown by Daskalakis, Goldberg, and Papadimitriou [12]. Several authors have explored sufficient conditions on payoffs to guarantee that BRD converges to a PNE: Monderer and Shapley [34] studied the potential game setting, while Friedman and Mezzetti [22] and Takahashi and Yamamori [47] investigated quasi-acyclic games. Potential games—games in which the incentive to change action can be described using a global function—have received special attention in the realm of BRD convergence; see Coucheney, Durand, Gaujal, and Touati [11], Durand and Gaujal [17] and Durand, Garin, and Gaujal [16]. See also Blume [7] and Young [48] for further studies on using BRD to locate PNEs.

Outside of these contexts however, convergence of BRD is not guaranteed, as iterative procedures can end up cycling indefinitely on a subset of action profiles. These subsets were introduced by Goemans, Mirrokni, and Vetta [23] as *sink equilibria* (in this thesis we refer to them as traps) and were further studied by Mirrokni and Skopalik [33]. In generic random games, Pangallo, Heinrich, and Farmer [35] demonstrated nonconvergence of various learning procedures through simulation. Following this, Pangallo et al. [36] explored how playing sequences affect the convergence of BRD. Recently, Mimun, Quattropani, and Scarsini [32] studied the impact of correlated payoffs within action profiles in a two-player game, and showed that BRD converges to a PNE.

Convergence of bRD is not as well-studied as that of BRD. Reny [39] introduces the concept of *better-reply secure* games and proves that compact, convex strategy space games with quasiconcave payoffs in each player's strategies admit PNEs. In this work, Reny also made significant progress in developing requirements needed for PNEs to exist in the presence of discontinuous payoffs. Kukushkin [28] was able to relax the assumption of better-reply security and show that, in the presence of discontinuities in the payoff distribution, bRD converges to a PNE in potential games, games with strategic complements, and aggregative games; Dindoš and Mezzetti [14] also investigated convergence of bRD in aggregative games. On the theme of studying games with specific properties, convergence in acyclic games was studied by Fabrikant, Jaggard, and Schapira [21].

1.2 Background

Here we will introduce several useful concepts that are used throughout this thesis.

1.2.1 Percolation

Diffusion processes are a widely studied phenomena that we can employ to understand how a fluid randomly flows through a medium. Conversely, a *percolation graph*, introduced by Broadbent and Hammersley [9] to study fluid flow through coal, approaches the problem from the opposite direction: instead of ascribing the randomness to the fluid, we instead consider the medium itself as being the source of randomness in the model. This simple change of perspective gave rise to the entire field of percolation theory, a source of deep mathematical results and significant physical implications alike.

To answer the question of how a fluid flows through a lump of coal, we can consider the coal as a finite subgraph of $\mathbb{Z}^3$. Independently for each edge inside this subgraph, with some parameter $p \in [0, 1]$, we leave the edge open with probability p, and close it with probability $1 - p$. Physically, an open edge is a tunnel through which the fluid can propagate, whereas a closed tunnel blocks the flow of fluid. Initially, this model was introduced in an effort to model the use of coal in gas masks [26]; in this context, we would be interested in whether or not a path of open edges exists from one "side" of the subgraph to the other. If such a path does exist, the coal admits gas through the mask, rendering it useless.

The model in which edges in a subgraph are randomly closed or left open is specifically known as *bond percolation*. An alternative model, *site percolation*, in which vertices are randomly designated as open or closed, is also widely studied, but will not be used in this thesis. Furthermore, many generalisations of both of these models exist, in which we can do away with the assumption of independence, or perhaps the edges/vertices are closed with probabilities that vary over the graph.

In this thesis, we will entirely concern ourselves with the specific case of percolation on the N-dimensional hypercube. Formally, this is the graph with vertex set $\{0,1\}^N$ and edge set defined such that an edge between two vertices exists if and only if the two vertices differ in exactly one coordinate. In this context, we will make use of classical results regarding monotonicity in the percolation parameter of the size of the largest connected component by Bollobás [8] and the coincidence of the set of isolated vertices and the set of inaccessible vertices in a graph obtained via percolation by Erdős and Spencer [20]. These combined with more modern results in the field describing the number of inaccessible vertices in a percolation graph as well as the number of inaccessible vertices in a ball of fixed radius by McDiarmid, Scott, and Withers [31] will help us reach our desired results in Chapter 2.

1.2.2 Borel-Cantelli Lemmas

The first and second Borel-Cantelli lemmas both deal with the limit supremum of a sequence of events, which is explicitly stated as

$$\limsup_{n\in\mathbb{N}} E_n = \bigcap_{n\in\mathbb{N}} \bigcup_{k\geq n} E_n.$$

An intuitive interpretation of this is no matter how many event outcomes you have already observed, there is always an event in the future with a positive outcome. An alternative way of describing the limit supremum is that the events E_n happen infinitely often. Note that the converse, that is, the events E_n occur finitely often, is a stronger statement than the complementary events ${E_n}^c$ occurring infinitely often.

Lemma 1.1 (First Borel-Cantelli Lemma). *Let $(E_n)_{n\in\mathbb{N}}$ be a sequence of events. If the probability of these events is summable, i.e.,*

$$\sum_{n=1}^{\infty} \mathbb{P}(E_n) < \infty,$$

then

$$\mathbb{P}\left(\limsup_{n\to\infty} E_n\right) = 0.$$

In words, if the sequence of probabilities of events is summable, then the events occur infinitely often with probability zero. Hence, at some random point in the sequence, the events will never occur again.

Lemma 1.2 (Second Borel-Cantelli Lemma). *Let* $(E_n)_{n\in\mathbb{N}}$ *be a sequence of independent events. If the probability of these events is not summable, i.e.,*

$$\sum_{n=1}^{\infty} \mathbb{P}\left(E_n\right) = \infty,$$

then

$$\mathbb{P}\left(\limsup_{n\to\infty} E_n\right) = 1.$$

Note that this lemma requires independence, but still completes a nice picture. For a sequence of independent events, we can explicitly determine whether or not they happen finitely often or infinitely often. This is not true in the case that they are dependent events. For example, consider a uniform random variable U over the interval $[0, 1]$, and the sequence of events $(E_n)_{n\in\mathbb{N}}$ defined such that E_n is true if $U < 1/n$. Once U is realised, there exists some n such that $1/n < U$ (unless $U = 0$, which occurs with probability 0), so the events cease to occur from this point onward in the sequence.

1.2.3 Berry-Esseen Theorem

The Berry-Esseen theorem tells us how quickly a sum of random variables converges to a normal distribution, under certain conditions. This theorem has been refined many times since its original statement, but we cite a version from Durrett [18] here. Let Φ denote the distribution function for a standard normal random variable.

Lemma 1.3 (Berry-Esseen theorem). *Let* $X_1, X_2, \ldots$ *be a sequence of i.i.d. zero-mean random variables with* $\mathbb{E}\left[X_i^2\right] = \sigma^2$ *and* $\mathbb{E}\left[|X_i|^3\right] = \rho < \infty$ *for all* $i \in \mathbb{N}$. *Denote by* $F_n(x)$ *the distribution function of* $\sum_{i=1}^{n} X_i/(\sigma\sqrt{n})$. *Then there exists a constant* C *such that*

$$|F_n(x) - \Phi(x)| \leq \frac{C\rho}{\sigma^3\sqrt{n}}.$$

The quest for finding a tight upper bound on the constant in the Berry-Esseen theorem is storied. A thorough narration of the history of this constant can be found in Korolev and Shevtsova [27]. For now, the best estimate is $C < 0.4748$, owing to Shevtsova [41].

1.3 Notation

Throughout this thesis, we will concern ourselves with games of the form

$$\Gamma = ([N], (S_i)_{i \in [N]}, (g_i)_{i \in [N]})$$

where $[N] := \{1, ..., N\}$ is the set of players and S_i is the set of actions from which player i can choose. Defining $\mathsf{S} := \times_{i \in [N]} S_i$, we can let $g_i : \mathsf{S} \to \mathbb{R}$ be the payoff function of player i. For convenience, we will write $\boldsymbol{s}_{-i}$ to be the action profile of all players except player i. With a slight abuse of notation, we will also write $(x, \boldsymbol{s}_{-i})$ to denote the action profile agreeing with $\boldsymbol{s}$ for all players except i, and player i adopts x as their action. If two action profiles $\boldsymbol{s}$ and $\boldsymbol{t}$ are such that $\boldsymbol{s} \neq \boldsymbol{t}$ and $\boldsymbol{t} = (x, \boldsymbol{s}_{-i})$ for some $x \in S_i$, then we write $\boldsymbol{s} \sim_i \boldsymbol{t}$.

Definition 1.4. An action profile $\boldsymbol{s}^*$ is a *pure Nash equilibrium* (PNE) of the game Γ if for all $i \in [N]$ and for all $s_i \in S_i$ we have

$$g_i(\boldsymbol{s}^*) \geq g_i(s_i, \boldsymbol{s}^*_{-i}). \tag{1.1}$$

The set of PNEs is denoted by E.

An action profile $\boldsymbol{s}^*$ is a *strict pure Nash equilibrium* (SPNE) if for all $i \in [N]$ and for all $s_i \in S_i$ we have

$$g_i(\boldsymbol{s}^*) > g_i(s_i, \boldsymbol{s}^*_{-i}). \tag{1.2}$$

The set of SPNEs is denoted by E.

If we are dealing with a sequence of games $(\Gamma_n)_{n \in \mathbb{N}}$, we will often refer to the associated sequence of sets of PNEs and SPNEs as $(\mathsf{E}_n)_{n \in \mathbb{N}}$ and $(\mathsf{E}^*_n)_{n \in \mathbb{N}}$, respectively.

Definition 1.5. For any $i \in [N]$ and $\boldsymbol{s} \in \mathsf{S}$, define the sets

$$\mathsf{P}^{(i)}_{\boldsymbol{s}} := \{\boldsymbol{t} \in \mathsf{S} : \boldsymbol{t} \sim_i \boldsymbol{s}, g_i(\boldsymbol{t}) > g_i(\boldsymbol{s})\},$$
$$\mathsf{M}^{(i)}_{\boldsymbol{s}} := \{\boldsymbol{t} \in \mathsf{P}^{(i)}_{\boldsymbol{s}} : g_i(\boldsymbol{t}) \geq g_i(\boldsymbol{u}) \text{ for all } \boldsymbol{u} \sim_i \boldsymbol{s}\}.$$

In words, $\mathsf{P}^{(i)}_{\boldsymbol{s}}$ is the set of all *profitable deviations* available to player i from $\boldsymbol{s}$, while $\mathsf{M}^{(i)}_{\boldsymbol{s}}$ is the set of all *maximally profitable deviations* available to player i from $\boldsymbol{s}$.

We now define the two processes on which this thesis will focus.

Definition 1.6. The *better-response dynamics* (bRD), denoted $\mathsf{bRD} = \{\mathsf{bRD}(n)\}_{n=0}^{\infty}$, is a discrete-time process on S that evolves as follows. At time $n+1$, pick a player at random, independently of the current value of the process and its past. Call this random variable $I \in [N]$, and choose $\mathsf{bRD}(n+1)$ uniformly at random from the set $\mathsf{P}^{(I)}_{\mathsf{bRD}(n)}$. If the latter set is empty, repeat the procedure with another player selected uniformly at random. If $\mathsf{P}^{(i)}_{\mathsf{bRD}(n)} = \varnothing$ for all $i \in [N]$ then set $\mathsf{bRD}(n+1) = \mathsf{bRD}(n)$; at this point, bRD has reached a PNE, and we say that the process has converged.

The *best-response dynamics* (BRD), denoted $\mathsf{BRD} = \{\mathsf{BRD}(n)\}_{n=0}^{\infty}$, is defined similarly, except it chooses its action profile at time $n+1$ from the set $\mathsf{M}^{(I)}_{\mathsf{BRD}(n)}$. Its convergence criterion follows mutatis mutandis. Importantly, for any $\boldsymbol{s} \in \mathsf{S}$ and $i \in [K]$, we have that $\mathsf{M}^{(i)}_{\boldsymbol{s}} \subseteq \mathsf{P}^{(i)}_{\boldsymbol{s}}$. Hence, BRD cannot consecutively jump between best responses for $\boldsymbol{s}_{-i}$.

These processes can behave in vastly different ways depending on the game we let them explore. In Chapter 2, we consider games with only two actions. On these games, bRD and BRD both behave identically, as there is at most one action that each player can change. Conversely, in Chapter 3 and 4, we consider games with only two players. Under the assumption that payoffs are drawn from a distribution without atoms, BRD is actually a deterministic process once the starting action profile is fixed and the payoffs are realised. Discontinuities, as we will see in Chapter 4, result in unintuitive behaviour.

Both BRD and bRD fail to converge in any games that do not contain PNE, but they could also infinitely cycle on a subset of action profiles. We define such subsets of vertices now. Denote by $|A|$ the cardinality of set A.

Definition 1.7. Let X denote either of the two processes bRD or BRD. A nonempty subset of action profiles $\tau \subset \mathsf{S}$ such that $|\tau| \geq 2$ is called an X-*trap* if $\{\mathsf{X}(n) \in \tau\} \subset \{\mathsf{X}(n+1) \in \tau\}$ and, for all $\boldsymbol{s} \in \tau$,

$$\{\mathsf{X}(n) \in \tau\} \subset \{\inf\{k \in (n, \infty) \cap \mathbb{N} \colon \mathsf{X}(k) = \boldsymbol{s}\} < \infty \text{ a.s.}\}.$$

We exclude traps of size 1 so that PNE are excluded from this definition. Moreover, it is easy to see that traps must be of at least size 4. The only way for either of our processes to not converge is to enter a trap.

We will also make use of the following notation to describe limiting behaviour of sequences. Given two positive-valued sequences $(a_n)_n$ and $(b_n)_n$, we say that:

- $a_n = O(b_n)$ if $\limsup_n a_n/b_n < \infty$;
- $a_n = o(b_n)$ if $\lim_{n\to\infty} a_n/b_n = 0$.

2

Multi-player Games with Two Actions

This chapter is based on joint work with Andrea Collevecchio, Marco Scarsini and Ziwen Zhong titled ‘Pure Nash Equilibria and Best-Response Dynamics in Random Games’, and has been published in Mathematics of Operations Research [3].

In this chapter, we will only consider random games consisting of N players, each of whom can pick from two actions. To facilitate this, consider a probability space $(\Omega, \mathscr{F}, \mathbb{P})$ upon which the following sequence of random games is defined. Let Γ_N be a game with N players, $S_i = \{0, 1\}$ for each $i \in [N]$, and i.i.d. payoffs. In particular, for each $\boldsymbol{s} \in \mathsf{S}$, the payoff $g_i(\boldsymbol{s})$ is a realisation of a random variable $Z_i^{\boldsymbol{s}}$, and the random variables $(Z_i^{\boldsymbol{s}})_{i\in[N],\boldsymbol{s}\in\mathsf{S}}$ are i.i.d. The symbol Z denotes a generic independent copy of $Z_i^{\boldsymbol{s}}$. We will also denote by E_N and E_N^* the set of PNEs and SPNEs of a game with N players.

Importantly, we need to define

$$\alpha := \mathbb{P}\left(Z_1 = Z_2\right), \qquad \beta := \mathbb{P}\left(Z_1 < Z_2\right) = \frac{1-\alpha}{2},$$

where Z_1 and Z_2 are i.i.d. copies of Z. All results in this chapter will depend on α, whereas most of the existing literature only deals with the case $\alpha = 0$. While it is true that we could have defined only one parameter, in the larger context of the chapter, α

and β tend to fulfill different roles. The probability of a tie is denoted as α, and is a more natural parameter when considering the game itself, whereas β is better suited to describing percolation results. We switch between these to hopefully achieve better clarity in our statements of results throughout the chapter.

2.1 Number of Nash Equilibria in Random Games

Observe that, by virtue of payoffs being i.i.d., the probability with which an action profile $\boldsymbol{s}$ is a PNE is $(1-\beta)^N$, and the probability that it is an SPNE is β^N. As a consequence, for any $N \geq 1$,

$$\mathbb{E}\left[|\mathsf{E}_N|\right] = 2^N(1-\beta)^N \qquad \text{and} \qquad \mathbb{E}\left[|\mathsf{E}_N^*|\right] = 2^N\beta^N.$$

This implies that the expected number of PNEs is always one when $\alpha = 0$ and diverges when $\alpha > 0$. Moreover, when $\alpha = 0$, the numbers of PNEs and SPNEs are almost surely equal, as any two payoffs are almost surely different. The following theorems elaborate on this theme, providing a description for the asymptotic behaviour of $|\mathsf{E}_N|$ and $|\mathsf{E}_N^*|$.

Theorem 2.1 (Behaviour of SPNE). *Consider a sequence of random games* $(\Gamma_N)_{N\in\mathbb{N}}$.

(a) If $\alpha = 0$ *then, for all* $k \in \mathbb{N} \cup \{0\}$,

$$\lim_{N\to\infty} \mathbb{P}\left(|\mathsf{E}_N^*| = k\right) = \frac{e^{-1}}{k!}.$$

(b) If $\alpha > 0$ *then*

$$\mathbb{P}\left(\lim_{N\to\infty} |\mathsf{E}_N^*| = 0\right) = 1.$$

Fig. 2.1 shows demonstrates not only the results described in Theorem 2.1 (b), but also the geometric upper bound in α that is used in the proof of this statement.

We will need some more machinery to prove the following two theorems, so we will return to these later in the chapter. Recall that Φ denotes the cumulative distribution of a standard normal random variable.

Theorem 2.2 (CLT for PNE). *Assume that* $\alpha > 0$. *Then there exists a constant* $K_\alpha > 0$, *dependent only on* α, *such that*

$$\sup_x \left| \mathbb{P}\left(\frac{|\mathsf{E}_N| - (1+\alpha)^N}{(1+\alpha)^{N/2}} \leq x\right) - \Phi(x) \right| \leq K_\alpha N \max\left(\frac{1+\alpha}{2}, \frac{1}{(1+\alpha)^{N/2}}\right)^N.$$

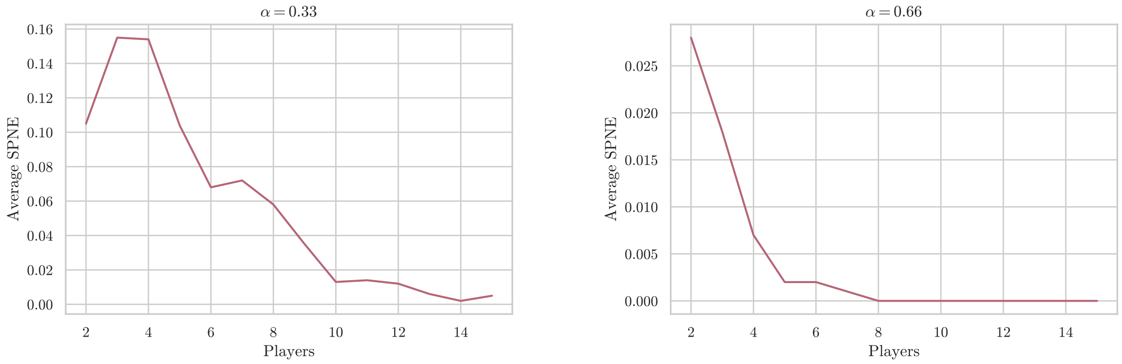

Figure 2.1: Average SPNE vs. number of players over 1000 simulations per player.

