AMSI VACATIONRESEARCH SCHOLARSHIPS 2019–20

EXPLORE THE MATHEMATICAL SCIENCES THIS SUMMER

DISCOVER RESEARCH THIS SUMMER

<The geometry of Nash Equilibria >

<Baihe Zhang>

Supervised by <Andrea Collevecchio>

<Monash University>

Vacation Research Scholarships are funded jointly by the Department of Education and Training and the Australian Mathematical Sciences Institute.

AMSI AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

Australian Government Department of Education

1 Introduction

Game theory is the study of strategic interaction between rational decision-makers. Nash equilibria plays an important role in game theory, and can be defined as the profile of actions where no player has an incentive to deviate from this strategy. In our research, we focus on games where a large number of players are involved and the payoffs are weakly dependent, with the possibility of ties. We study a multitude of phase transitions in terms of the correlation between payoffs and probability of ties. The asymptotic results about the random number of pure Nash equilibria will also be provided. The methods used in our analysis include large deviations, combinatorics, mechanical statistics, and computer simulation. Indeed,a link between percolation models on hypercube and game theory was already established by many authors. In particular, Amiet et al. (2019) studied the case of independent and identically distributed payoffs, with a chance of ties. They introduced a new approach for the analysis of Nash Equilibria, establishing a connection with percolation theory. The aim of the research is to describe the geometry of Nash Equilibria in detail when the assumption of identically distribution is tweaked and show evidence that the players are able to reach a Nash equilibrium with a probability that quickly approaches one as the number of players grows.

2 Definitions

Firstly, Some concepts related to Game theory will be introduced.

Game

A game can be described as a collection of players, a set of strategies available to each of them, and payoffs. Label the strategies using consecutive positive integer values, and identify with the set $[N] := 1, 2, ...N$ the collection of players. Let S_i be the set of actions of each player $i \in [N]$, and $S = \otimes_i S_i$, i.e. the cartesian product of the space of actions. We let $g_i : S \to R$ be the payoff function of player i. It is common in game theory to denote by s_{-i} the action of profile of all players except i. A game can be represented as

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

Pure Nash Equilibrium (PNE)

Pure Nash Equilibria are those strategies (if any) with the property that no single player has an incentive to deviate from a strategy. It is not necessarily an optimal outcome, but is a very natural concept.

More precisely, an action profile $\mathbf{s}^*$ is a pure Nash Equilibrium (PNE) of the game τ_N if for all $i \in [N]$ and for all $s_i \in S_i$, we have

$$g_i(\mathbf{s}^*) \geqslant g_i(s_i, s*_{-i})$$

3 Examples of Nash Equilibrium

3.1 A two-player example

The goal of this section is to illustrate the concept of PNE with the easiest possible example: two players and two strategies for each player.

Consider the case of two gangsters, that have been caught by the police. Each criminal are questioned, in different rooms. Each of them has two strategies, either confess or keep silent. The payoff of each criminal is shown in figure 1. For example, if both criminals confess, both of them will spend 8 years in a jail.

Criminal 2 / Criminal 1	Confess	Keep silent
Confess	$(-8,-8)$	$(0,-10)$
Keep silent	$(-10,0)$	$(-1,-1)$

Figure 1: the payoff for criminal 1 and criminal 2

It is easily to see that there is one Nash equilibrium, at the strategy [confess, confess],in this game.If either player one or two deviate from confess to keep silent, he will spend ten years in jail, which is a worse outcome. Therefore, when reaching this point, no player has an incentive to deviate. Similarly, by interpreting the other three strategies, it is concluded that only profile (confess, confess) is the pure Nash Equilibrium. Notice that in terms of payoffs, is not optimal, at least when compared with strategy profile (silent, silent).

3.2 Non-existence of PNE example

Unfortunately, even though the concept of PNE is natural, such equilibria not always exist in a game. The example provides in Figure 2 below has no PNE.

Player 2 / Player 1	0	1
0	$(1,7)$	$(4,1)$
1	$(4,1)$	$(1,7)$

Figure 2: A Game with no Nash Equilibrium

A fundamental idea of Nash was to introduce a different, more general concept of Equilibria, by looking at the so-called mixed strategy, where the set of strategies is enlarged, to include randomlised strategies. In the game appearing in Figure 2, a mixed strategy equilbria is achieved if both players flip an independent fair coin to decide which strategy to play. This concept has not a direct "natural" interpretation.

So a question raises naturally: under what conditions PNE exist? Few authors chose to answer this question by using randomness. If we pick a game uniformly at random, how many PNE are there?

4 A new model to represent : Best-Response Dynamics(BRD)

It is also natural to ask how to reach PNE, if they exist. It is well-known that PNE are hard to compute (Daskalakis et al., 2009). One way to address the issue is to devise iterative procedures that converge to a PNE. Suppose that at each step, a player is chosen at random and asked if she/he wants to change strategy. This defines a stochastic process. In the case each player chosen will consistently choose among the very best alternatives (vs. just better alternative), we speak about Best-Response Dynamics (Better-Response Dynamics respectively). In the case when each player has exactly two actions available, the two processes just described coincide. We use the acronym BRD to denote Best Response Dynamics. BRD helps the player to find PNE in a very natural way. PNE are absorbing states of BRD. To summarise, BRD a single player is picked at random and allowed to deviate from another more profitable strategy. If no such action exists, the player will not move and a different player is chosen randomly. When a new action profile is reached, the process is repeated. If we reach a profile which no player has a profitable deviation, then the process has reached a Pure Nash equilibrium.