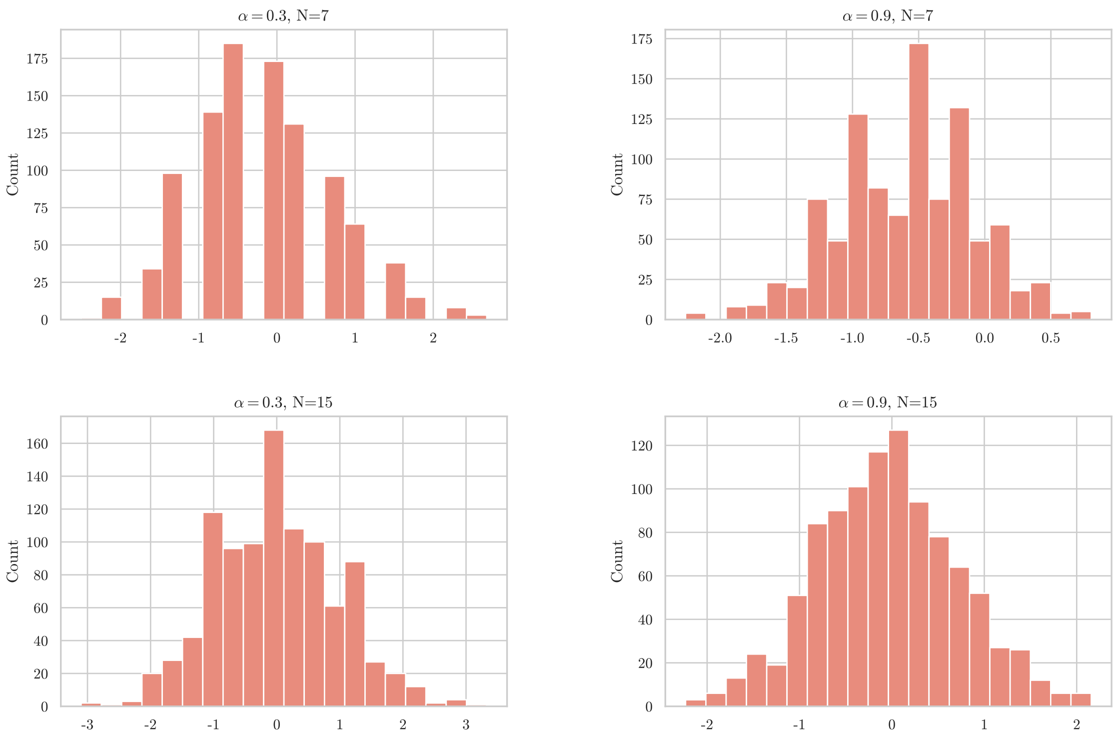

Figure 2.2: Scaled number of PNE in 1000 simulations.

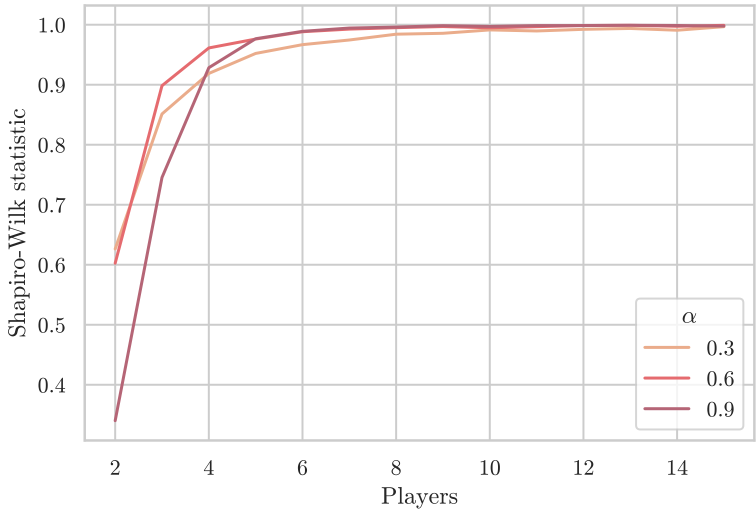

Figure 2.3: Shapiro-Wilk statistic for varying α. 1000 simulations per combination of N and α.

See Fig. 2.2 for a demonstration of the normality of the distribution of the number of PNE. In a similar vein, Fig. 2.3 displays both convergence to normality as N increases and speed of convergence to normality as α increases, both of which are predicted by Theorem 2.2.

A direct consequence of the above theorem is that the number of PNEs grows exponentially in N whenever $\alpha > 0$. The following is a more precise statement.

Theorem 2.3. *If* $\alpha > 0$ *then*

$$\mathbb{P}\left(\lim_{N\to\infty} \frac{|\mathsf{E}_N|}{(1+\alpha)^N} = 1\right) = 1.$$

2.2 Best-response Dynamics

We will denote by BRD_N the BRD on a game with N players. For simplicity, and without loss of generality, we will always set $\mathsf{BRD}_N(0) = \mathbf{0}$, where $\mathbf{0}$ is the N-dimensional zero vector.

It will be helpful in this chapter to consider graphical representations of our games. To this end, we associate with our game Γ_N the following graph. To each $\boldsymbol{s} \in \mathsf{S}$, assign a vertex $\boldsymbol{v} \in \mathcal{V}$. Two vertices $\boldsymbol{u}, \boldsymbol{v}$ are connected by an edge in $\mathcal{E}_N$ if and only if $u_i \neq v_i$ for exactly one $i \in [N]$ and $u_j = v_j$ for all $j \neq i$. We then say $\boldsymbol{u}$ and $\boldsymbol{v}$ are *neighbours*, denoting this relationship by $\boldsymbol{u} \sim_i \boldsymbol{v}$. In general, if there exists some i such that $\boldsymbol{u} \sim_i \boldsymbol{v}$, we write $\boldsymbol{u} \sim \boldsymbol{v}$. Notice that this exactly coincides with our use of $\sim$ in the game theoretic sense.

We will then associate to $\mathcal{H}_N$ a random partially oriented graph with parameter β $\overrightarrow{\mathcal{H}}_N^{\beta} = \left(\mathcal{V}_N, \overrightarrow{\mathcal{E}}_N^{\beta}\right)$ as follows. For convenience, given a vertex $\boldsymbol{u} \in \mathcal{V}_N$ and its corresponding action profile $\boldsymbol{s} \in \mathsf{S}$, we will write $Z_i^{\boldsymbol{u}} = Z_i^{\boldsymbol{s}}$. For two vertices $\boldsymbol{u} \sim_i \boldsymbol{v}$, the oriented edge from $\boldsymbol{u}$ to $\boldsymbol{v}$, denoted $\overrightarrow{[\boldsymbol{u}, \boldsymbol{v}]}$, is in $\overrightarrow{\mathcal{E}}_N^{\beta}$ if and only if $Z_i^{\boldsymbol{u}} > Z_i^{\boldsymbol{v}}$. An unoriented edge between the two vertices, denoted $[\boldsymbol{u}, \boldsymbol{v}]$, exists between them if and only if $Z_i^{\boldsymbol{u}} = Z_i^{\boldsymbol{v}}$. In the context $\alpha = 0$ (or $\beta = 1/2$), every edge in the graph is oriented almost surely, as ties occur with probability zero. Hence, $\overrightarrow{\mathcal{H}}_N^{1/2}$ is a *proper* orientation of $\mathcal{H}_N$ where each edge is independently oriented in one direction or the other with probability $1/2$.

With an abuse of notation, we will use E_N to refer to the set of vertices in $\mathcal{V}_N$ corresponding to action profiles that are PNEs in Γ_N. Note that graphically, these are vertices adjacent to only incoming or undirected edges.

Definition 2.4. We say that $\boldsymbol{v}$ is *directly accessible* from $\boldsymbol{u}$ if $\overrightarrow{[\boldsymbol{u}, \boldsymbol{v}]} \in \overrightarrow{\mathcal{E}}_N^{\beta}$. We say that $\boldsymbol{v}$ is *accessible* from $\boldsymbol{u}$ if there exists a finite sequence $\boldsymbol{u}_0, ..., \boldsymbol{u}_k$ such that $\boldsymbol{u} = \boldsymbol{u}_{\mathbf{0}}$, $\boldsymbol{u}_{\boldsymbol{k}} = \boldsymbol{v}$ and, for all $i \in \{0, ..., k-1\}$, we have $\overrightarrow{[\boldsymbol{u}_i, \boldsymbol{u}_{i+1}]} \in \overrightarrow{\mathcal{E}}_N^{\beta}$.

The notion of accessibility has a natural interpretation in terms of the BRD, insofar that a vertex $\boldsymbol{u}$ is accessible from $\mathbf{0}$ if and only if there is a positive probability that BRD reaches $\boldsymbol{u}$. Also of note is the fact that any vertex in a trap must be accessible from every other vertex in the same trap.

On $\overrightarrow{\mathcal{H}}_N^{\beta}$, define the partition

$$\mathcal{V}_N = \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \,\dot{\cup}\, \overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$$

where $\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$ is the set of all vertices accessible from $\mathbf{0}$ and $\overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$ is the set of all inaccessible vertices from $\mathbf{0}$. The following theorem demonstrates a phase transition in accessibility of vertices, dependent only on α.

Theorem 2.5. *Let $(\Gamma_N)_{N\in\mathbb{N}}$ be a sequence of random games, with probability of tie $\alpha = 1 - 2\beta$.*

(a) If $0 \leq \alpha < 1/2$ then

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\left|\mathsf{E}_N \cap \overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}\right| > 0\right) < \infty.$$

(b) If $\alpha = 1/2$ then

$$\lim_{N\to\infty} \mathbb{P}\left(\left|\mathsf{E}_N \cap \overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}\right| > 0\right) \geq 1 - e^{-1}.$$

(c) *If $\alpha > 1/2$ then, for any $K > 0$,*

$$\lim_{N\to\infty} \mathbb{P}\left(\left|\mathsf{E}_N \cap \overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}\right| > K\right) = 1.$$

This theorem can be interpreted as follows:

(a) if $0 \leq \alpha < 1/2$, we can use the first Borel-Cantelli lemma to conclude that there exists a finite random N^* such that, for all $N \geq N^*$, each of the PNE in Γ_N is potentially reachable by BRD_N;

(b) if $\alpha = 1/2$, PNE exist that are unreachable by BRD_N with positive probability;

(c) if $\alpha > 1/2$, the number of PNE that are not reachable by BRD_N grows to infinity with probability approaching one.

See Fig. 2.4 for graphical representations of these behaviours.

To complete this narrative, we give a result on the probability of convergence of BRD_N. We use $\lfloor x \rfloor$ to denote the floor of x, i.e., the largest integer smaller than or equal to x.

Theorem 2.6 (Convergence of BRD)**.** *If $0 \leq \alpha < 2^{3/4} - 1$ then*

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathsf{BRD}_N \textit{ does not converge to a PNE}\right) < \infty.$$

Using the first Borel-Cantelli lemma, we can conclude that there exists a finite random N^* such that, for all $N \geq N^*$, the process BRD_N converges to a PNE. Equivalently, as entering a trap is the only way for BRD_N to fail to converge, we can reinterpret this result as BRD_N only enters a trap finitely often. To prove this theorem, we will instead show the stronger result that traps exist only finitely often. Fig. 2.5 shows the number of vertices in traps for varying values of α and N.

2.3 Bond Percolation

Most of the proofs of the theorems stated earlier in this chapter rely on a coupling between between random games and a bond percolation model. This coupling will be such that the set of strategies that are accessible by BRD_N coincide with the connected component containing $\mathbf{0}$ on the hypercube. Once established, we will use results on the geometry of the cluster, obtained by McDiarmid, Scott, and Withers [31], to infer limit theorems for the number of PNEs and the BRD.

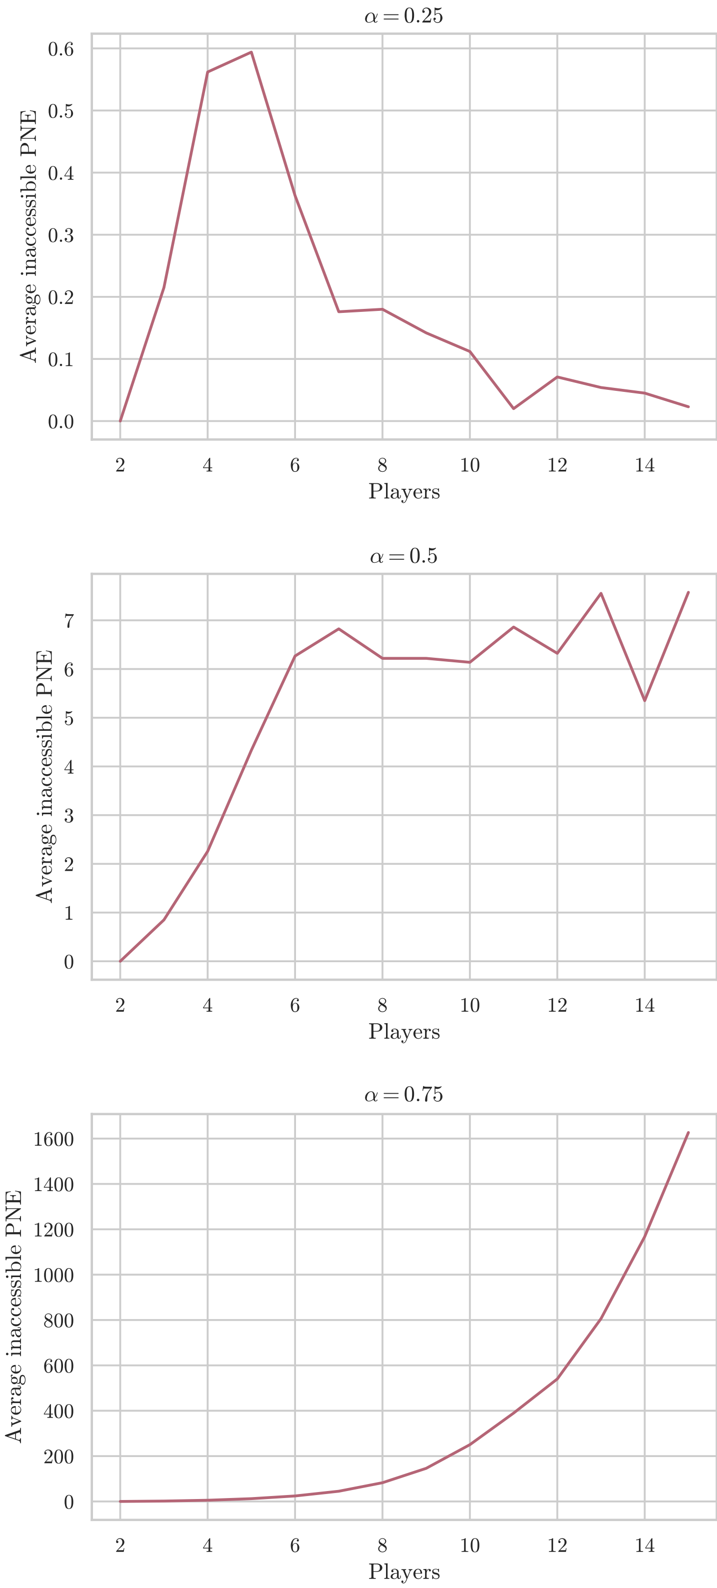

Figure 2.4: Number of inaccessible PNE in 1000 simulations.

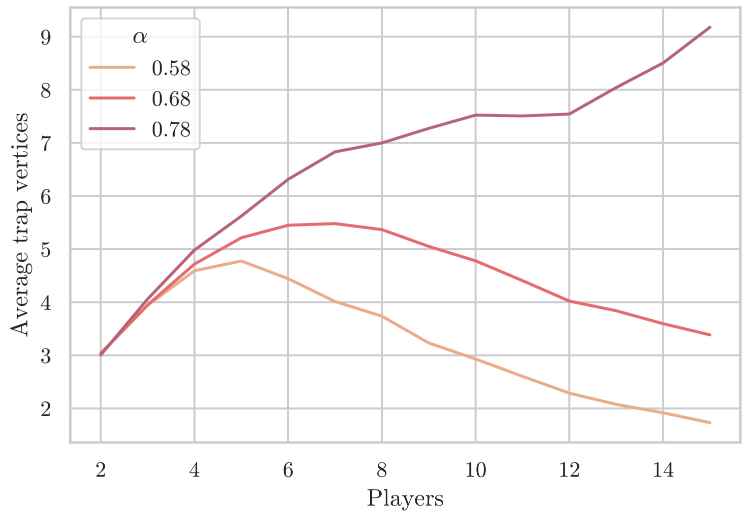

Figure 2.5: Number of trap vertices in 1000 simulations.

Independent bond percolation on $\mathcal{H}_N$ is defined as follows. For each edge in $\mathcal{E}_N$, flip an independent coin with probability β of showing heads. If the result is heads, then declare the edge to be *open*; otherwise, the edge is *closed*. The random subgraph obtained by deleting the closed edges is our independent bond percolation. It includes all vertices in $\mathcal{V}_N$, but could be disconnected. Whenever we mention a connected component of the percolation that has a certain property, we implicitly refer to the largest connected component that satisfies that property. Moreover, we use the term *giant component* to refer to the connected component with the largest number of vertices.

The next result relates $\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$ to a percolation on the hypercube.

Proposition 2.7. *For any $\beta \in [0, 1/2]$, there exists a percolation $\mathcal{B}_N^\beta$ such that the vertex set of its connected component containing $\mathbf{0}$, denoted by $\mathcal{L}_N^{\beta,\mathbf{0}}$, coincides with $\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$.*

Proof. First, define the event

$$\boldsymbol{u} \to \boldsymbol{v} := \left\{ \overrightarrow{[\boldsymbol{u}, \boldsymbol{v}]} \in \overrightarrow{\mathcal{E}}_N^\beta \right\}.$$

Because each player only has two strategies, we have that $\boldsymbol{r} \to \boldsymbol{t}$ is independent of $\boldsymbol{u} \to \boldsymbol{v}$ for every $\boldsymbol{u} \to \boldsymbol{v} \notin \{\boldsymbol{r} \to \boldsymbol{t}, \boldsymbol{t} \to \boldsymbol{r}\}$. For any subset $\mathcal{U} \in \mathcal{V}_N$, define

$$\overrightarrow{\partial \mathcal{U}} := \{\boldsymbol{v} \in \mathcal{V}_N \setminus \mathcal{U} : \exists \boldsymbol{u} \in \mathcal{U} \text{ such that } \boldsymbol{u} \to \boldsymbol{v}\},$$

$$\partial \mathcal{U} := \{\boldsymbol{v} \in \mathcal{V}_N \setminus \mathcal{U} : \exists \boldsymbol{u} \in \mathcal{U} \text{ such that } \boldsymbol{u} \sim \boldsymbol{v}\}.$$

Next, we construct an exploration process on $\overrightarrow{\mathcal{H}}_N$ that will act as one of the two pillars to help establish a coupling between our game and percolation. To this end, we define the process $\mathcal{Q} = (\mathcal{Q}_n)_{n\in\mathbb{N}}$ with $\mathcal{Q}_1 = \{\mathbf{0}\}$ by the recursive relation

$$\mathcal{Q}_{n+1} = \mathcal{Q}_n \cup \overrightarrow{\partial \mathcal{Q}_n}.$$

Our second pillar requires us to construct a finite sequence of unoriented random graphs such that each graph of the sequence is a bond percolation with parameter β. The last percolation in this finite sequence is such that the vertex set of its connected component containing $\mathbf{0}$ is equal to $\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$. Start with a bond percolation on $\mathcal{H}_N$ with parameter β that is independent of the random variables $\{Z_i^{\boldsymbol{s}} : i \in [N], \boldsymbol{s} \in \mathcal{V}_N\}$. Call the resulting graph $\mathcal{B}_1$. We can now define recursively our process $\mathcal{B} = (\mathcal{B}_n)_{n\in\mathbb{N}}$. For any $e \in \mathcal{E}_N$, let $\mathcal{B}_k\{e\}$ be the status (either open or closed) of edge e in $\mathcal{B}_k$. We obtain $\mathcal{B}_{n+1}$ from $\mathcal{B}_n$ by updating edges in $\mathcal{E}_N$ if and only if they connect a vertex from $\mathcal{Q}_k$ to a vertex in $\partial\mathcal{Q}_k$. More precisely, for any pair of vertices $\boldsymbol{u}, \boldsymbol{v} \in \mathcal{V}_N$ such that $\boldsymbol{u} \sim \boldsymbol{v}$,

$$\mathcal{B}_{k+1}\{[\boldsymbol{u}, \boldsymbol{v}]\} = \begin{cases} \text{open}, & \boldsymbol{u} \in \mathcal{Q}_k, \boldsymbol{v} \in \partial\mathcal{Q}_k, \boldsymbol{u} \to \boldsymbol{v}, \\ \text{closed}, & \boldsymbol{u} \in \mathcal{Q}_k, \boldsymbol{v} \in \partial\mathcal{Q}_k, \{\boldsymbol{u} \to \boldsymbol{v}\}^c, \\ \mathcal{B}_k\{[\boldsymbol{u}, \boldsymbol{v}]\}, & \text{otherwise.} \end{cases}$$

As the status of the edges is updated independently of the original configuration and with i.i.d. Bernoulli random variables with parameter β, we have that $\mathcal{B}_{k+1}$ is still a bond percolation with parameter β.

Notice that in the worst case scenario, each of these processes explores the entirety of $\mathcal{H}_N$ in 2^N iterations. That is, $\mathcal{B}_{k+1} = \mathcal{B}_k$ and $\mathcal{Q}_{k+1} = \mathcal{Q}_k$ for all $k \geq 2^N$. Hence, set $\mathcal{B}_N^\beta = \mathcal{B}_{2^N}$. By construction, $\mathcal{Q}_{2^N} = \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$, but by the above procedure, we also know that $\mathcal{Q}_{2^N}$ is exactly the set of vertices in the connected component of the percolation graph $\mathcal{B}_{2^N}$ that contains $\mathbf{0}$. Recalling that $\mathcal{L}_N^{\beta,\mathbf{0}}$ is the set of vertices in the connected component that contains $\mathbf{0}$ in the percolation graph $\mathcal{B}_N^\beta$, we have that $\mathcal{L}_N^{\beta,\mathbf{0}} = \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$. □

Define $\mathcal{L}_N^\beta$ to be the set of vertices in the giant component of the percolation $\mathcal{B}_N^\beta$ introduced in Proposition 2.7. We call the complement of the giant component the *fragment* of the percolation, and denote it by $\mathcal{M}_N^\beta$. With an abuse of notation, we will also use $\mathcal{M}_N^\beta$ to denote the set of vertices of the fragment, i.e., $\mathcal{V}_N \setminus \mathcal{L}_N^\beta$. We will also make use of the following lemma.

Lemma 2.8 (McDiarmid, Scott, and Withers [31, Theorem 1(a)]). *For each $\varepsilon > 0$ and for a fixed $\beta \in (0, 1/2)$, the percolation $\mathcal{B}_N^\beta$ satisfies*

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\left| \left|\mathcal{M}_N^\beta\right| - \mathbb{E}\left[\left|\mathcal{M}_N^\beta\right|\right] \right| \geq \varepsilon\sqrt{N(2(1-\beta))^N} \right) < \infty,$$

where $\mathbb{E}\left[\left|\mathcal{M}_N^\beta\right|\right] = (2(1-\beta))^N(1+O(N(1-\beta)^N))$.

Proposition 2.9. *For any* $\beta \in (0, 1/2]$, *we have*

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathcal{L}_N^\beta \neq \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}\right) < \infty.$$

Proof. We first prove the proposition for $\beta \in (0, 1/2)$. In virtue of Proposition 2.7, we have $\mathcal{L}_N^{\beta,\mathbf{0}} = \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$, and in order to obtain our result it is enough to prove

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathcal{L}_N^\beta \neq \mathcal{L}_N^{\beta,\mathbf{0}}\right) < \infty.$$

By symmetry, we have

$$\mathbb{P}\left(\mathbf{0} \in \mathcal{M}_N^\beta \;\middle|\; \left|\mathcal{M}_N^\beta\right|\right) = \frac{\left|\mathcal{M}_N^\beta\right|}{2^N}.$$

Hence,

$$\mathbb{P}\left(\mathcal{L}_N^\beta \neq \mathcal{L}_N^{\beta,\mathbf{0}} \;\middle|\; \left|\mathcal{M}_N^\beta\right|\right) = \mathbb{P}\left(\mathbf{0} \in \mathcal{M}_N^\beta \;\middle|\; \left|\mathcal{M}_N^\beta\right|\right) = \frac{\left|\mathcal{M}_N^\beta\right|}{2^N}. \tag{2.1}$$

Set

$$\zeta_N := \mathbb{E}\left[\left|\mathcal{M}_N^\beta\right|\right] + \varepsilon\sqrt{N(2(1-\beta))^N}.$$

Using Lemma 2.8 in conjunction with Eq. (2.1), we obtain

$$\begin{aligned}\sum_{N=1}^{\infty} \mathbb{P}\left(\mathcal{L}_N^\beta \neq \mathcal{L}_N^{\beta,\mathbf{0}}\right) &\leq \sum_{N=1}^{\infty} \mathbb{P}\left(\left\{\mathcal{L}_N^\beta \neq \mathcal{L}_N^{\beta,\mathbf{0}}\right\} \cap \left\{\left|\mathcal{M}_N^\beta\right| < \zeta_N\right\}\right) + \sum_{N=1}^{\infty} \mathbb{P}\left(\left|\mathcal{M}_N^\beta\right| \geq \zeta_N\right) \\ &\leq \sum_{N=1}^{\infty} \frac{\zeta_N}{2^N} + \sum_{N=1}^{\infty} \mathbb{P}\left(\left|\mathcal{M}_N^\beta\right| \geq \zeta_N\right) < \infty.\end{aligned}$$

To complete our result, we turn our attention to the $\beta = 1/2$ case. It is well known that there exists a coupling such that if $\beta < \beta'$ then $\mathcal{L}_N^\beta \subseteq \mathcal{L}_N^{\beta'}$ (Bollobás [8, Theorem 2.1]). As $\mathcal{L}_N^\beta$ does not coincide with $\mathcal{L}_N^{\beta,\mathbf{0}}$ if and only if $\mathbf{0}$ is not an element of $\mathcal{L}_N^\beta$, it follows that $\mathbb{P}\left(\mathcal{L}_N^\beta \neq \mathcal{L}_N^{\beta,\mathbf{0}}\right)$ is nonincreasing in β. Thus, for any $\beta \in (0, 1/2)$,

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathcal{L}_N^{1/2} \neq \mathcal{L}_N^{1/2,\mathbf{0}}\right) \leq \sum_{N=1}^{\infty} \mathbb{P}\left(\mathcal{L}_N^\beta \neq \mathcal{L}_N^{\beta,\mathbf{0}}\right) < \infty.$$

□

Hence, for large enough N, we have that the giant component and the component containing $\mathbf{0}$ coincide. Combining this with Proposition 2.7, our bond percolation's giant component only disagrees with the vertex set of accessible strategies from $\mathbf{0}$ finitely often.

We will make use of the following lemma several times throughout our proofs.

Lemma 2.10 (McDiarmid, Scott, and Withers [31, Theorem 2(a)]). *Fix $\beta \in (0, 1/2]$. For any $r \in \mathbb{N}$ and any $\boldsymbol{v} \in \mathcal{V}_N$, call*

$$B_r(\boldsymbol{v}) := \{\boldsymbol{u} : h(\boldsymbol{v}, \boldsymbol{u}) \leq r\}$$

where h is the Hamming distance. Set

$$m_\beta := \left\lfloor \frac{1}{-\log_2(1-\beta)} \right\rfloor.$$

Then there exists $\bar{\delta} > 0$ such that, for any $\delta < \bar{\delta}$, we have

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\exists \boldsymbol{v} : \left| B_{\lceil \delta N \rceil}(\boldsymbol{v}) \setminus \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \right| > m_\beta \right) < \infty.$$

We say that a vertex is *isolated* if it has degree zero in the graph induced by the percolation. Erdős and Spencer [20] analysed the asymptotic behaviour of $\mathcal{B}_M^{1/2}$ (the nonatomic case, or $\alpha = 0$) and showed that the random graph is connected with probability tending to one. On further inspection of their proof, it is evident that, still with probability tending to one, the giant component of this percolation contains all the vertices in $\mathcal{V}_N$ with the exception of some isolated vertices in the random graph $\mathcal{B}_N^{1/2}$. As the following result was not explicitly stated in the aforementioned paper, we include a proof for the sake of completeness.

Proposition 2.11 (Erdős and Spencer [20]). *Let $\mathcal{I}_N$ denote the set of isolated vertices in $\mathcal{B}_N^{1/2}$. Then*

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathcal{M}_N^{1/2} \neq \mathcal{I}_N\right) < \infty. \tag{2.2}$$

Moreover, for all $k \in \mathbb{N}$,

$$\lim_{N \to \infty} \mathbb{P}\left(|\mathcal{I}_N| = k\right) = \frac{e^{-1}}{k!}. \tag{2.3}$$

Proof. Choose δ as in Lemma 2.10. If $\{\mathcal{M}_N^{1/2} \neq \mathcal{I}_N\}$ then there exists $\boldsymbol{v}$ such that $\left| B_{\lceil \delta N \rceil}(\boldsymbol{v}) \setminus \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \right| > 1$. Combining Lemma 2.10 (case $\beta = 1/2$) with Proposition 2.7, we have

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\exists \boldsymbol{v} : \left| B_{\lceil \delta N \rceil}(\boldsymbol{v}) \setminus \overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \right| > 1 \right) < \infty,$$

which proves Eq. (2.2). To prove Eq. (2.3), we reason as follows. The PNEs are the vertices of $\overrightarrow{\mathcal{H}}_N$ that are incident to only incoming edges. When $\alpha = 0$, the number of PNEs has the same distribution as the number of vertices in $\overrightarrow{\mathcal{H}}_N$ that are incident only to outgoing edges. Call $\overrightarrow{\mathcal{I}}$ the set of such vertices. It is well known that, when $\alpha = 0$, the number of PNEs converges in distribution to a **Poisson**(1) random variable (Rinott and Scarsini [40]). Hence, for $k \in \mathbb{N}$,

$$\lim_{N \to \infty} \mathbb{P}\left(\left|\overrightarrow{\mathcal{I}}_N\right| = k\right) = \frac{e^{-1}}{k!}. \tag{2.4}$$

Now we prove that, for all N large enough, $\mathcal{I}_N = \overrightarrow{\mathcal{I}}_N$. If $\mathbf{0} \notin \overrightarrow{\mathcal{I}}_N$, then all the vertices in $\overrightarrow{\mathcal{I}}_N$ are not accessible from $\mathbf{0}$. As

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathbf{0} \in \overrightarrow{\mathcal{I}}_N\right) = \sum_{N=1}^{\infty} \frac{1}{2^N} < \infty,$$

we can use the first Borel-Cantelli lemma to conclude that for all N large enough, $\overrightarrow{\mathcal{I}}_N \subseteq \overrightarrow{\mathcal{M}}_N^{1/2,\mathbf{0}}$. Using Proposition 2.7, we have that $\overrightarrow{\mathcal{I}} \subseteq \mathcal{V}_N \backslash \mathcal{L}_N^{1/2,\mathbf{0}}$. Then, using Proposition 2.9 and the first Borel-Cantelli lemma, we have that for N large enough, $\overrightarrow{\mathcal{I}}_N \subseteq \mathcal{M}_N^{1/2}$. Conversely, each isolated point $\boldsymbol{v}$ in $\mathcal{B}_N^{1/2}$ whose neighbours are all in $\mathcal{L}_N^{1/2,\mathbf{0}}$ satisfies $\boldsymbol{v} \in \overrightarrow{\mathcal{I}}_N$. Hence, $\overrightarrow{\mathcal{I}}_N \neq \mathcal{M}_N^{1/2}$ only if there exist two elements $\boldsymbol{u}, \boldsymbol{v} \in \mathcal{M}_N^{1/2}$ such that $h(\boldsymbol{u}, \boldsymbol{v}) = 1$, that is, $\boldsymbol{u}$ and $\boldsymbol{v}$ are neighbours in $\mathcal{H}_N$. Using Section 2.3, we have

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\overrightarrow{\mathcal{I}}_N \neq \mathcal{M}_N^{1/2}\right) < \infty.$$

Finally, combining this with Eq. (2.2) gives

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\overrightarrow{\mathcal{I}}_N \neq \mathcal{M}_N^{1/2}\right) < \infty.$$

We once again use the first Borel-Cantelli lemma, and in conjunction with Eq. (2.4), obtain our desired result in Eq. (2.3). ☐

2.4 Proofs

We begin this section with a proof of Theorem 2.1.

Proof of Theorem 2.1.

(a) When $\alpha = 0$, convergence of the number of SPNEs to a Poisson distribution as the number of players increases was proved by Rinott and Scarsini [40] for any fixed number of strategies.

(b) We can write

$$|\mathsf{E}_N^*| = \sum_{\boldsymbol{s} \in \mathsf{S}} \mathbb{1}_{\{\boldsymbol{s} \in \mathsf{E}_N^*\}}$$

where $\mathbb{1}_A$ denotes the indicator function of the event A. Because $\mathbb{P}(\boldsymbol{s} \in \mathsf{E}_N^*) = \beta^N$ for every $\boldsymbol{s} \in \mathcal{V}_N$, we have

$$\mathbb{E}[|\mathsf{E}_N^*|] = (2\beta)^N,$$

from which Markov's inequality implies

$$\mathbb{P}(|\mathsf{E}^*| \geq 1) \leq \mathbb{E}[|\mathsf{E}_N^*|] = (2\beta)^N.$$

As $\alpha > 0$, $\beta < 1/2$, so the upper bound converges to 0 geometrically fast. Hence,

$$\sum_{N=1}^{\infty} \mathbb{P}(|\mathsf{E}^*| \geq 1) < \infty$$

and by the first Borel-Cantelli lemma, the event $\{|\mathsf{E}_N^*| < 1\}$ holds true for all N large enough. As cardinality is a nonnegative integer, it must be zero for all large N.

□

Now that we have established a connection between percolation and our random games, we can return to the results stated earlier in this chapter. We will soon address Theorem 2.2, but first need to introduce the notion of distance between random variables.

Definition 2.12. Given two probability measures $\mathbb{P}, \mathbb{Q}$ on $\mathbb{N}$, their *total variation distance*, denoted ρ_{TV}, is defined as

$$\rho_{TV}(\mathbb{P}, \mathbb{Q}) := \sup_{A \subseteq \mathbb{N}} |\mathbb{P}(A) - \mathbb{Q}(A)|.$$

Definition 2.13. Given two probability measures $\mathbb{P}, \mathbb{Q}$ on $\mathbb{R}$, their *Kolmogorov distance*, denoted κ, is defined as

$$\kappa(\mathbb{P}, \mathbb{Q}) := \sup_{x \in \mathbb{R}} |\mathbb{P}((-\infty, x]) - \mathbb{Q}((-\infty, x])|.$$

Remark. Notice that Kolmogorov distance is a supremum over a collection of sets $[0, k]$ for $k \in \mathbb{N}$ when dealing with integer-valued distributions, whereas total variation distance is a supremum over a richer collection of sets. It immediately follows that, for two probability measures $\mathbb{P}, \mathbb{Q}$ on $\mathbb{N}$,

$$\kappa(\mathbb{P}, \mathbb{Q}) \leq \rho_{TV}(\mathbb{P}, \mathbb{Q}). \tag{2.5}$$

With an abuse of language and notation, we will often speak of total variation distance and Kolmogorov distance of two random variables to indicate the total variation distance of their laws.

In order to prove Theorem 2.2, we will need the following lemma.

Lemma 2.14 (McDiarmid, Scott, and Withers [31, Theorem 4(c)]). *Call $\mathcal{C}_j$ the number of connected components with exactly j vertices in the percolation $\mathcal{B}_N^\beta$, and let $\mu_j = \mathbb{E}[\mathcal{C}_j]$. Then for $\beta \in (0, 1/2)$,*

$$\rho_{TV}(\mathcal{C}_j, \mathbf{Poisson}(\mu_j)) = O\big(N^j(1-\beta)^{jN}\big).$$

Proof of Theorem 2.2. Consider the random partially oriented hypercube $\overrightarrow{\mathcal{H}}_N^\beta$ and invert the orientation of each edge, while leaving the unoriented edges unchanged. The new random partially oriented hypercube $\overleftarrow{\mathcal{H}}_N^\beta$ has the same law as $\overrightarrow{\mathcal{H}}_N^\beta$ by symmetry. Moreover, a vertex in $\overleftarrow{\mathcal{H}}_N^\beta$ is a PNE if and only if it is an isolated point in $\overrightarrow{\mathcal{H}}_N^\beta$.

Consider the percolation associated with $\overleftarrow{\mathcal{H}}_N^\beta$. In virtue of our precious reasoning, it is enough to establish a limit theorem for the number of isolated vertices of this percolation. By applying Lemma 2.14 with $j = 1$, we obtain that there exists a constant $K_{1,\alpha}$, which depends on α but not on N, such that

$$\rho_{TV}\big(|\mathsf{E}_N|, \mathbf{Poisson}\big(2^N(1-\beta)^N\big)\big) \leq K_{1,\alpha}\big(N(1-\beta)^N\big),$$

which by Eq. (2.5) gives

$$\kappa\big(|\mathsf{E}_N|, \mathbf{Poisson}\big(2^N(1-\beta)^N\big)\big) \leq K_{1,\alpha}\big(N(1-\beta)^N\big). \tag{2.6}$$

Notice that a $\mathbf{Poisson}\big(2^N(1-\beta)^N\big)$ random variable can be expressed as the sum of $\lfloor 2^N(1-\beta)^N \rfloor$ independent $\mathbf{Poisson}(1)$ random variables plus a small remainder. Consider a sequence of i.i.d. random variables $(X_n)_{n\in\mathbb{N}}$ where $X_n = \mathbf{Poisson}(1) - 1$. Clearly $\mathbb{E}[X_n] = 0$, $\mathbb{E}\left[{X_n}^2\right] = 1$ and $\mathbb{E}[|X_n|^3] = 1 + 2e^{-1}$. Hence, we can use the Berry-Esseen theorem and obtain

$$\kappa\left(\frac{\mathbf{Poisson}\big(2^N(1-\beta)^N\big) - 2^N(1-\beta)^N}{\sqrt{2^N(1-\beta)^N}}, \mathbf{Normal}(0,1)\right) \leq \frac{K_{2,\alpha}}{\sqrt{2^N(1-\beta)^N}}.$$

Note that we follow the convention that a $\mathbf{Normal}(\mu, \sigma^2)$ random variable has mean μ and variance σ^2. Call F_n the distribution function of $\mathbf{Poisson}\big(2^N(1-\beta)^N\big)/\sqrt{2^N(1-\beta)^N}$. Note that

$$\begin{aligned} F_n(x) &= \mathbb{P}\left(\frac{\mathbf{Poisson}\big(2^N(1-\beta)^N\big) - 2^N(1-\beta)^N}{\sqrt{2^N(1-\beta)^N}} \leq x\right) \\ &= \mathbb{P}\left(\mathbf{Poisson}\big(2^N(1-\beta)^N\big) \leq \sqrt{2^N(1-\beta)^N}x + 2^N(1-\beta)^N\right). \end{aligned}$$

Meanwhile,

$$
\begin{aligned}
\Phi(x) &= \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right) du \\
&= \int_{-\infty}^{\sqrt{2^N(1-\beta)^N}x+2^N(1-\beta)^N} \frac{1}{\sqrt{2\pi 2^N(1-\beta)^N}} \exp\left(-\frac{\big(u-2^N(1-\beta)^N\big)^2}{2(2^N(1-\beta)^N)}\right) du
\end{aligned}
$$

which is the distribution function of a $\mathbf{Normal}\big(2^N(1-\beta)^N, 2^N(1-\beta)^N\big)$ random variable evaluated at $\sqrt{2^N(1-\beta)^N}x + 2^N(1-\beta)^N$. Hence,

$$
\begin{aligned}
&\kappa\left(\frac{\mathbf{Poisson}\big(2^N(1-\beta)^N\big) - 2^N(1-\beta)^N}{\sqrt{2^N(1-\beta)^N}}, \mathbf{Normal}(0,1)\right) \\
&\qquad = \kappa\big(\mathbf{Poisson}\big(2^N(1-\beta)^N\big), \mathbf{Normal}\big(2^N(1-\beta)^N, 2^N(1-\beta)^N\big)\big) \\
&\qquad \leq \frac{K_{2,\alpha}}{\sqrt{2^N(1-\beta)^N}}.
\end{aligned} \tag{2.7}
$$