4.1 Related work

Other than the recent paper by Amiet et al., several papers have considered aspects related to the number of PNE in games with random payoffs. In many of the papers that we consider below, and unless otherwise stated, the random payoffs are i.i.d. from a continuous distribution.

Goldman (1957) showed that for a zero-sum two-person games, the probability of having a PNE converges to zero as the number of strategies grows. The payoffs following a Bernoulli distribution was also briefly discussed in his research.

For a general two-person games,It was showed by Goldberg et al. (1968) that the probability of having at least one PNE converges to $1 - e^{-1}$ as the number of strategies diverges. Dresher (1970) generalized this result to the case of an arbitrary finite number of players.

The asymptotic distribution of the number of PNE can be found through recent papers. For example, Powers (1990) discovered that, when the number of strategies of at least two players goes to infinity, the distribution of the number of PNE converges to a Poisson(1).

Stanford (1995) derived an exact formula for the distribution of PNE in random games and obtained the result in Powers (1990) as a corollary. Stanford (1996) dealt with the case of two-person symmetric games and obtained Poisson convergence for the number of both symmetric and asymmetric PNE.

In all the above models, the expected number of PNE is found to be 1. Under different hypotheses, this expected number diverges. For instance, Stanford (1997, 1999) showed that this is the case for games with vector payoffs and for games of common interest, respectively. Stanford (1999) studied both strictly and weakly ordinal preferences. Raič (2003) used Chen-Stein method to bound the distance between the distribution of the normalized number of pPNE and a normal distribution, of which the result does not assume continuity of the payoff distributions.

Rinott and Scarsini (2000) weakened the hypothesis of i.i.d. payoffs; that is, they assumed that payoff vectors corresponding to different strategy profiles are i.i.d., but they allowed some dependence within the same payoff vector. In this setting, the asymptotic results were seen when either the number of players or the number of strategies diverges. More precisely, if each payoff vector has a multinormal exchangeable distribution with correlation coefficient ρ, then the following hold: for negative ρ, the probability of the number of PNE goes to zero; for $\rho = 0$ it converges to a Poisson(1), and; for ρ positive, it diverges and a central limit theorem can be used as a tool to do further analysis.

Takahashi (2008) considered the distribution of the number of PNE in a random game with two players, with a condition of the game having non-decreasing best-response functions. This assumption greatly increases the expected number of PNE. Daskalakis et al. (2011) extended the framework of games with random payoffs to graphical games. Strategy profiles are vertices of a graph and players' strategies are binary, like in our model. Moreover, their payoff depends only on their strategy and the strategies of their neighbors. The authors studied how the structure of the graph affects existence of

PNE and they examined both deterministic and random graphs.

A myriad of authors has explored the solutions of the random payoffs aspect in games from a range of angles.

For instance, Cohen (1998) studied the probability that Nash equilibria (both pure and mixed) maximize the sum of the players' payoffs in a finite random game. Pei and Takahashi (2019) devoted their attention to rational strategies in two-person games with random payoffs and performed an asymptotic analysis in the number of strategies.

Finding a PNE in a game is PPAD-complete (Daskalakis et al. (2009)). Therefore, given this computational difficulty, several learning procedures have been proposed to reach an equilibrium by playing the game several times (see, e.g., Tardos and Vazirani, 2007, Blum and Mansour, 2007). Probably the simplest such procedure is BRD. This approach has been taken, among others, by Blume (1993), Young (1993), Friedman and Mezzetti (2001), Takahashi and Yamamori (2002), Fabrikant et al. (2013). The main problem that arises is that BRD is guaranteed to converge to a PNE only when the game is of some specific type, for instance, a potential game (Monderer and Shapley (1996)). Bernheim (1984) introduced the concept of point rationalizable strategies, i.e., strategies that survive iterated elimination of never best responses against pure strategies. Pei and Takahashi (2019) pointed out an interesting connection between BRD and point rationalizability, showing that BRD converges to a PNE almost surely if and only if point rationalizable strategies consist only of PNE.

The performance of BRD in randomly drawn potential games has been studied in Coucheney et al. (2014), Durand and Gaujal (2016) and Durand et al. (2019). To be able to deal also with games for which BRD does not converge to a PNE, Goemans et al. (2005) defined the concept of *sink equilibria* A trap is a strongly connected set of two or more vertices with no edges leading out of the set. If players are selected at random and asked to choose a best response, the process may eventually reach a sink equilibrium and wander on the game's components permanently. Christodoulou et al. (2012) applied a similar model , with a focus on the rate of convergence to approximate solutions of the game. Dütting and Kesselheim (2017) analyzed BRD in the context of combinatorial auctions.