Using the triangular inequality with Eqs. (2.6) and (2.7) yields

$$
\begin{aligned}
&\kappa\big(|\mathsf{E}_N|, \mathbf{Normal}\big(2^N(1-\beta)^N, 2^N(1-\beta)^N\big)\big) \\
&\quad \leq \kappa\big(|\mathsf{E}_N|, \mathbf{Poisson}\big(2^N(1-\beta)^N\big)\big) \\
&\qquad + \kappa\big(\mathbf{Poisson}\big(2^N(1-\beta)^N, \mathbf{Normal}\big(2^N(1-\beta)^N, 2^N(1-\beta)^N\big)\big)\big) \\
&\quad \leq K_{1,\alpha}\big(N(1-\beta)^N\big) + \frac{K_{2,\alpha}}{\sqrt{2^N(1-\beta)^N}}.
\end{aligned}
$$

After making the substitution $1+\alpha = 2(1-\beta)$ and some slight rearranging, we obtain our desired result. □

Proof of Theorem 2.3. We first want to show that, for any $\varepsilon > 0$ small enough,

$$
\sum_{N=1}^{\infty} \mathbb{P}\left(|\mathsf{E}_N| \leq (1+\alpha)^N - (1+\alpha)^{N/2}(1+\varepsilon)^N\right) < \infty. \tag{2.8}
$$

Choosing $\varepsilon < \sqrt{1+\alpha} - 1$ gives

$$
(1+\alpha)^N - (1+\alpha)^{N/2}(1+\varepsilon)^N = (1+\alpha)^N(1+o(1)).
$$

Therefore by establishing Eq. (2.8), we would prove that the number of PNEs grows at least as fast as $(1+\alpha)^N$. Recall the standard inequality $\Phi(-x) \leq \phi(x)/x$ where ϕ is the density function of a standard normal random variable. We then have that

$$
\Phi\big(-(1+\varepsilon)^N\big) \leq \frac{1}{\sqrt{2\pi}(1+\varepsilon)^N} \exp\left(-\frac{(1+\varepsilon)^{2N}}{2}\right) < \frac{1}{(1+\varepsilon)^N}.
$$

Hence,

$$\begin{aligned}
&\sum_{N=1}^{\infty} \mathbb{P}\left(|\mathsf{E}_N| \leq (1+\alpha)^N - (1+\alpha)^{N/2}(1+\varepsilon)^N\right) \\
&\quad = \sum_{N=1}^{\infty}\left(\mathbb{P}\left(\frac{|\mathsf{E}_N| - (1+\alpha)^N}{(1+\alpha)^{N/2}} \leq -(1+\varepsilon)^N\right) - \Phi\big(-(1+\varepsilon)^N\big)\right) + \sum_{N=1}^{\infty}\Phi\big(-(1+\varepsilon)^N\big) \\
&\quad < \sum_{N=1}^{\infty} K_\alpha N \max\left(\frac{1+\alpha}{2}, \frac{1}{(1+\alpha)^{N/2}}\right)^N + \sum_{N=1}^{\infty}\frac{1}{(1+\varepsilon)^N} < \infty.
\end{aligned}$$

To complete the proof, we also need to show that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(|\mathsf{E}_N| \geq (1+\alpha)^N + (1+\alpha)^{N/2}(1+\varepsilon)^N\right) < \infty,$$

as this would show that the number of PNEs grows no faster than $(1+\alpha)^N$. However, notice that

$$\mathbb{P}\left(|\mathsf{E}_N| \geq (1+\alpha)^N + (1+\alpha)^{N/2}(1+\varepsilon)^N\right) = 1 - \mathbb{P}\left(\frac{|\mathsf{E}_N| - (1+\alpha)^N}{(1+\alpha)^{N/2}} < (1+\varepsilon)^N\right).$$

Hence,

$$\begin{aligned}
&\sum_{N=1}^{\infty} \mathbb{P}\left(|\mathsf{E}_N| \geq (1+\alpha)^N + (1+\alpha)^{N/2}(1+\varepsilon)^N\right) \\
&\quad = \sum_{N=1}^{\infty}\left(1 - \mathbb{P}\left(\frac{|\mathsf{E}_N| - (1+\alpha)^N}{(1+\alpha)^{N/2}} < (1+\varepsilon)^N\right) - \big(1 - \Phi\big((1+\varepsilon)^N\big)\big)\right) \\
&\quad\quad + \sum_{N=1}^{\infty}\big(1 - \Phi\big((1+\varepsilon)^N\big)\big) \\
&\quad < \sum_{N=1}^{\infty} K_\alpha N \max\left(\frac{1+\alpha}{2}, \frac{1}{(1+\alpha)^{N/2}}\right)^N + \sum_{N=1}^{\infty}\frac{1}{(1+\varepsilon)^N} < \infty,
\end{aligned}$$

the last line coming from the fact that $\Phi(x) = 1 - \Phi(-x)$. □

We now turn our attention to results regarding the BRD.

Proof of Theorem 2.5.

(a) We first deal with the case $\alpha = 0$. Here, $\beta = 1/2$, so $m_\beta = 1$. In this case, there are no unoriented edges in $\overrightarrow{\mathcal{H}}_N^{\beta}$. If a PNE belongs to $\overrightarrow{\mathcal{M}}_N^{1/2,\mathbf{0}}$, then all of its neighbours must also belong to $\overrightarrow{\mathcal{M}}_N^{1/2,\mathbf{0}}$; if a neighbour was in fact accessible from $\mathbf{0}$, then so too would be the PNE. This means there exists a ball of radius 1 with $N+1$ vertices in $\overrightarrow{\mathcal{M}}_N^{1/2,\mathbf{0}}$. Hence, by Lemma 2.10,

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathsf{E}_N \cap \overrightarrow{\mathcal{M}}_N^{1/2,\mathbf{0}} \neq \emptyset\right) \leq \sum_{N=1}^{\infty} \mathbb{P}\left(\exists \boldsymbol{v} : \left|B_{\lceil \delta N \rceil}(\boldsymbol{v}) \setminus \overrightarrow{\mathcal{L}}_N^{1/2,\mathbf{0}}\right| > 1\right) < \infty.$$

Now consider the case $0 < \alpha < 1/2$, and suppose that $\boldsymbol{v}$ both is a PNE and belongs to $\overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$. Either of the two following conditions must then hold:

(i) $\boldsymbol{v}$ is adjacent to more than m_β oriented edges, which implies that the size of the connected component of $\overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$ containing $\boldsymbol{v}$ is larger than m_β; or

(ii) $\boldsymbol{v}$ is adjacent to at most m_β oriented edges.

Clearly, (i) can be treated by Lemma 2.10. To resolve (ii), let Y_N be the number of vertices that are incident to at least $N - m_\beta$ unoriented edges. The number of unoriented edges a given vertex is adjacent to is binomially distributed, as each edge is oriented independently of one another. Markov's inequality yields

$$\mathbb{P}\left(Y_N \geq 1\right) \leq 2^N \sum_{k=N-m_\beta}^{N} \binom{N}{k} \alpha^k (1-\alpha)^{N-k} \leq K_{3,\alpha} N^{m_\beta} 2^N \alpha^{N-m_\beta}$$

for some constant $K_{3,\alpha}$. As $\alpha < 1/2$,

$$\sum_{N=1}^{\infty} \mathbb{P}\left(Y_N \geq 1\right) < \infty.$$

Hence, by a union bound,

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\mathsf{E}_N \cap \overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}} \neq \emptyset\right)$$
$$\leq \sum_{N=1}^{\infty} \mathbb{P}\left(\exists \boldsymbol{v} : \left|B_{\lceil \delta N \rceil}(\boldsymbol{v}) \setminus \overrightarrow{\mathcal{L}}_N^{1/2,\mathbf{0}}\right| > m_\beta\right) + \sum_{N=1}^{\infty} \mathbb{P}\left(Y_N \geq 1\right) < \infty.$$

(b) For the case $\alpha = 1/2$, we introduce a new percolation on $\mathcal{H}_N$, denoted by $\widetilde{\mathcal{B}}_N$, defined such that an edge is open if and only if it is oriented in $\overrightarrow{\mathcal{H}}_N^{1/2}$. Call $\widetilde{\mathcal{L}}_N^{\mathbf{0}}$ the connected component of $\widetilde{\mathcal{B}}_N$ that contains $\mathbf{0}$. Any isolated vertex in $\widetilde{\mathcal{B}}_N$ is a PNE in $\overrightarrow{\mathcal{H}}_N^{1/2}$, as it is incident to nonoriented edges only. Notice also that the parameter of this percolation is $1/2$, so we can apply Proposition 2.11, and obtain that the number of PNEs outside of $\widetilde{\mathcal{L}}_N^{\mathbf{0}}$ asymptotically follows a Poisson distribution with parameter 1. Owing to the fact that any edge that is open in $\mathcal{B}_N^{1/2}$ is also open in $\widetilde{\mathcal{B}}_N$, we have that $\mathcal{L}_N^{1/2,\mathbf{0}} \subseteq \widetilde{\mathcal{L}}_N^{\mathbf{0}}$. Finally, using Proposition 2.7, we have that $\overrightarrow{\mathcal{L}}_N^{1/2,\mathbf{0}} \subseteq \widetilde{\mathcal{L}}_N^{\mathbf{0}}$, and it immediately follows that the number of PNEs outside of $\overrightarrow{\mathcal{L}}_N^{1/2,\mathbf{0}}$ is stochastically larger than a Poisson random variable with parameter 1.

(c) For $\boldsymbol{v} \in \mathcal{V}_N$, let $\theta_{\boldsymbol{v}}$ be the event in which $\boldsymbol{v}$ is incident to only unoriented edges in $\overrightarrow{\mathcal{H}}_N^{\beta}$, and define

$$\Theta_N = \sum_{\boldsymbol{v} \in \mathcal{V}_N} \mathbb{1}_{\theta_{\boldsymbol{v}}}.$$

Note that $\mathcal{H}_N^\beta$ can be decomposed as

$$\mathcal{V}_N = \mathcal{V}_N^{\text{even}} \dot{\cup} \mathcal{V}_N^{\text{odd}},$$

where $\boldsymbol{v} \in \mathcal{V}_N^{\text{even}}$ if and only if $h(\boldsymbol{v}, \mathbf{0})$ is even. By symmetry, we know that $|\mathcal{V}_N^{\text{even}}| = |\mathcal{V}_N^{\text{odd}}| = 2^{N-1}$. Edges only connect vertices from $\mathcal{V}_N^{\text{even}}$ to $\mathcal{V}_N^{\text{odd}}$ or vice versa, so no two vertices belonging to exactly one of the sets can be neighbours.

Our first goal is to prove that $\{\theta_{\boldsymbol{v}} : \boldsymbol{v} \in \mathcal{V}_N^{\text{even}}\}$ is a collection of independent events. We will proceed inductively, as if we can show that for any $\mathcal{U} \subseteq \mathcal{V}_N^{\text{even}}$ and any $\boldsymbol{u} \in \mathcal{U}$ that

$$\mathbb{P}\left(\bigcap_{\boldsymbol{v} \in \mathcal{U}} \theta_{\boldsymbol{u}}\right) = \mathbb{P}(\theta_{\boldsymbol{u}})\, \mathbb{P}\left(\bigcap_{\boldsymbol{v} \in \mathcal{U} \setminus \{\boldsymbol{u}\}} \theta_{\boldsymbol{v}}\right), \tag{2.9}$$

then it is easy to conclude that

$$\mathbb{P}\left(\bigcap_{\boldsymbol{v} \in \mathcal{U}} \theta_{\boldsymbol{v}}\right) = \prod_{\boldsymbol{v} \in \mathcal{U}} \mathbb{P}(\theta_{\boldsymbol{v}}).$$

For each $i \in [N]$, denote by $\boldsymbol{u}^{(i)}$ the unique vertex satisfying $\boldsymbol{u} \sim_i \boldsymbol{u}^{(i)}$. Not only is the orientation of $\left[\boldsymbol{u}, \boldsymbol{u}^{(i)}\right]$ determined exclusively by the random variables $Z_i^{\boldsymbol{u}}$ and $Z_i^{\boldsymbol{u}^{(i)}}$, but neither of these random variables are used in determining the orientation for any other edges in our graph. Furthermore, as each of these random variables are independent, we immediately obtain our desired result in Eq. (2.9).

Now that $\{\theta_{\boldsymbol{v}} : \boldsymbol{v} \in \mathcal{V}_N^{\text{even}}\}$ has been proven to be a collection of independent events, we know that $\sum_{\boldsymbol{v} \in \mathcal{V}_N^{\text{even}}} \mathbb{1}_{\theta_{\boldsymbol{v}}}$ is a **Binomial**$\left(2^{N-1}, \alpha\right)$ random variable. Moreover, $\Theta_N \geq \sum_{\boldsymbol{v} \in \mathcal{V}_N^{\text{even}}} \mathbb{1}_{\theta_{\boldsymbol{v}}}$, so Θ_N is stochastically larger than a **Binomial**$\left(2^{N-1}, \alpha\right)$ random variable. Finally, as any vertex (excluding $\mathbf{0}$) that is incident to only unoriented edges is both a PNE and in $\overrightarrow{\mathcal{M}}_N^{\beta, \mathbf{0}}$, we have that for any fixed $K > 0$,

$$\lim_{N \to \infty} \mathbb{P}\left(\left|\mathsf{E}_N \cap \overrightarrow{\mathcal{M}}_N^{\beta, \mathbf{0}}\right| > K\right) = 1.$$

□

Recall the definition of a trap from Definition 1.7. Notice that a trap cannot contain a PNE as every vertex must be accessible from every other vertex in a trap.

Proof of Theorem 2.6. As BRD_N fails to converge if and only if it enters a trap, it is enough to prove the stronger result

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\text{a trap exists in } \overrightarrow{\mathcal{H}}_N^\beta\right) < \infty.$$

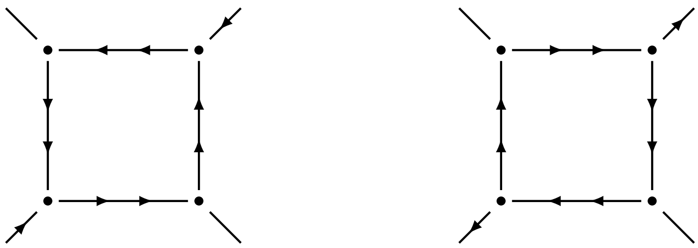

Figure 2.6: A trap of size 4 (left) and the corresponding subgraph after reversing the orientation on the edges (right). Notice that, after reversing the edges, there is no way for the process to enter the cycle from the outside.

In order to obtain this result, we need to investigate how traps behave under orientation reversal. To this end, invert the orientation of $\overrightarrow{\mathcal{H}}_N^{\beta}$ while leaving the unoriented edges unchanged to obtain $\overleftarrow{\mathcal{H}}_N^{\beta}$, and consider the partition

$$\mathcal{V}_N = \overleftarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \dot{\cup} \overleftarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$$

where $\overleftarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$ is the set of all vertices that are accessible from $\mathbf{0}$ in the oriented graph $\overleftarrow{\mathcal{H}}_N^{\beta}$. Conversely, all vertices that are not accessible from $\mathbf{0}$ are contained within $\overleftarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$.

Fix a trap $\overrightarrow{\mathcal{T}} \subseteq \overrightarrow{\mathcal{H}}_N^{\beta}$, noting that each edge connecting $\overrightarrow{\mathcal{T}}$ to its boundary $\partial\overrightarrow{\mathcal{T}}$ is either unoriented or points toward $\overrightarrow{\mathcal{T}}$. Let $\overleftarrow{\mathcal{T}}$ be the subgraph in $\overleftarrow{\mathcal{H}}_N^{\beta}$ corresponding to $\overrightarrow{\mathcal{T}}$. More precisely:

- $\mathcal{V}\left(\overleftarrow{\mathcal{T}}\right) = \mathcal{V}\left(\overrightarrow{\mathcal{T}}\right)$;
- if $[\boldsymbol{u}, \boldsymbol{v}] \in \overrightarrow{\mathcal{E}}_N^{\beta}$ then $[\boldsymbol{u}, \boldsymbol{v}] \in \overleftarrow{\mathcal{E}}_N^{\beta}$; and
- if $\overrightarrow{[\boldsymbol{u}, \boldsymbol{v}]} \in \overrightarrow{\mathcal{E}}_N^{\beta}$ then $\overleftarrow{[\boldsymbol{u}, \boldsymbol{v}]} \in \overleftarrow{\mathcal{E}}_N^{\beta}$.

We call $\overleftarrow{\mathcal{T}}$ a reversed trap. All edges connecting $\partial\overleftarrow{\mathcal{T}}$ and $\overleftarrow{\mathcal{T}}$ are either unoriented or oriented away from $\overleftarrow{\mathcal{T}}$ (see Fig. 2.6). Hence, we have the following behaviours:

(i) if $\mathbf{0} \notin \mathcal{V}\left(\overleftarrow{\mathcal{T}}\right)$ then $\mathcal{V}\left(\overleftarrow{\mathcal{T}}\right) \subseteq \overleftarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$; or

(ii) if $\mathbf{0} \in \mathcal{V}\left(\overleftarrow{\mathcal{T}}\right)$ then $\mathcal{V}\left(\overleftarrow{\mathcal{T}}\right) = \overleftarrow{\mathcal{L}}_N^{\beta,\mathbf{0}}$.

To treat (i), we will show that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\text{a reversed trap } \overleftarrow{\mathcal{T}} \text{ exists such that } \mathcal{V}(\overleftarrow{\mathcal{T}}) \subseteq \overleftarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}\right). \tag{2.10}$$

Notice that any trap (and thus, reverse trap) must have at least four vertices. Moreover, $\overleftarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$ and $\overrightarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}$ are identically distributed. We can apply Lemma 2.10 in this setting. Due to our restriction on α, we have that

$$m_\beta = \left\lfloor \frac{1}{-\log_2(1-\beta)} \right\rfloor = \left\lfloor \frac{1}{-\log_2(1/2+\alpha/2)} \right\rfloor \leq 3.$$

Hence,

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\text{a connected component of at least size 4 exists in } \overleftarrow{\mathcal{M}}_N^{\beta,\mathbf{0}}\right) < \infty,$$

which implies Eq. (2.10).

We now focus on (ii) and prove that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \text{ is a trap}\right) < \infty.$$

Since a subgraph is a trap if and only if it contains no PNEs, this is equivalent to showing that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right) < \infty.$$

To prove this, we will find an upper bound for the terms in the above sum via a union bound. If we define, for $\varepsilon < \sqrt{1+\alpha} - 1$,

$$l_N = (1+\alpha)^N - (1+\varepsilon)^N (1+\alpha)^{N/2},$$

then we can write

$$\mathbb{P}\left(\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right) \leq \mathbb{P}\left(|\mathsf{E}_N| < l_N\right) + \mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\}\right). \tag{2.11}$$

By using Eq. (2.8), we immediately get that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(|\mathsf{E}_N| < l_N\right) < \infty, \tag{2.12}$$

so we turn our attention to the remaining probability above.

Let $H_{\boldsymbol{v}}$ be the event in which $\boldsymbol{v}$ is incident to three or more directed edges, and write

$$\begin{aligned} &\mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\}\right) \\ &\qquad \leq \mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcup_{\boldsymbol{v}\in\mathsf{E}_N} H_{\boldsymbol{v}}\right\}\right) \\ &\qquad\quad + \mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcap_{\boldsymbol{v}\in\mathsf{E}_N} {H_{\boldsymbol{v}}}^c\right\}\right). \end{aligned} \tag{2.13}$$

The event

$$\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcup_{\boldsymbol{v}\in\mathsf{E}_N} H_{\boldsymbol{v}}\right\}$$

implies that there exists a PNE in the fragment, say $\boldsymbol{v}$, that is adjacent to at least three directed edges. However, owing to $\boldsymbol{v}$ being a PNE, all of these adjacent edges must be directed towards $\boldsymbol{v}$, so $\boldsymbol{v}$ must be part of a connected component in the fragment of size larger than m_β. Following the proof for (i) by way of Lemma 2.10, we know that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcup_{\boldsymbol{v}\in\mathsf{E}_N} H_{\boldsymbol{v}}\right\}\right) < \infty. \tag{2.14}$$

In words,

$$\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcap_{\boldsymbol{v}\in\mathsf{E}_N} {H_{\boldsymbol{v}}}^c\right\}$$

is the event that every PNE, of which there are at least l_N, exists in the fragment and is adjacent to at most two directed edges. This event implies that there are at least l_N vertices that are adjacent to at most two directed edges, so we have

$$\mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcap_{\boldsymbol{v}\in\mathsf{E}_N} {H_{\boldsymbol{v}}}^c\right\}\right) \leq \mathbb{P}\left(\sum_{\boldsymbol{v}\in\mathcal{V}_N} \mathbb{1}_{H_{\boldsymbol{v}}} \geq l_N\right). \tag{2.15}$$

To deal with this final probability, we can employ Markov's inequality and obtain, for some $\boldsymbol{u} \in \mathcal{V}_N$,

$$\mathbb{P}\left(\sum_{\boldsymbol{v}\in\mathcal{V}_N} \mathbb{1}_{H_{\boldsymbol{v}}} \geq l_N\right) \leq \frac{\mathbb{E}\left[\sum_{\boldsymbol{v}\in\mathcal{V}_N} H_{\boldsymbol{v}}\right]}{l_N} = \frac{2^N \mathbb{P}\left(H_{\boldsymbol{u}}\right)}{l_N}.$$

By direct calculation,

$$\mathbb{P}\left(H_{\boldsymbol{u}}\right) = \sum_{k=0}^{2} \binom{N}{k} (1-\alpha)^k \alpha^{N-k} \leq K_{4,\alpha} N^2 \alpha^N$$

for some constant $K_{4,\alpha}$. It follows that

$$\sum_{N=1}^{\infty} \mathbb{P}\left(\sum_{\boldsymbol{v}\in\mathcal{V}_N} \mathbb{1}_{H_{\boldsymbol{v}}} \geq l_N\right) \leq \sum_{N=1}^{\infty} \frac{K_{4,\alpha} N^2 (2\alpha)^N}{l_N} < \infty. \tag{2.16}$$

Combining Eqs. (2.11), (2.13) and (2.15) and then using the results from Eqs. (2.12), (2.14) and (2.16) finally gives

$$\begin{aligned}
&\sum_{N=1}^{\infty} \mathbb{P}\left(\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right) \\
&\quad \leq \sum_{N=1}^{\infty} \Bigg[\mathbb{P}\left(|\mathsf{E}_N| < l_N\right) + \mathbb{P}\left(\left\{\overrightarrow{\mathcal{L}}_N^{\beta,\mathbf{0}} \cap \mathsf{E}_N = \emptyset\right\} \cap \{|\mathsf{E}_N| \geq l_N\} \cap \left\{\bigcup_{\boldsymbol{v}\in\mathsf{E}_N} H_{\boldsymbol{v}}\right\}\right) \\
&\qquad + \mathbb{P}\left(\sum_{\boldsymbol{v}\in\mathcal{V}_N} \mathbb{1}_{H_{\boldsymbol{v}}} \geq l_N\right)\Bigg] < \infty.
\end{aligned}$$

□

3

Multi-action Games with Two Players

This chapter is based on joint work with Andrea Collevecchio and Kais Hamza titled 'When "Better" is better than "Best"', and has been published in Operations Research Letters. We will expand on the best-response dynamics result included in that paper, as well as provide further clarity for the calculations included in the appendix.

Contrasting to the previous chapter, we now consider games consisting of two players, each of whom can choose from one of K actions. Again denote by $[K]$ the set of integers $\{1, ..., K\}$; we will also write $[k, K]$ to denote the set of integers $\{k, k+1, ..., K\}$. As there are only two players in this model, the set of all action profiles in our game is $\mathsf{S} := [K]^2$. For a action profile $\boldsymbol{s} := (s_1, s_2) \in \mathsf{S}$, the payoff of player i is $Z_i^{\boldsymbol{s}} = Z_i^{s_1,s_2}$, and we will assume for each $i \in \{1, 2\}$ and $\boldsymbol{s} \in \mathsf{S}$ that the payoffs are independent and follow a continuous distribution. Recall that the set of all PNE is denoted E, and that $\mathsf{P}_{\boldsymbol{s}}^{(i)}$ ($\mathsf{M}_{\boldsymbol{s}}^{(i)}$) is the set of all (maximally) profitable deviations available to player i from $\boldsymbol{s}$. Due to our assumption that payoffs are drawn independently from a continuous distribution, $\mathsf{M}_{\boldsymbol{s}}^{(i)}$ will contain at most one action profile almost surely (and is an empty set if there are no profitable deviations for player i from $\boldsymbol{s}$).

The question we will address in this chapter is: with what likelihood do both bRD and BRD converge to a PNE?

3.1 Best Response Dynamics

Theorem 3.1. *For a two-player, K-action game where payoffs are drawn independently from a continuous distribution,*

$$\mathbb{P}\left(\mathsf{BRD}\ \textit{converges to a PNE}\right) \leq \frac{1}{K} + \sqrt{\frac{\pi}{K}}.$$

Proof. For aesthetic purposes, in what follows, we implicitly condition on the event in which $\mathsf{BRD}(0)$ is a best response for **exactly** one player. If it is instead not a best response for either player, after one step it will arrive at a best response for the action which remained constant, and the process behaves as if it were under our implied condition from then on. Finally, the probability that the starting vertex is a best response for both players – i.e., a PNE – is $1/K^2$, and does not influence our final result.

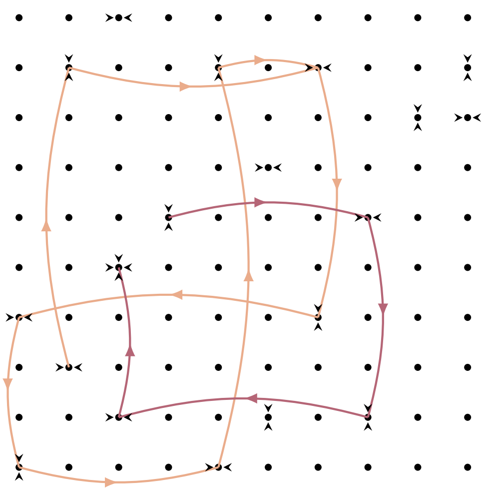

Figure 3.1: Two BRD paths, one enters a trap (peach), the other reaches a PNE (maroon).

First we note that, because the payoffs follow a continuous distribution, with probability one there exists exactly one best response for each action either player chooses. It follows that BRD is trapped if and only if it visits a row or column it has visited previously. Moreover, BRD can only visit at most two action profiles along any given row or column: the action profile from which it enters a row or column; and the corresponding best response. As such, the maximum number of steps BRD can make before revisiting a row or column is at most $2K-2$. Hence, BRD is trapped if and only if it has not reached a PNE by this time.

Due to the fact that there is exactly one best response for each action, BRD must alternate moving along rows or columns at each step (see Fig. 3.1 for an example of this behaviour). It follows that, at time t, BRD must avoid $\lfloor t/2 \rfloor$ rows or columns. As each step is to a best response, each new action profile that BRD reaches is a PNE with probability $1/K$. Letting $\mathcal{R}(t)$ be the set of all action profiles in the rows and columns that BRD has visited by time t, we have

$$\begin{aligned}
&\mathbb{P}\left(\mathsf{BRD}(t+1) \in \mathsf{E}\right) \\
&\qquad = \mathbb{P}\left(\mathsf{BRD}(t) \in \mathsf{E}\right) + \mathbb{P}\left(\mathsf{BRD}(t) \notin \mathsf{E} \cup \mathcal{R}(t-1)\right) \frac{K-1-\lfloor t/2 \rfloor}{K-1} \frac{1}{K}. \quad (3.1)
\end{aligned}$$

In order for BRD to not be in a trap or a PNE by time $t+1$, it must: not be in a trap or PNE at time t; avoid all previously visited rows and columns; and not step to a PNE. It follows that

$$\mathbb{P}\left(\mathsf{BRD}(t+1) \notin \mathsf{E} \cup \mathcal{R}(t)\right) = \mathbb{P}\left(\mathsf{BRD}(t) \notin \mathsf{E} \cup \mathcal{R}(t-1)\right) \frac{K-1-\lfloor t/2 \rfloor}{K-1} \frac{K-1}{K}. \quad (3.2)$$

Applying Eqs. (3.1) and (3.2) repeatedly, we obtain

$$\begin{aligned}
&\mathbb{P}\left(\mathsf{BRD} \text{ converges to a PNE}\right) \\
&\quad = \mathbb{P}\left(\mathsf{BRD}(2K-2) \in \mathsf{E}\right) \\
&\quad = \mathbb{P}\left(\mathsf{BRD}(1) \in \mathsf{E}\right) \\
&\qquad + \mathbb{P}\left(\mathsf{BRD}(1) \notin \mathsf{E} \cup \mathcal{R}(0)\right) \left(\frac{1}{K-1} \sum_{t=1}^{2K-3} \prod_{j=1}^{t} \frac{K-1-\lfloor j/2 \rfloor}{K} \right) \\
&\quad = \frac{1}{K} + \frac{1}{K} \sum_{t=1}^{2K-3} \prod_{j=1}^{t} \frac{K-1-\lfloor j/2 \rfloor}{K}.
\end{aligned}$$

To find an upper bound for the above sum, we make use of the inequality $e^x \geq 1+x$, giving

$$\sum_{t=1}^{2K-3} \prod_{j=1}^{t} \frac{K-1-\lfloor j/2 \rfloor}{K} \leq \sum_{t=1}^{2K-3} \exp\left(-\frac{1}{K} \sum_{j=3}^{t+2} \lfloor j/2 \rfloor \right) \leq 2 \sum_{t=1}^{K-2} \exp\left(-\frac{t^2}{K} \right).$$

Finally, we can bound this quantity above by making use of a Riemann sum:

$$\sqrt{K} \sum_{t=1}^{K-2} \frac{1}{\sqrt{K}} \exp\left(-\frac{t^2}{K} \right) \leq \sqrt{K} \int_0^\infty e^{-x^2} dx = \frac{\sqrt{K\pi}}{2}.$$

This in turn yields

$$\mathbb{P}\left(\mathsf{BRD} \text{ converges to a PNE}\right) \leq \frac{1}{K} + \sqrt{\frac{\pi}{K}}.$$

□

The computations to find a lower bound are quite involved, and exponential inequalities like the one used above are only able to achieve loose lower bounds. Fortunately, Pangallo et al. [36] were able to prove a lower bound for this probability which is also asymptotically $\sqrt{\pi/K}$ using different methods.

3.2 Better Response Dynamics

While there are clear-cut conditions that indicate when BRD has entered a trap, that luxury unfortunately does not extend to the domain of bRD. To determine when a bRD has entered a trap based solely on its past, it must exhaust all movement options from every action profile in the suspected trap. This definition of a trap is unwieldy at best, so we require a different approach to address the question of convergence for bRD. Our main goal in this section is to prove that

$$\mathbb{P}(\mathsf{bRD} \text{ converges to a PNE}) = 1 - \mathrm{e}^{-1} + d_K \tag{3.3}$$

where the sequence d_K is non-negative and $d_K = O\left(1/K\right)$. Powers [38] proved that, in a strictly ordinal normal-form game with i.i.d. payoffs, when the number of actions of both players is sent to infinity, the number of PNE present in the game approaches a Poisson(1) distribution. This, in conjunction with our result, gives that

$$\mathbb{P}\left(\mathsf{bRD} \text{ converges to a PNE} \mid \mathsf{E}_K \neq \varnothing\right) = 1 - o(1).$$

3.2.1 Combinatorial bounds

For the vector $\boldsymbol{c} = (c_1, \ldots, c_K)$, we set $\ell(\boldsymbol{c}) = \sum_{s=1}^{K} c_s$. For intuitive purposes, consider $\boldsymbol{c}$ to be the vector describing the number of action profiles in each column belonging to a given trap; $\ell(\boldsymbol{c})$ is then the number of vertices in said trap.

Proposition 3.2. *Let $m \in [K]$. For fixed $\boldsymbol{c} \in [0, m]^K$,*

$$\prod_{i=1}^{K} c_i!(K - c_i)! \leq (m!)^{\lfloor \ell(\boldsymbol{c})/m \rfloor}(K - m)!^{\lfloor \ell(\boldsymbol{c})/m \rfloor} K!^{K - \lfloor \ell(\boldsymbol{c})/m \rfloor}.$$

Proof. Fix indices j, k such that $c_j \geq c_k \geq 1$ and define $\tilde{\boldsymbol{c}} \in [0, m]^K$ by $\tilde{c_j} = c_j + 1$, $\tilde{c_k} = c_k - 1$, and $\tilde{c_l} = c_l$ for all $l \notin \{j, k\}$. We have that

$$\prod_{i=1}^{K} \frac{\tilde{c_i}!(K - \tilde{c_i})!}{c_i!(K - c_i)!} = \frac{(c_j + 1)(K - (c_k - 1))}{c_k(K - c_j)} > 1.$$

Hence, moving weight from a column to one with equal or larger weight increases the desired value. If each column has an upper bound of m and a lower bound of 0, then this weight-shifting device leads to an optimized $\tilde{c}$ made up of 0's, $\lfloor \ell(\boldsymbol{c})/m \rfloor$ m's and possibly one column in $(0, m)$. Since $(K-c)!c! \leq K!$ for all $c \in [K]$, the "remainder" column can be bounded by $K!$. The result follows. □

Proposition 3.3. *Let $j \in [K]$. For fixed $\boldsymbol{c} \in [0, K]^K$ with $\ell(\boldsymbol{c}) < K$ and exactly j nonzero elements,*

$$\prod_{i=1}^{K} c_i!(K-c_i)! \leq (\ell(\boldsymbol{c}) - j + 1)!(K - \ell(\boldsymbol{c}) + j - 1)(K-1)!^{j-1} K!^{\ell(\boldsymbol{c})-j}.$$

Proof. As with Proposition 3.2, we know that the maximal value for this product is obtained when we can no longer shift weight from a smaller column to a larger column. If $\ell(\boldsymbol{c}) < K$ and j elements of $\boldsymbol{c}$ are nonzero, then this arrangement is obtained when $j-1$ elements of $\boldsymbol{c}$ are 1 and there is a unique element equal to $\ell(\boldsymbol{c}) - (j-1)$. □

Proposition 3.4. *Fix $m \in [K]$. For any $\boldsymbol{c} \in [0, m]^K$ with $\ell(\boldsymbol{c}) < K$,*

$$\prod_{i=1}^{K} \frac{\binom{m}{c_i}}{\binom{K}{c_i}} \leq \left(\frac{m}{K}\right)^{\ell(\boldsymbol{c})}.$$

Proof. Fix indices j, k such that $c_j \geq c_k \geq 1$, and define $\tilde{\boldsymbol{c}} \in [0, m]^K$ as in the proof of Proposition 3.2. We have that

$$\prod_{i=1}^{K} \frac{\binom{m}{\tilde{c}_i}\binom{K}{c_i}}{\binom{K}{\tilde{c}_i}\binom{m}{c_i}} = \frac{(K-\tilde{c}_j)!(K-\tilde{c}_k)!(m-c_j)!(m-c_k)!}{(m-\tilde{c}_j)!(m-\tilde{c}_k)!(K-c_j)!(K-c_k)!} = \frac{(m-c_j)(K-c_k+1)}{(m-c_k+1)(K-c_j)}.$$

As the function $x/(x+n)$ increases in x for any $n > 0$, the ratio above is less than 1. Hence, moving weight from one column to another with equal or larger weight decreases the product in question, and the maximum value is obtained when all columns have equal weighting. As $\ell(\boldsymbol{c}) < K$, this means that $\ell(\boldsymbol{c})$ elements of $\boldsymbol{c}$ have a value of 1, and all other elements are 0. □

3.2.2 Convergence of better response dynamics

We return our focus to the main task at hand. For a bRD-trap τ, denote by $\mathsf{R}(\tau)$ (resp. $\mathsf{C}(\tau)$) the K-dimensional vector whose j-th entry is the number of action profiles in τ for which the first player (resp. second player) may choose action j. The length of a trap is the number of action profiles that it contains. We denote by T_n the collection of bRD-traps of length n; further, let $\mathsf{T} = \bigcup_{n=4}^{K^2} \mathsf{T}_n$. For action $i \in [K]$, let

$$M_{i,u} := \{k\colon Z_2^{i,k} \text{ is among the largest } u \text{ payoffs in row } i\},$$

and define $\Delta_1^\sigma := \bigcup_{i,j} \bigcap_{k \in M_{i,K^\sigma}} \{Z_2^{i,k} > Z_2^{j,k}\}$. In words, Δ_1^σ is the event in which there exists two rows, i and j, such that the K^σ action profiles in i with the largest payoffs are all profitable deviations from their neighbouring profiles in j.

Lemma 3.5. *Fix a parameter $\sigma \in (0,1)$. We have $\mathbb{P}(\Delta_1^\sigma) \le K^2 \left(\frac{1}{2}\right)^{K^\sigma}$.*

Proof. A union bound allows us to compare two fixed distinct rows, say i and j, which gives rise to the K^2 factor. For the remaining part of the upper bound, notice that for fixed $k \in [K]$, the events $\{Z_2^{i,k} > Z_2^{j,k}\}$, for $k \in M_{i,K^\sigma}$, are independent and share the same probability of 1/2. □

Instead of proving Eq. (3.3), we will aim to instead prove a stronger result, that traps only exist for finitely many K. Clearly, the nonexistence of traps implies that bRD will converge. Set $J_K := \{\mathsf{E}_K \neq \varnothing\}$, i.e., the event in which no PNE exists, and let $\overline{J_K}$ be its complement. We have

$$\mathbb{P}(\mathsf{bRD} \text{ does not converge to a PNE}) \le \mathbb{P}(\{\mathsf{T} \neq \varnothing\} \cap J_K) + \mathbb{P}\big(\,\overline{J_K}\,\big).$$

It thus suffices to prove the following theorem.

Theorem 3.6. *For any $\alpha \in (0,1)$, we have that there exists a positive constant, denoted by cnst, such that for all large enough K,*

$$\mathbb{P}(\{\mathsf{T} \neq \varnothing\} \cap J_K) \le \mathit{cnst}\, K^{-3+2\alpha} + K^2 \left(\frac{1}{2}\right)^{K^{\alpha/2}} + K^{-K^\alpha}.$$

Proof. In what follows, *cnst* denotes a generic constant that may change from line to line. We have that

$$\mathbb{P}\left(\left\{\bigcup_{n=4}^{K^2} \mathsf{T}_n \neq \varnothing\right\} \cap J_K\right) \le \mathbb{P}\left(\bigcup_{n=4}^{\lfloor K^\alpha \rfloor} \mathsf{T}_n \neq \varnothing\right) + \mathbb{P}\left(\left\{\bigcup_{n=\lfloor K^\alpha \rfloor+1}^{K^2} \mathsf{T}_n \neq \varnothing\right\} \cap J_K\right). \quad (3.4)$$

For a given subset of action profiles V, let $\mathsf{F}(\mathsf{V})$ be the projection of V on the second coordinate (i.e., $i \in \mathsf{F}(\mathsf{V})$ if and only if V contains action profiles from the i^{th} row). Observe that for V to be a trap, in every row and column intersecting with V, there must not exist action profiles that are better responses than those in V. If there are v vertices of V in a given row or column, the probability that these v vertices attain the highest v payoffs is $\binom{K}{v}$. Moreover, this event is independent from the payoffs in every other row and column. Hence, if V has length n, then

$$\mathbb{P}(\mathsf{V} \in \mathsf{T}_n) \leq \left(\prod_{i=1}^{K} \binom{K}{\mathsf{R}_i(\mathsf{V})} \binom{K}{\mathsf{C}_i(\mathsf{V})} \right)^{-1}.$$

The inequality arises from the fact that certain orderings may not produce a trap, e.g., a row and column may have their best responses on the same action profile, resulting in a PNE.

Recall that $\ell(\boldsymbol{c})$ represents the number of vertices in a vertex set represented by the vector $\boldsymbol{c}$. To bound the first term in Eq. (3.4), we apply a union bound:

$$\begin{aligned}
\mathbb{P}\left(\bigcup_{n=4}^{\lfloor K^\alpha \rfloor} \mathsf{T}_n \neq \varnothing \right) &\leq \sum_{n=4}^{\lfloor K^\alpha \rfloor} \sum_{j=2}^{n-2} \mathbb{P}\big(\exists \mathsf{V} \subset [K]^2 : \mathsf{V} \in \mathsf{T}_n, |\mathsf{F}(\mathsf{V})| = j \big) \\
&\leq \sum_{n=4}^{\lfloor K^\alpha \rfloor} \sum_{j=2}^{n-2} \binom{K}{j} \sum_{\substack{\boldsymbol{r} \in [0,K]^K: \\ \ell(\boldsymbol{r}) = n, \max(r_i) \leq j}} \sum_{\substack{\mathsf{V} \subset [j]^K: \\ \mathsf{R}(\mathsf{V}) = \boldsymbol{r}, \mathsf{F}(\mathsf{V}) = [j]}} \frac{1}{\prod_{i=1}^{K} \binom{K}{\mathsf{R}_i(\mathsf{V})} \binom{K}{\mathsf{C}_i(\mathsf{V})}} \\
&\leq \sum_{n=4}^{\lfloor K^\alpha \rfloor} \sum_{j=2}^{n-2} \binom{K}{j} \sum_{\substack{\boldsymbol{r} \in [0,K]^K: \\ \ell(\boldsymbol{r}) = n, \max(r_i) \leq j}} \frac{1}{\prod_{i=1}^{K} \binom{K}{r_i}} \sum_{\substack{\mathsf{V} \subset [j]^K: \\ \mathsf{R}(\mathsf{V}) = \boldsymbol{r}, \mathsf{F}(\mathsf{V}) = [j]}} \frac{1}{\prod_{i=1}^{j} \binom{K}{\mathsf{C}_i(\mathsf{V})}}.
\end{aligned}$$

We use Proposition 3.3 to obtain the upper bound

$$\prod_{i=1}^{j} (K - \mathsf{C}_i(\mathsf{V}))! \mathsf{C}_i(\mathsf{V})! \leq (\ell(\mathsf{C}(\mathsf{V})) - j + 1)!(K - \ell(\mathsf{C}(\mathsf{V})) + j - 1)!(K-1)!^{j-1}.$$

Moreover, for a given vector $\boldsymbol{r}$ with each element being at most j, there are $\prod_{i=1}^{K} \binom{j}{r_i}$ sets of action profiles V satisfying $\mathsf{R}(\mathsf{V}) = \boldsymbol{r}$. This gives

$$\mathbb{P}\left(\bigcup_{n=1}^{\lfloor K^\alpha \rfloor} \mathsf{T}_n \neq \varnothing \right) \leq \sum_{n=4}^{\lfloor K^\alpha \rfloor} \sum_{j=2}^{n-2} \frac{(K-n+j-1)!(n-j+1)!}{K^{j-1} j!(K-j)!} \sum_{\substack{\boldsymbol{r} \in [0,K]^K: \\ \ell(\boldsymbol{r}) = n, \max(r_i) \leq j}} \prod_{i=1}^{K} \frac{\binom{j}{r_i}}{\binom{K}{r_i}}.$$

Proposition 3.4 allows us to bound the product above, and a loose bound for the number of vectors $\boldsymbol{r}$ with $\ell(\boldsymbol{r}) = n$ is $\binom{n+K-1}{K}$. Hence,

$$\mathbb{P}\left(\bigcup_{n=1}^{\lfloor K^\alpha \rfloor} \mathsf{T}_n \neq \varnothing \right) \leq \sum_{n=4}^{\lfloor K^\alpha \rfloor} \sum_{j=2}^{n-2} \frac{(K-n+j-1)!(n-j+1)!}{K^{j-1} j!(K-j)!} \binom{n+K-1}{K} \left(\frac{j}{K} \right)^n. \quad (3.5)$$

We will use Stirling's approximation to upper bound this sum. To this end, let $\phi(x) = x^x = e^{\ln(x)}$; after using the approximation the term inside the sum becomes

$$\sqrt{\frac{(K-n+j-1)(n-j+1)(n+K-1)}{2\pi j(K-j)K(n-1)}} \frac{\phi(K-n+j-1)\phi(n-j+1)\phi(n+K-1)}{K^{K+n+j-1} j^{j-n} \phi(K-j)\phi(n-1)}.$$

Now, let

$$G(K,n,j) := \frac{\phi(K-n+j-1)\phi(n-j+1)\phi(n+K-1)}{K^{K+n+j-1}j^{j-n}\phi(K-j)\phi(n-1)}.$$

Differentiating with respect to n yields

$$\frac{\partial G}{\partial n}(K,n,j) = G(K,n,j)\ln\left(\frac{(n-j+1)(n+K-1)j}{(K-n+j-1)(n-1)K}\right),$$

and as $j+2 \le n \le \lfloor K^\alpha \rfloor$, it is easily seen that, for large enough K, G is decreasing in n, so

$$G(K,n,j) \le G(K,j+2,j) = \frac{27\phi(K-3)\phi(j+k+1)j^3}{K^{K+2j+1}\phi(j)\phi(K-j)\phi(j+1)}.$$

With a minor rearrangement, and using the bound $(1+a/x)^{a+x} \le \exp(a+a^2/x)$, we see that

$$G(K,j+2,j) \le \mathit{cnst}\,\frac{\phi(K-3)j^2}{K^j\phi(j+1)\phi(K-j)} \le \mathit{cnst}\,\frac{K^{K-j-3}}{(j+1)^{j-1}\phi(K-j)}.$$

Let $H(K,j) = \ln(G(K,j+2,j))$; differentiating with respect to j gives

$$\frac{\partial H}{\partial j}(K,j) = \frac{2}{j+1} - \ln\left(j+1+\frac{j(j+1)}{K-j}\right) \le \frac{2}{j+1} - \ln(j+1)$$

which is negative for $j \ge 2$. Hence, H is decreasing and is thus maximised at $j = 2$; it immediately follows that $G(K,j+2,j) \le \mathit{cnst}\,K^{-3}$. Turning our attention to the square root term that arises from Stirling's approximation, we can see that

$$\begin{aligned}
&\frac{(K-n+j-1)(n-j+1)(n+K-1)}{j(K-j)K(n-1)} \\
&\qquad = \frac{(K-n+j-1)(n-j+1)}{j(K-j)K} + \frac{(K-n+j-1)(n-j+1)}{j(K-j)(n-1)} \\
&\qquad \le \frac{(n-j+1)}{j(K-j)} + \frac{(K-n+j-1)}{j(K-j)} \\
&\qquad = \frac{K}{j(K-j)} \le 1,
\end{aligned}$$

so the entire summand is upper bounded by $\mathit{cnst}\,K^{-3}$. Hence, for large enough K,

$$\mathbb{P}\left(\bigcup_{n=1}^{\lfloor K^\alpha \rfloor} \mathsf{T}_n \neq \varnothing\right) \le \sum_{n=4}^{\lfloor K^\alpha \rfloor}\sum_{j=2}^{n-2}\frac{\mathit{cnst}}{K^3} \le \mathit{cnst}\,K^{-3+2\alpha}. \tag{3.6}$$

To bound the remaining term in Eq. (3.4), we make use of the fact that for a trap τ to be larger than K^α, either: there must exist a row r or a column c containing at least $K^{\alpha/2}$ action profiles of τ; or no such row/column exists and the trap spans more than $K^{\alpha/2}$ rows and columns. These events will be denoted as A_1 and A_2 respectively. Note that any action profile neighbouring a PNE cannot be part of a trap; hence, if $\boldsymbol{s}$ and $\boldsymbol{t}$ are such that $\boldsymbol{s} \sim_i \boldsymbol{t}$, $\boldsymbol{s}$ neighbours a PNE, and $\boldsymbol{t}$ is in a trap, then $Z_i^{\boldsymbol{s}} < Z_i^{\boldsymbol{t}}$. Under event A_1, we can find at least $K^{\alpha/2}$ such pairs of action profiles: $\boldsymbol{t}$ chosen from row r and $\boldsymbol{s}$ chosen from any row containing a PNE. Hence, we are in the framework of Lemma 3.5, and by a union bound on rows and columns we have

$$\mathbb{P}\left(\left\{\bigcup_{n=\lfloor K^\alpha\rfloor+1}^{K^2} \mathsf{T}_n \neq \varnothing\right\} \cap J_K \cap A_1\right) \leq 2K^2\left(\frac{1}{2}\right)^{K^{\alpha/2}}. \tag{3.7}$$

We turn our attention to event A_2. By taking a union bound over all trap sizes larger than K^α, and applying Proposition 3.2, we get

$$\begin{aligned}\mathbb{P}\left(\left\{\bigcup_{n=\lfloor K^\alpha\rfloor+1}^{K^2} \mathsf{T}_n \neq \varnothing\right\} \cap J_K \cap A_2\right) &\leq \sum_{n=\lfloor K^\alpha\rfloor}^{K^2} \sum_{\mathsf{V}:|\mathsf{V}|=n} \frac{1}{(K!)^K \prod_{i=1}^K \binom{K}{\mathsf{R}_i(\mathsf{V})}} \prod_{i=1}^K c_i!(K-c_i)! \\ &\leq \sum_{n=\lfloor K^\alpha\rfloor}^{K^2} \sum_{\mathsf{V}:|\mathsf{V}|=n} \frac{(\lfloor K^{\frac{\alpha}{2}}\rfloor!(K-\lfloor K^{\frac{\alpha}{2}}\rfloor)!)^{\lfloor n/\lfloor K^{\alpha/2}\rfloor\rfloor} K!^{K-\lfloor n/\lfloor K^{\alpha/2}\rfloor\rfloor}}{(K!)^K \prod_{i=1}^K \binom{K}{\mathsf{R}_i(\mathsf{V})}}. \end{aligned}\tag{3.8}$$

For any $\boldsymbol{r} \in [0,K]^K$ with $\ell(\boldsymbol{r}) = n$, there are $\prod_{i=1}^K \binom{K}{r_i}$ action profile sets V of size n which satisfy $\mathsf{R}(\mathsf{V}) = \boldsymbol{r}$; moreover, there are $\binom{n+K-1}{K}$ such vectors. Hence the right hand side of Eq. (3.8) becomes

$$\sum_{n=\lfloor K^\alpha\rfloor}^{K^2} \binom{n+K-1}{K}\binom{K}{\lfloor K^{\alpha/2}\rfloor}^{-\lfloor n/\lfloor K^{\alpha/2}\rfloor\rfloor}. \tag{3.9}$$

With some careful manipulation of the many floor functions present and by making use of the bound

$$\left(\frac{n}{k}\right)^k \leq \binom{n}{k},$$

we can see that

$$\binom{n+K-1}{K}\binom{K}{\lfloor K^{\alpha/2}\rfloor}^{-\lfloor n/\lfloor K^{\alpha/2}\rfloor\rfloor} \leq \binom{n+K-1}{K} K^{(1-\alpha/2)(K^{\alpha/2}-1)(1-nK^{-\alpha/2})}.$$

Denote by $G(K,n)$ the rightmost side of the above equation; for large enough values of K, this function is decreasing in n, since

$$\begin{aligned}\frac{G(K,n+1)}{G(K,n)} &= \left(1+\frac{K}{n}\right) K^{-(1-\alpha/2)(K^{\alpha/2}-1)K^{-\alpha/2}} \\ &\leq K^{K^{-\alpha/2}}\left(\frac{1}{K^{1-\alpha/2}}+\frac{1}{K^{\alpha/2}}\right) < 1.\end{aligned}$$

By Stirling's approximation and the above result,

$$\begin{aligned}
&\sum_{n=\lfloor K^\alpha \rfloor}^{K^2} \binom{n+K-1}{K} \binom{K}{\lfloor K^{\alpha/2} \rfloor}^{-\lfloor n/\lfloor K^{\alpha/2} \rfloor \rfloor} \\
&\qquad \leq \sum_{n=\lfloor K^\alpha \rfloor}^{K^2} G(K, K^\alpha - 1) \\
&\qquad \leq cnst\, K^2 \sqrt{\frac{K^\alpha + K - 2}{K(K^\alpha - 2)}} \frac{\phi(K^\alpha + K - 2)}{\phi(K)\phi(K^\alpha - 2)} K^{(1-\alpha/2)(K^{\alpha/2}-1)(1-K^{\alpha/2}-K^{-\alpha/2})} \\
&\qquad \leq cnst\, K^{2-\alpha/4} \frac{\phi(K^\alpha + K - 2)}{\phi(K)\phi(K^\alpha - 2)} K^{(1-\alpha/2)(K^{\alpha/2}-1)(1-K^{\alpha/2}-K^{-\alpha/2})} \\
&\qquad \leq K^{-\beta K^\alpha}
\end{aligned}$$

for any choice of $\beta < \alpha/2$. This result, in conjunction with Eqs. (3.4), (3.6) and (3.7), concludes the proof. □

Remark. Notice that the upper bound Theorem 3.6 gives a better rate when α gets closer to zero. A "best rate" in the upper bound is never achieved.

4

Multi-action Games with Ties

The preceding two chapters present a pleasant duality in the structures of games that they address while still offering interesting differences in the hurdles they encounter. However, they also leave some obvious holes in the content explored, and this thesis would not be complete without at least beginning to address them.

Before we explore these new games, we can reduce our workload by noticing the following: if a bRD-trap exists, then so too does a BRD-trap. This fact is very easy to see—the BRD-trap in question is simply a subset of maximal payoff strategies in the bRD-trap. Thus, BRD traces out a path contained entirely within the bRD-trap, and since no PNE exist in the bRD-trap, this path must too be a trap. See Fig. 4.1 for a depiction of this behaviour. Because most of our convergence results will instead be stated in terms of the existence (or nonexistence) of traps, we can now just show that, for example, BRD-traps do not exist and immediately get that bRD-traps also do not exist.

In this thesis, we only deal with normal-form representations of games. For any two strategies $\boldsymbol{s}$ and $\boldsymbol{t}$ with $\boldsymbol{s} \sim_i \boldsymbol{t}$, player i either prefers $\boldsymbol{s}$ to $\boldsymbol{t}$, prefers $\boldsymbol{t}$ to $\boldsymbol{s}$, or (in the case of a tie in payoffs) is indifferent to the two. This behaviour is invariant to any strictly

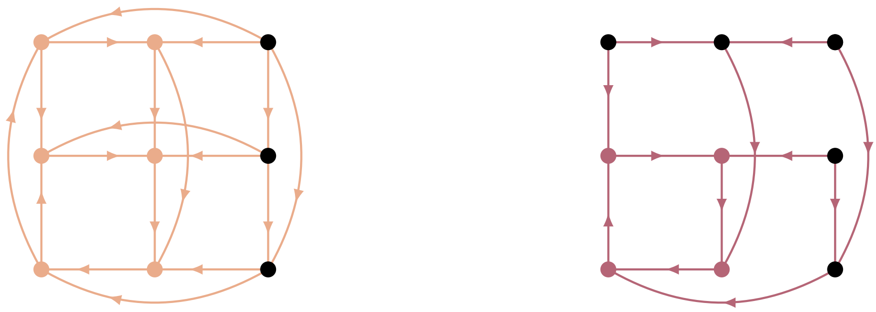

Figure 4.1: Graphical representations of bRD (left) and BRD (right) with the same payoffs. The bRD-trap strategies are depicted in peach, while BRD-trap strategies are depicted in maroon.

increasing transformations applied to the payoffs, but also implies that regardless of the underlying payoff distribution, whenever we are only comparing two payoffs at a time, the behaviour of the set of PNEs remains the same given the probability of ties between payoffs remains constant. In Chapter 2, this allowed us to only consider the probability of ties between payoff distributions, and disregard the underlying distribution governing these payoffs. Conversely, in Chapter 3, as we strictly concerned ourselves with continuous i.i.d. payoff distributions, we always had symmetry in our dependent payoffs allowing an even stronger abstraction than we use in Chapter 2.

Complications arise when we introduce atoms into the payoffs of a game with more than two actions and consider BRD on this game. A significant reason that calculations are simplified in the two-action model is that the direction of edges is calculated using only two payoffs. If we consider two games in which the payoffs are distributed following either a geometric distribution with mean $3/2$ or a Bernoulli distribution with parameter $1/2$, the probability of ties in both cases is $1/2$. When we change the model to consider K actions and we try to calculate the probability of the top two payoffs being tied, things change considerably. For the aforementioned Bernoulli distribution, to not have a tie we need exactly one player to have a payoff of 1, and every other player to have a payoff of 0, giving a probability of $1/2^K$, so immediately we see that the probability of a tie is a function of the number of actions. The geometric distribution on the other hand has probability of no ties given by the expression

$$\sum_{x=1}^{\infty} K \frac{2}{3^x} \left(1 - \frac{1}{3^x}\right)^{K-1}.$$

Interestingly, not only does this sum not have a closed form expression, but it also fails to converge [19].

This section will explicitly look at two-player, K-action games. Continuing with the notation from Chapter 3, we will denote by Z a generic independent copy of the payoff random variables. We will also introduce $F_Z(z) = \mathbb{P}(Z \leq z)$, as well as $f_Z(z) = F_Z'(z)$ wherever F_Z is differentiable. Finally, define $z_\infty := \inf\{z : F_Z(z) = 1\}$ and $f_\infty := \mathbb{P}(Z = z_\infty)$.

4.1 Continuity at the Top

We will first prove that the limiting distribution of the number of PNEs in this setting is Poisson with mean 1. In order to do this, we need the following lemma.

Lemma 4.1 (Arratia, Goldstein, and Gordon [5])**.** *For some index set I and each $n \in I$, let $X_n \sim$ **Bernoulli**(p_n) where $p_n = \mathbb{P}(X_n = 1)$, and define $X = \sum_{n \in I} X_n$ with $\lambda = \mathbb{E}[X]$. Furthermore, for each $n \in I$, define $B_n \subset I$ such that, for all $m \in I \backslash B_n$, X_n is independent of all X_m. Finally, set*

$$b_1 := \sum_{n \in I} \sum_{m \in B_n} p_m p_n \qquad \text{and} \qquad b_2 := \sum_{n \in I} \sum_{m \in B_n \backslash \{n\}} \mathbb{E}[X_m X_n].$$

Then, letting $\mathcal{L}(\cdot)$ denote the law of a random variable and recalling Definition 2.12,

$$\rho_{TV}(\mathcal{L}(X) - \mathcal{L}(\textbf{Poisson}(\lambda))) \leq \frac{2(b_1 + b_2)(1 - e^{-\lambda})}{\lambda}.$$

Remark. The requirement that all random variables outside of B_n need to be independent of X_n is not necessary, and Arratia, Goldstein, and Gordon [5] relax this requirement in the original result statement. We do not need this generalisation however, and it has been omitted for brevity.

Theorem 4.2. *Let $\varepsilon > 0$ such that the distribution of Z is continuous for all $z \in (z_\infty - \varepsilon, z_\infty]$. Then the limiting distribution of the number of PNEs is* **Poisson**(1)*.*

Proof. Let q be such that $1 - q = F_Z(z_\infty - \varepsilon)$, define

$$Q_1 = \{\forall t_1 \in [K], \exists t_2 \in [K] : Z_1^{t_1 t_2} \geq z_\infty - \varepsilon\},$$
$$Q_2 = \{\forall t_2 \in [K], \exists t_1 \in [K] : Z_2^{t_1 t_2} \geq z_\infty - \varepsilon\},$$

and let $Q = Q_1 \cap Q_2$. Note that when conditioning on Q, by symmetry every action profile has probability $1/K^2$ of being a PNE as we cannot observe ties. Moreover, neighbouring vertices cannot both be PNEs. Hence,

$$\begin{aligned}\mathbb{P}(\boldsymbol{s} \in \mathsf{E}) &\leq \mathbb{P}(\boldsymbol{s} \in \mathsf{E} \mid Q) + \mathbb{P}(Q^c) \leq \frac{1}{K^2} + 2K(1-q)^K, \\ \mathbb{P}(\boldsymbol{s} \in \mathsf{E}, \boldsymbol{t} \in \mathsf{E}) &\leq \mathbb{P}(\boldsymbol{s} \in \mathsf{E}, \boldsymbol{t} \in \mathsf{E} \mid Q) + \mathbb{P}(Q^c) \leq 2K(1-q)^K.\end{aligned} \tag{4.1}$$

We can also see that $\mathbb{P}(\boldsymbol{s} \in \mathsf{E}) \geq 1/K^2$ by symmetry: if the maximal payoff in $\mathsf{N}_1(\boldsymbol{s})$ is unique, $\boldsymbol{s}$ achieves that maximal payoff with probability $1/K$; otherwise if a tie exists, this probability increases. As this is the same when we consider $\mathsf{N}_2(\boldsymbol{s})$ and the two events are independent, we obtain our lower bound.

We can now use Lemma 4.1 with index set S, $X_{\boldsymbol{s}} = \mathbb{1}_{\{\boldsymbol{s} \in \mathsf{E}\}}$ and $B_{\boldsymbol{s}} = \mathsf{N}(\boldsymbol{s}) \cup \{\boldsymbol{s}\}$. We can upper bound b_1 and b_2 as

$$\begin{aligned}b_1 &\leq K^2(2K-1)\left(\frac{1}{K^2} + 2K(1-q)^K\right)^2, \\ b_2 &\leq K^2(2K-2)2K(1-q)^K.\end{aligned}$$

Hence,

$$\rho_{TV}\big(\mathcal{L}(|\mathsf{E}|) - \mathcal{L}\left(\mathbf{Poisson}\big(K^2\mathbb{P}(\boldsymbol{s} \in \mathsf{E})\big)\right)\big) \leq \frac{2}{K} + 4K^3q^K.$$

Since $\mathbb{P}(\boldsymbol{s} \in \mathsf{E}) \geq 1/K^2$, we can use a more general version of an upper bound on the total variation distance between two Poisson random variables by Adell and Jodrá [1, Eq. 2.2] in conjunction with Eq. (4.1) and get

$$\rho_{TV}\big(\mathcal{L}\left(\mathbf{Poisson}\big(K^2\mathbb{P}(\boldsymbol{s} \in \mathsf{E})\big)\right) - \mathcal{L}(\mathbf{Poisson}(1))\big) \leq K^2\mathbb{P}(\boldsymbol{s} \in \mathsf{E}) - 1 \leq 2K^2q^K.$$

Finally, by the triangular inequality,

$$\begin{aligned}&\rho_{TV}(\mathcal{L}(|\mathsf{E}|) - \mathcal{L}(\mathbf{Poisson}(1))) \\ &\quad\leq \rho_{TV}\big(\mathcal{L}(|\mathsf{E}|) - \mathcal{L}\left(\mathbf{Poisson}\big(K^2\mathbb{P}(\boldsymbol{s} \in \mathsf{E})\big)\right)\big) \\ &\qquad + \rho_{TV}\big(\mathcal{L}\left(\mathbf{Poisson}\big(K^2\mathbb{P}(\boldsymbol{s} \in \mathsf{E})\big)\right) - \mathcal{L}(\mathbf{Poisson}(1))\big) \\ &\quad\leq \frac{2}{K} + \big(4K^3 + 2K^2\big)q^K.\end{aligned}$$

□

4.2 Atom at the Top

In this section we consider when the payoff distribution has a positive probability of achieving the maximum. Recall the definition of the Kolmogorov distance from Definition 2.13, and that Φ denotes the standard normal distribution function.

Theorem 4.3. *If $f_\infty > 0$ then the following hold.*

(a) Both BRD and bRD fail to converge to a PNE only finitely often.

(b) For all $K \in \mathbb{N}$ and for some constant C_{f_∞} dependent only on f_∞,

$$\kappa\left(\frac{|\mathsf{E}| - K^2 f_\infty^2}{\sqrt{K^2 f_\infty^2 (1 - f_\infty^2)}}, \Phi\right) \leq \frac{C_{f_\infty}}{K} + 2K(1 - f_\infty)^K.$$

Proof. We begin with part (a). For action profile $\boldsymbol{s} \in \mathsf{S}$ and for $i \in \{1, 2\}$, let $Q_i^\infty(\boldsymbol{s})$ be the number of action profiles $\boldsymbol{t} \in \mathsf{N}_i(\boldsymbol{s})$ with payoff vector (z_∞, z_∞). Clearly,

$$\mathbb{P}\left(Q_1^\infty(\boldsymbol{s}) + Q_2^\infty(\boldsymbol{s}) = 0\right) = \left(1 - f_\infty^2\right)^{2K-2}.$$

If an action profile has payoff (z_∞, z_∞), it must necessarily be a PNE. Moreover, if an action profile is adjacent to a PNE, it cannot be part of a trap. We can upper bound the probability that there exists an action profile that is not adjacent to a PNE by

$$\mathbb{P}\left(\bigcup_{\boldsymbol{s} \in \mathsf{S}} \{Q_1^\infty(\boldsymbol{s}) + Q_2^\infty(\boldsymbol{s}) = 0\}\right) \leq K^2\left(1 - f_\infty^2\right)^{2K-2}.$$

Hence by the first Borel-Cantelli Lemma, both bRD-traps and BRD-traps exist finitely often, so they can only fail to converge to a PNE finitely often.

Moving our attention to part (b), we will define Q^∞ to be the number of action profiles with payoff vector (z_∞, z_∞). Because action profiles achieve this payoff vector independently of each other, Q^∞ follows a **Binomial**(K^2, f_∞^2) distribution. Clearly,

$$\kappa(|\mathsf{E}|, Q^\infty) = \kappa\left(\frac{|\mathsf{E}| - K^2 f_\infty^2}{\sqrt{K^2 f_\infty^2 (1 - f_\infty^2)}}, \frac{Q^\infty - K^2 f_\infty^2}{\sqrt{K^2 f_\infty^2 (1 - f_\infty^2)}}\right),$$

and then by the triangular inequality,

$$\kappa\left(\frac{|\mathsf{E}| - K^2 f_\infty^2}{\sqrt{K^2 f_\infty^2 (1 - f_\infty^2)}}, \Phi\right) \leq \kappa(|\mathsf{E}|, Q^\infty) + \kappa\left(\frac{Q^\infty - K^2 f_\infty^2}{\sqrt{K^2 f_\infty^2 (1 - f_\infty^2)}}, \Phi\right). \tag{4.2}$$

Since the indicator function of the event in which an action profile has payoff vector (z_∞, z_∞) is a Bernoulli random variable with parameter f_∞^2, we can view $Q^\infty - K^2 f_\infty^2$ as the sum of K^2 zero-mean random variables. We can then use the Berry-Esseen Theorem to obtain

$$\kappa\left(\frac{Q^\infty - K^2 f_\infty^2}{\sqrt{K^2 f_\infty^2 (1 - f_\infty^2)}}, \Phi\right) \leq \frac{C_{f_\infty}}{K} \tag{4.3}$$

for some constant C_{f_∞}.

If an action profile has payoff vector (z_∞, z_∞) then it is a PNE, so

$$\mathbb{P}\left(|\mathsf{E}| \leq x\right) \leq \mathbb{P}\left(Q^\infty \leq x\right).$$

Let H be the event in which there exists a PNE in the game that does not have payoff vector (z_∞, z_∞). For such a PNE to exist, we also need a row or column to exist along which no action profile attains a payoff of z_∞. Hence,

$$\mathbb{P}\left(H\right) \leq 2K(1 - f_\infty)^K.$$

Finally, conditional on H^c, we know that $|\mathsf{E}| = Q^\infty$, so we can write

$$\begin{aligned}
\mathbb{P}\left(|\mathsf{E}| \leq x\right) &= \mathbb{P}\left(\{|\mathsf{E}| \leq x\} \cap H\right) + \mathbb{P}\left(\{|\mathsf{E}| \leq x\} \cap H^c\right) \\
&\geq \mathbb{P}\left(\{Q^\infty \leq x\} \cap H^c\right) \\
&\geq \mathbb{P}\left(Q^\infty \leq x\right) - \mathbb{P}\left(H\right) \\
&\geq \mathbb{P}\left(Q^\infty \leq x\right) - 2K(1 - f_\infty)^K.
\end{aligned}$$

Therefore,

$$\kappa(|\mathsf{E}|, Q^\infty) \leq 2K(1 - f_\infty)^K,$$

and in conjunction with Eqs. (4.2) and (4.3), we obtain our desired result. □

4.3 Infinite Sequence of Atoms

The final type of payoff distribution we will investigate is those with an infinite sequence of atoms at the top of the distribution (e.g., Poisson and geometric). We will assume our payoff distribution has support $\mathbb{N} \cup \{0\}$; let $f_n = \mathbb{P}\left(Z = n\right)$ and $F_n = \mathbb{P}\left(Z \leq n\right)$. Note that $f_\infty = 0$.

The difficulty in working with these distributions lies in the probability of the maximal sample from K observations not being unique, i.e., a tie exists at the top of our sample. Letting T_K be the event in which we observe a tie at the top of a sample of size K, we can easily see that

$$1 - \mathbb{P}\left(T_K\right) = \sum_{n=0}^{\infty} K f_{n+1} F_n^{K-1}. \tag{4.4}$$

Interestingly, the behaviour of this series varies wildly based on the rate at which the distribution decays to 0. To describe this behaviour, we need the following definition.

Definition 4.4. For a sequence $(a_n)_{n\in\mathbb{N}}$ to *converge in logarithmic mean* to some value l, we require

$$\lim_{m\to\infty}\frac{1}{\ln(m)}\sum_{n=1}^{m}\frac{a_n}{n}=l.$$

In this case, we write $a_n \to_{\log} l$.

In some sense, convergence in logarithmic mean to l is akin to converging "on average" to l. The following lemma is a collection of several results from Eisenberg, Stengle, and Strang [19] that details these differences.

Lemma 4.5 (Eisenberg, Stengle, and Strang [19])**.**

*a) If $f_{j+1}/f_j = q$ (i.e., $Z \sim$ **Geometric**$(1-q)$) then the following hold.*

(i) $\lim_{K\to\infty}\mathbb{P}(T_K)$ *does not exist, but*

$$1-\mathbb{P}(T_K)\to_{\log}\frac{1-q}{\ln(1/q)}.$$

(ii) Fix $x\in\mathbb{R}$. If $K(m)=\lfloor q^{-(m+x)}\rfloor$ then

$$\lim_{m\to\infty}1-\mathbb{P}\left(T_{K(m)}\right)=(1-q)\sum_{t=-\infty}^{\infty}q^{t-x}\exp\left(-q^{t-x}\right)=:L_q(x).$$

L_q is non-constant, infinitely differentiable, and 1-periodic in x. Moreover, there exist $x,y\in[0,1]$ such that

$$L_q(y)\le\frac{1-q}{\ln(1/q)}\le L_q(x).$$

b) If $f_{j+1}/f_j\to 1$ then $\lim_{K\to\infty}\mathbb{P}(T_K)=0$.

c) If $f_{j+1}/f_j\to 0$ then $\limsup_{K\to\infty}\mathbb{P}(T_K)=1$ *but* $\liminf_{K\to\infty}\mathbb{P}(T_K)=1-e^{-1}$.

These results suggest some interesting behaviours in our games. Firstly, in regards to geometric payoffs, we obtain the following result.

Corollary 4.6. *Consider a sequence of games $(\Gamma_K)_{K\ge 2}$ with payoffs distributed as a* **Geometric**$(1-q)$ *random variable.*

(i) $\lim_{K\to\infty}\mathbb{E}[|\mathsf{E}_K|]$ *does not exist, but*

$$\mathbb{E}[|\mathsf{E}_K|]\to_{\log}\left(\frac{1-q}{q\ln(1/q)}\right)^2.$$

(ii) Fix $x \in \mathbb{R}$. If $K(m) = \lfloor q^{-(m+x)} \rfloor$ then

$$\lim_{m\to\infty} \mathbb{E}\left[\left|\mathsf{E}_{K(m)}\right|\right] = \left(\frac{L_q(x)}{q}\right)^2.$$

(iii) There exist $x, y \in [0,1]$ such that

$$\lim_{m\to\infty} \mathbb{E}\left[\left|\mathsf{E}_{\underline{K}(m)}\right|\right] \leq \left(\frac{1-q}{q\ln(1/q)}\right)^2 \leq \lim_{m\to\infty} \mathbb{E}\left[\left|\mathsf{E}_{\overline{K}(m)}\right|\right]$$

where $\overline{K}(m) = \lfloor q^{-(m+x)} \rfloor$ and $\underline{K}(m) = \lfloor q^{-(m+y)} \rfloor$.

Proof. Observe that

$$\begin{aligned}
\mathbb{E}\left[\left|\mathsf{E}_K\right|\right] &= K^2 \left(\sum_{n=0}^{\infty} f_n F_n^{K-1}\right)^2 \\
&= \left(\sum_{n=0}^{\infty} K(1-q)q^n \left(1-q^{n+1}\right)^{K-1}\right)^2 \\
&= \frac{1}{q^2}(1 - \mathbb{P}(T_K))^2.
\end{aligned}$$

The results immediately follow. □

It is possible to numerically solve for when $L_q(x)$ attains a maximum or minimum (see Fig. 4.2 for the behaviour of $L_q(x)$ over a single period). Hence, we can choose a subsequence of games to ensure a higher or lower expected number of PNE than if we were to simply observe the entire sequence as a whole. Fig. 4.2 also seems to suggest that if the maximum of L_q occurs at x, the minimum occurs at $x + 1/2$. It follows that if we simulated a game and obtained an average number of PNEs less than $(1-q)/\ln(1/q)$, we could in fact *reduce* the number of actions by a factor of $q^{-1/2}$ and the resulting game would have a higher expected number of PNEs while also being computationally much more efficient. Finally, the amplitude of the oscillations L_q exhibits drastically decreases as q increases. It follows that for particularly small q, this optimisation technique may result in much larger increases (or decreases) in expected PNEs while also significantly cutting down the required computation time.

See Fig. 4.3 for simulation results regarding average number of PNEs for various values of q. Note that $((1-q)/(q\ln(1/q))^2$ is roughly 16.98, 2.08 and 1.11 for q at 0.1, 0.5 and 0.9 respectively, which seems to match the midpoints in each of the respective line plots. For large values of q, we know from Eisenberg, Stengle, and Strang [19] that the magnitude of oscillations is remarkably small. But for q small, these oscillations become increasingly

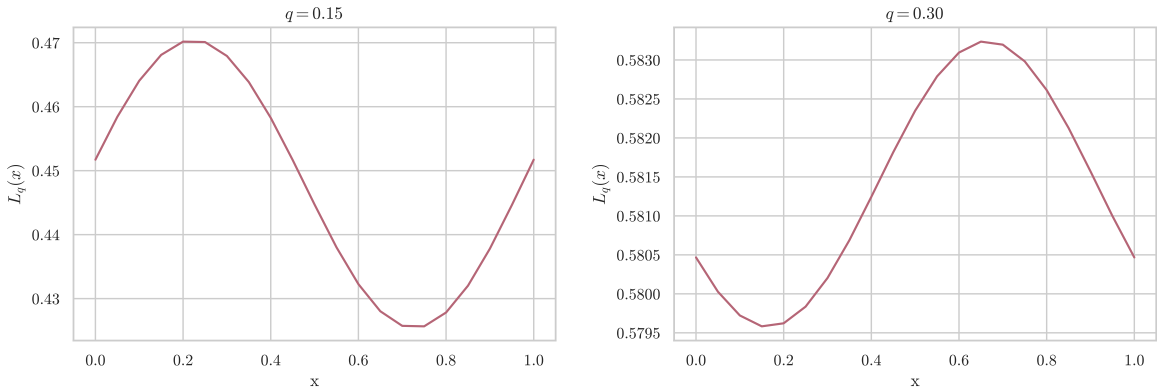

Figure 4.2: Behaviour of the function L_q over $[0, 1]$.

large. This behaviour can be seen in these plots; for the plots corresponding to $q = 0.5$ and $q = 0.9$, we see roughly constant expected value, whereas $q = 0.1$ displays the oscillatory behaviour Corollary 4.6 predicts. Moreover, the period of oscillation seems to be multiplicatively increasing, which aligns with how Corollary 4.6 predicts the minimal and maximal subsequences should be constructed.

A nice parallel to Lemma 4.5 is the following result.

Proposition 4.7. *Consider a game with K actions and payoff distribution satisfying $f_n/f_{n+1} \leq c$. Then $\mathbb{E}[|\mathsf{E}|] \leq c^2$.*

Proof. This proof is inspired by Eisenberg, Stengle, and Strang [19, Lemma 2]. Begin by writing

$$\mathbb{E}[|\mathsf{E}|] = \left(\sum_{n=0}^{\infty} K f_n F_n^{K-1} \right)^2 .$$

We turn our attention to the sum inside the brackets. By our assumption, $f_n \leq c f_{n+1}$. We can then substitute $f_{n+1} = F_{n+1} - F_n$ and get

$$\sum_{n=0}^{\infty} K f_n F_n^{K-1} \leq c \sum_{n=0}^{\infty} K f_{n+1} F_n^{K-1} = c \sum_{n=0}^{\infty} K(F_{n+1} - F_n) F_n^{K-1}.$$

This last sum is a lower Riemann sum for $\int_0^1 K x^{K-1} dx$, so

$$c \sum_{n=0}^{\infty} K(F_{n+1} - F_n) F_n^{K-1} \leq c$$

and the result immediately follows. □

Remark. From Lemma 4.5, distributions satisfying the assumption in Proposition 4.7 (distributions with tails at least as heavy as a geometric distribution) also have that $\limsup_{K\to\infty} \mathbb{P}(T_K)$ is bounded away from 1. In other words, there exists no subsequence of games such that the probability of ties converges to 1.

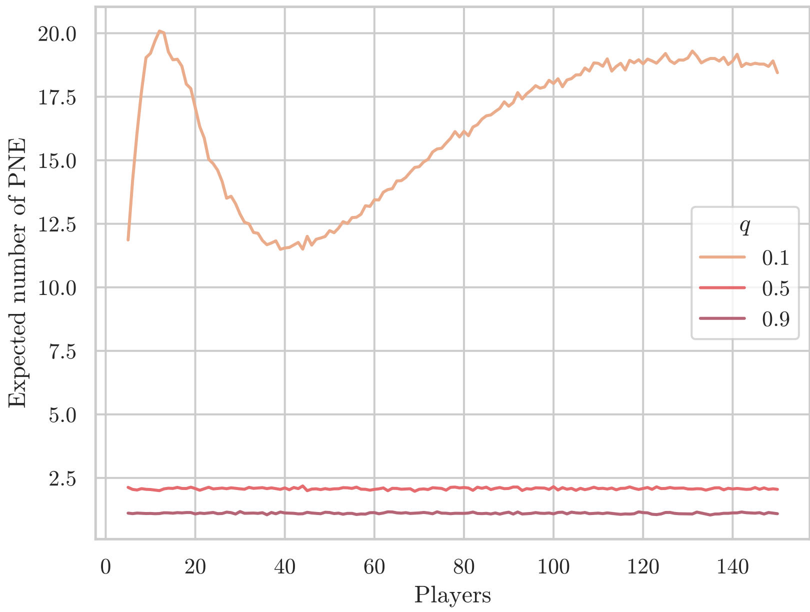

Figure 4.3: Average number of PNEs in 1500 simulations of games with geometrically distributed payoffs of parameter $1 - q$.

An example of a heavy-tailed discrete distribution is the distribution of the form

$$\mathbb{P}(Z_\gamma = z) = \frac{1}{z^\gamma} - \frac{1}{(z+1)^\gamma}$$

for $\gamma > 0$. We refer to this as the *discretised Pareto* distribution. Figs. 4.4a and 4.4b show simulation results for when payoffs follow a discretised Pareto distribution and a Poisson distribution with varying parameters, respectively. The number of PNEs when payoffs follow a discretised Pareto distribution are bounded as predicted by Proposition 4.7, while they appear unbounded when the payoffs follow a Poisson distribution, and look to grow logarithmically.

The discretised Pareto plot suggests a stronger result than just a bounded number of PNEs however. If BRD remains untrapped, it means it must visit a new row or column at every step. It stands to reason that if the probability of ties is sufficiently small, BRD may fail to encounter a tie and simply behave as if the distributions were continuous. Moreover, if ties exist finitely often for sufficiently heavy-tailed distributions, choosing a large enough K will result in games that behave identically to those with continuous payoffs.

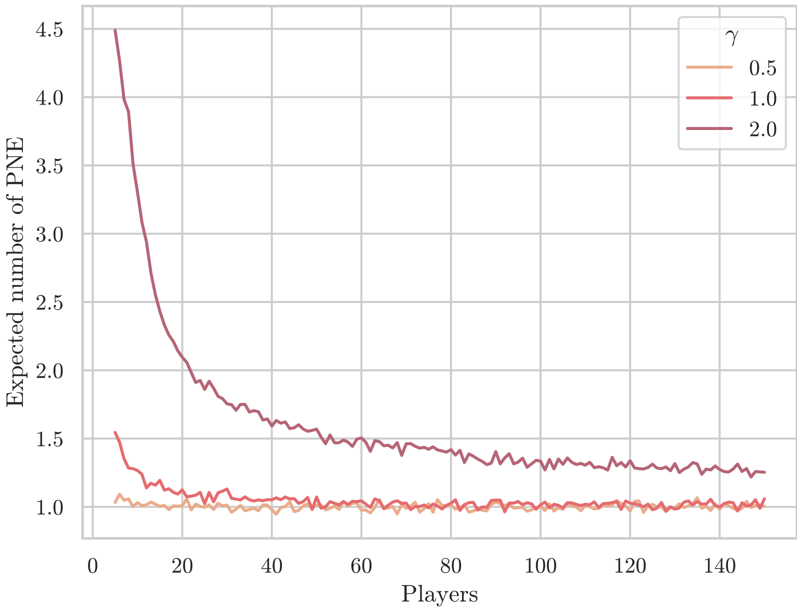

(a) Payoffs follow a discretised Pareto distribution with parameter γ.

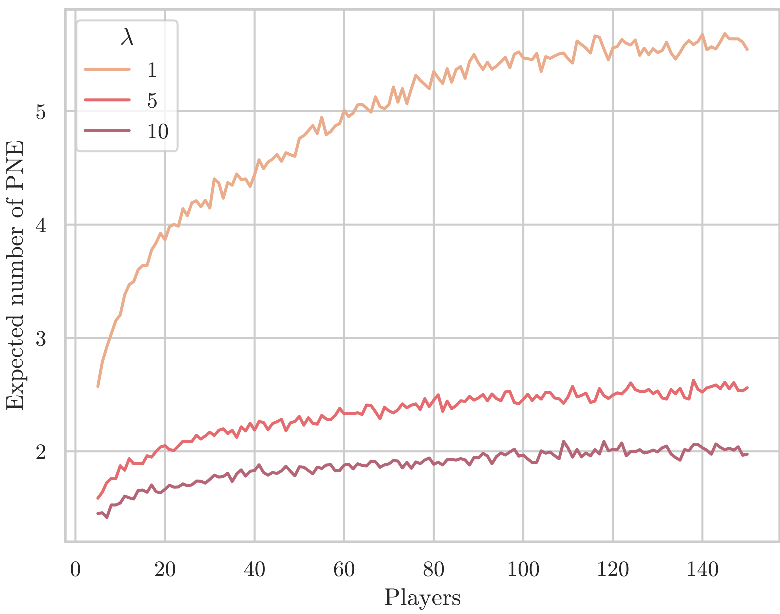

(b) Payoffs follow a Poisson distribution with parameter λ.

Figure 4.4: Expected number of PNE in 1500 simulations.

5

Conclusions

5.1 Summary

This thesis has explored the frequency with which PNEs exist in various types of random games. We have also investigated the difference between BRD and bRD, and identified situations in which each should be the preferred choice of iterative procedure to locate PNEs.

Chapter 2 focused on games in which N players could pick from two actions. A key extension in this model to the standard literature was allowing the existence of ties between payoffs. Graphically, this game could be represented as a hypercube where every edge was directed (or in the case of a tie, undirected) independently of every other edge. Owing to this, we were able to introduce a percolation model that would help answer questions about which action profiles were accessible by BRD. We were able to prove a central limit theorem result concerning the number of PNEs in the game, as well as the fact that the number of PNEs grows geometrically as the number of players increases. Furthermore, we showed that regarding the number of PNEs inaccessible by

BRD, there exists a phase transition in the probability of ties in payoffs: when ties occur with probability less than 1/2, PNEs are accessible by BRD in large games; when the probability of ties is 1/2, inaccessible PNEs occur with positive probability; and when ties occur with probability greater than 1/2, the number of inaccessible PNEs is unbounded. Finally, we showed that traps, the structures that prevent BRD from reaching a PNE, fail to exist in large games when the probability of ties is less than $2^{3/4} - 1$.

Chapter 3 focused on games in which two players could choose from K actions. Graphically, we were instead investigating games on a grid (in which each row and column is a complete graph of K vertices) rather than a hypercube. As the edges on this grid were no longer independently oriented, we had to discard the independent bond percolation approach. The first key result in this chapter was identifying that BRD fails to converge to a PNE as the number of actions grows, and was obtained via analysing the path of BRD in these games. We then showed that, in the presence of PNEs, structures that can trap bRD fail to exist in games with a large number of actions.

Finally, Chapter 4 extended the models we looked at in Chapter 3 to include a positive probability of ties occurring between payoffs. Initially, this was just a natural extension of the model in Chapter 3 when considered in context with the model from Chapter 2, but this simple tweak gave rise to a very interesting problem with unintuitive conclusions. The first section investigated payoff distributions that were continuous at the top of the support, and showed that the behaviour of the set of PNEs is the same as the case in which the entire distribution is continuous. The second section addressed payoff distributions with an achievable upper bound. In this context, both bRD and BRD converge to a PNE for large enough games. Moreover, we were able to show a central limit theorem result for the number of PNEs in these games. To conclude, we also investigated when the payoff distribution consisted of an infinite sequence of atoms and the support did not have a maximum, such as a geometric or Poisson distribution. Thanks to the contributions made by Eisenberg, Stengle, and Strang [19], under these assumptions we showed that the number of PNEs is bounded for distributions with sufficiently heavy tails. Moreover, specifically for geomterically distributed payoffs, we showed that the number of PNEs fails to converge, and that there exist subsequences of games that maximise and minimise the number of expected PNEs.

5.2 Current and Future Work

Describing the existence of PNEs and methods of finding them are both significant topics in the fields of game theory and economics. The 2023 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel was awarded to Ben Bernanke, Douglas Diamond and Philip Dybvig for improving our understanding of the role banks play in society, especially in the context of financial crises. Diamond and Dybvig specifically used PNEs to describe how the behaviour of bank customers varies between ordinary circumstances and in a situation in which a bank run is feared [13]. Being able to locate these PNEs in complex games with certainty, as well as knowing under what circumstances they exist, may uncover further meaning in the Diamond-Dybvig model and help us further understand financial instability in the future.

Currently, describing the existence of and locating PNEs is an active research topic. One such line of investigation by Mimun, Quattropani, and Scarsini [32] studies two-player, multi-action games when, for a given action profile, player 2 has probability p of attaining the same payoff as player 1. This model begins to bridge the gap between the games we have investigated in this thesis and random potential games. Another recent body of work by Pangallo et al. [36] extends into the realm of multi-player, multi-action games with no ties. However, instead of just looking at BRD as we have here, they analyse how different playing sequences impact convergence results. Specifically, they look at random playing sequences (our model), fixed cyclic order playing sequences, and even letting all players update their action at the same time.

5.2.1 Increasing the Number of Actions or Players

The most natural extensions from this thesis are investigating multi-player games with three or more actions, or multi-action games with three or more players with positive probability of ties. As soon as we introduce this generalisation to our games however, the bulk of the approaches we used in previous chapters become effectively useless:

- In the case of three-player, K-action games, we can no longer rely on independent bond percolation to obtain results, as the directed edges in the graph lose independence.
- In the case of three-action, N-player games, the BRD is no longer deterministic once the game is realised, so we cannot explicitly calculate the path probabilities.

5.2.2 Beyond Better and Best

So far we have only looked at two search algorithms to find PNEs in our games. Accordingly, BRD chooses out of only the best adjacent payoffs and moves to the corresponding action profile, while bRD picks from any payoff that yields an improvement and moves to that action profile. In Chapter 3, we saw that in two-player, K-action games with continuous payoffs, BRD fails to converge to a PNE whereas bRD will find them. An interesting line of inquisition is to investigate how we can relax the rules BRD follows (or tighten them for bRD) such that it successfully locates PNE in these games. For example, we could investigate a response process that could change action to achieve either the highest or second-highest payoff, and investigate the existence of traps in this setting.

An avenue that has remained unaddressed in this thesis is the question of speed of convergence. What is the expected number of steps required for either BRD, bRD or some other iterative process to reach a PNE? We can actually adapt the proof of Theorem 3.1 to show that BRD has an expected path length of $\sqrt{K}$ before it either reaches a PNE or enters a trap. Results on BRD convergence speed in the two-player game, as well as bRD convergence in any of our settings, would nicely complement the probability of convergence results detailed in this thesis.

5.2.3 Optimal Subsequences of Games

We have only scratched the surface of the games detailed in Chapter 4.3. The ramifications of subsequences of games achieving a higher expected number of PNEs have yet to be explored, and is an exciting avenue of research from a practical standpoint. Moreover, we predict there exist multiple phase transitions with respect to game behaviour in how quickly payoff distributions decay to 0. For sufficiently fast decay, we predict that games will exhibit the same behaviour that we see in finite payoff distributions. For sufficiently slow decay, we expect them to behave more like games with continuous payoffs. We also do not think that games with geometrically distributed payoffs are the only class of games that fail to converge in some sense to either of these extremes, as suggested by Fig. 4.4a.

Bibliography

[1] J. A. Adell and P. Jodrá. “Exact Kolmogorov and total variation distances between some familiar discrete distributions”. In: *Journal of Inequalities and Applications* (2006), Art. ID 64307, 8.

[2] B. Amiet, A. Collevecchio, and K. Hamza. “When “Better” is better than “Best””. In: *Operations Research Letters* 49(2) (2021), pp. 260–264.

[3] B. Amiet, A. Collevecchio, M. Scarsini, and Z. Zhong. “Pure Nash equilibria and best-response dynamics in random games”. In: *Mathematics of Operations Research* 46(4) (2021), pp. 1552–1572.

[4] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. “The price of stability for network design with fair cost allocation”. In: *SIAM Journal on Computing* 38(4) (2008), pp. 1602–1623.

[5] R. Arratia, L. Goldstein, and L. Gordon. “Two moments suffice for Poisson approximations: the Chen-Stein method”. In: *The Annals of Probability* 17(1) (1989), pp. 9–25.

[6] B. D. Bernheim. “Rationalizable strategic behavior”. In: *Econometrica. Journal of the Econometric Society* 52(4) (1984), pp. 1007–1028.

[7] L. E. Blume. “The statistical mechanics of strategic interaction”. In: *Games and Economic Behavior* 5(3) (1993), pp. 387–424.

[8] B. Bollobás. “Random Graphs”. In: *Modern Graph Theory.* Springer, 1998, pp. 215–252.

[9] S. R. Broadbent and J. M. Hammersley. “Percolation processes. I. Crystals and mazes”. In: *Proceedings of the Cambridge Philosophical Society* 53 (1957), pp. 629–641.

[10] J. E. Cohen. "Cooperation and self-interest: Pareto-inefficiency of Nash equilibria in finite random games". In: *Proceedings of the National Academy of Sciences of the United States of America* 95(17) (1998), pp. 9724–9731.

[11] P. Coucheney, S. Durand, B. Gaujal, and C. Touati. "General revision protocols in best response algorithms for potential games". In: *Netwok Games, Control and OPtimization (NetGCoop)*. Trento, Italy: IEEE Explore, Oct. 29, 2014.

[12] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. "The complexity of computing a Nash equilibrium". In: *SIAM Journal on Computing* 39(1) (2009), pp. 195–259.

[13] D. W. Diamond and P. H. Dybvig. "Bank runs, deposit insurance, and liquidity". In: *Journal of Political Economy* (1983), pp. 401–419.

[14] M. Dindoš and C. Mezzetti. "Better-reply dynamics and global convergence to Nash equilibrium in aggregative games". In: *Games and Economic Behavior* 54(2) (2006), pp. 261–292.

[15] M. Dresher. "Probability of a pure equilibrium point in n-person games". In: *Journal of Combinatorial Theory* 8(1) (1970), pp. 134–145.

[16] S. Durand, F. Garin, and B. Gaujal. "Efficiency of best response dynamics with high playing rates in potential games". In: *IFIP WG 7.3 Performance 2018 - 36th International Symposium on Computer Performance, Modeling, Measurements and Evaluation*. Toulouse, France, Dec. 2018, pp. 1–2.

[17] S. Durand and B. Gaujal. "Complexity and optimality of the best response algorithm in random potential games". In: *Algorithmic Game Theory*. Vol. 9928. Lecture Notes in Computer Science. Springer, Berlin, 2016, pp. 40–51.

[18] R. Durrett. *Probability—Theory and Examples*. 5th ed. Cambridge University Press, Cambridge, 2019, pp. xii+419.

[19] B. Eisenberg, G. Stengle, and G. Strang. "The asymptotic probability of a tie for first place". In: *The Annals of Applied Probability* 3(3) (1993), pp. 731–745.

[20] P. Erdős and J. Spencer. "Evolution of the n-cube". In: *Computers & Mathematics with Applications* 5(1) (1979), pp. 33–39.

[21] A. Fabrikant, A. D. Jaggard, and M. Schapira. "On the structure of weakly acyclic games". In: *Theory of Computing Systems* 53(1) (2013), pp. 107–122.

[22] J. W. Friedman and C. Mezzetti. “Learning in games by random sampling”. In: *Journal of Economic Theory* 98(1) (2001), pp. 55–84.

[23] M. Goemans, V. Mirrokni, and A. Vetta. “Sink equilibria and convergence”. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05).* Oct. 2005, pp. 142–151.

[24] K. Goldberg, A. J. Goldman, and M. Newman. “The probability of an equilibrium point”. In: *Journal of Research of the National Bureau of Standards* 72(2) (1968), pp. 93–101.

[25] A. J. Goldman. “The probability of a saddlepoint”. In: *The American Mathematical Monthly* 64(10) (1957), pp. 729–730.

[26] J. M. Hammersley and D. J. A. Welsh. “Percolation theory and its ramifications”. In: *Contemporary Physics* 21(6) (1980), pp. 593–605.

[27] V. Y. Korolev and I. G. Shevtsova. “On the upper bound for the absolute constant in the Berry–Esseen inequality”. In: *Theory of Probability & Its Applications* 54(4) (2010), pp. 638–658.

[28] N. S. Kukushkin. “Better response dynamics and Nash equilibrium in discontinuous games”. In: *Journal of Mathematical Economics* 74 (2018), pp. 68–78.

[29] K. Kultti, H. Salonen, and H. Vartiainen. *Distribution of pure Nash equilibria in n-person games with random best replies.* Discussion Papers 71. Aboa Centre for Economics, Dec. 2011.

[30] J. Maynard Smith. “The theory of games and the evolution of animal conflicts”. In: *Journal of Theoretical Biology* 47(1) (1974), pp. 209–221.

[31] C. McDiarmid, A. Scott, and P. Withers. “The component structure of dense random subgraphs of the hypercube”. In: *Random Structures & Algorithms* 59(1) (2021), pp. 3–24.

[32] H. A. Mimun, M. Quattropani, and M. Scarsini. “Best-Response dynamics in two-person random games with correlated payoffs”. In: *arXiv preprint arXiv:2209.12967* (2022).

[33] V. S. Mirrokni and A. Skopalik. “On the complexity of Nash dynamics and sink equilibria”. In: *Proceedings of the 10th ACM conference on Electronic commerce.* 2009, pp. 1–10.

[34] D. Monderer and L. S. Shapley. "Potential games". In: *Games and Economic Behavior* 14(1) (1996), pp. 124–143.

[35] M. Pangallo, T. Heinrich, and J. D. Farmer. "Best reply structure and equilibrium convergence in generic games". In: *Science Advances* 5(2) (2019).

[36] M. Pangallo, T. Heinrich, Y. Jang, A. Scott, B. Tarbush, S. Wiese, and L. Mungo. *Best-response dynamics, playing sequences, and convergence to equilibrium in random games*. INET Oxford Working Papers 2021-02. Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Jan. 2021.

[37] T. Pei and S. Takahashi. "Rationalizable strategies in random games". In: *Games and Economic Behavior* 118 (2019), pp. 110–125.

[38] I. Y. Powers. "Limiting distributions of the number of pure strategy Nash equilibria in N-person games". In: *International Journal of Game Theory* 19(3) (1990), pp. 277–286.

[39] P. J. Reny. "On the existence of pure and mixed strategy Nash equilibria in discontinuous games". In: *Econometrica. Journal of the Econometric Society* 67(5) (1999), pp. 1029–1056.

[40] Y. Rinott and M. Scarsini. "On the number of pure strategy Nash equilibria in random games". In: *Games and Economic Behavior* 33(2) (2000), pp. 274–293.

[41] I. Shevtsova. "On the absolute constants in the Berry-Esseen type inequalities for identically distributed summands". In: *Informatika i Ee Primeneniya* 7 (1 2013), pp. 124–125.

[42] W. Stanford. "A note on the probability of k pure Nash equilibria in matrix games". In: *Games and Economic Behavior* 9(2) (1995), pp. 238–246.

[43] W. Stanford. "On the distribution of pure strategy equilibria in finite games with vector payoffs". In: *Mathematical Social Sciences* 33(2) (1997), pp. 115–127.

[44] W. Stanford. "On the number of pure strategy Nash equilibria in finite common payoffs games". In: *Economics Letters* 62(1) (1999), pp. 29–34.

[45] W. Stanford. "The limit distribution of pure strategy Nash equilibria in symmetric bimatrix games". In: *Mathematics of Operations Research* 21(3) (1996), pp. 726–733.

[46] S. Takahashi. “The number of pure Nash equilibria in a random game with nondecreasing best responses”. In: *Games and Economic Behavior* 63(1) (2008), pp. 328–340.

[47] S. Takahashi and T. Yamamori. “The pure Nash equilibrium property and the quasi-acyclic condition”. In: *Economics Bulletin* (2002).

[48] H. P. Young. “The evolution of conventions”. In: *Econometrica. Journal of the Econometric Society* 61(1) (1993), pp. 57–84.