The idea of generating games at random to check properties of learning procedures was adopted by Galla and Farmer (2013) and, more recently, by Pangallo et al. (2019), who studied—mainly through simulations—the behavior of various learning procedures in games whose payoffs are drawn at random from a multinormal distribution.

5 Random games

In this section, the random game model will be considered. The model studied in Amiet et al (2019) considers independent and identically distributed payoffs, in the case of N players having 2 actions each. The probability of ties in the payoff, which we denote by α, is the important parameter in the model. Different values of α produce different possible behaviours in the number of PNE. Amiet, Collevecchio, Scarsini and Zhong proved that the number of PNE is roughly $(1+\alpha)^N$ and they also provide a very precise central limit theorem. They also give a description of the geometry of the PNE, and the question asked is "are they reachable by BRD?" They proved that BRD converges to PNE in the case for $\alpha < 0.68$ and conjectured that this is true for all $\alpha > 0$. The graphs below illustrate how concentrated the number of PNE is $(1+\alpha)^N$.

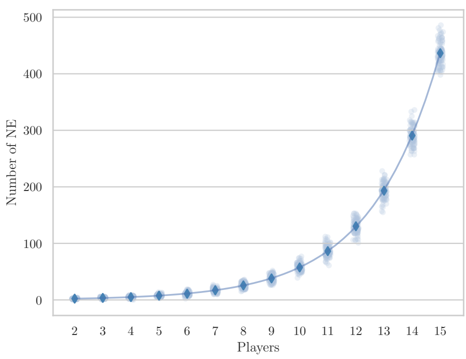

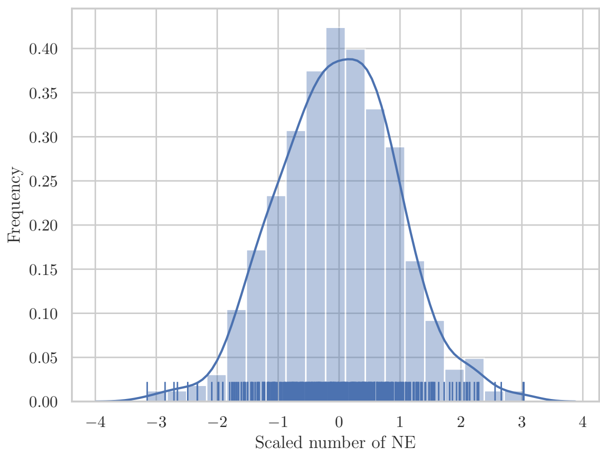

The method used in Amiet et al. rely on the connection between random games and random oriented graphs, which in turn are connected to percolation. We explain these connection using an example below.

Consider a game with three players, where the payoffs are described in the table below. This game can be described in terms orientation of the edges of the hypercube by repeating the following reasoning.

It is found that The two vertices in red $(0,0,0)$ and $(1,1,0)$ are the two pure Nash equilibria of the game.

Player 3 - Strategy 0		
Player 1 \ Player 2	0	1
0	$(0.52, 0.7, 0.42)$	$(0.2, 0.7, 0.56)$
1	$(0.22, 0.64, 0.12)$	$(0.85, 0.77, 0.58)$

Player 3 -Strategy 1		
Player 1 \ Player 2	0	1
0	$(0.2, 0.54, 0.17)$	$(0.19, 0.09, 0.39)$
1	$(0.11, 0.57, 0.79)$	$(0.95, 0.53, 0.05)$

Consider the vectors of strategies. As each of the coordinate can only take values in the set $\{0,1\}$, then we can identify this collection with the vertex set of the hypercube. For example consider the strategy $(0,0,1)$. The third player, does not like it. It is convenient for her/him to deviate from this strategy. We orient the edge connecting the vertices of the hypercube $(0,0,1)$ to $(0,0,0)$ in the direction of the latter, to indicate this preference. The presence of ties produces a partial orientation of the hypercube. In fact the edge connecting $(0,0,0)$ to $(0,1,0)$ is undirected. This representation points out a very simple fact, that the values of the payoffs are irrelevant, as long as we know they ranking. We need to know that $(0,0,0)$ is preferred to $(0,1,0)$ and we do not care if the increment is by 1 dollar or by 100 dollars. A random game, results in a random orientation. A random orientation is naturally connected with percolation, which can be described as follows. Independent bond percolation on hypercube is defined as follows. For each edge flip a coin having probability β of showing heads. If the toss shows heads, then declare the edge to be *open*; otherwise the edge is *closed*. The subgraph obtained by deleting (from the hypercube) the closed edges is a random graph , called percolation, that includes all vertices in the hypercube, but could be disconnected. Amiet et al. showed the largest connected component of the percolation, when $\beta = (1-\alpha)/2$ coincides (almost) with the set of strategies that can be reached by the BRD starting from a fixed non-random strategy. In the following picture we show what can be reached by BRD if starting from $(1,1,1)$ by continuing our examples with three players. The orange corresponds to the strategies that are accessible by Best-Response Dynamics.

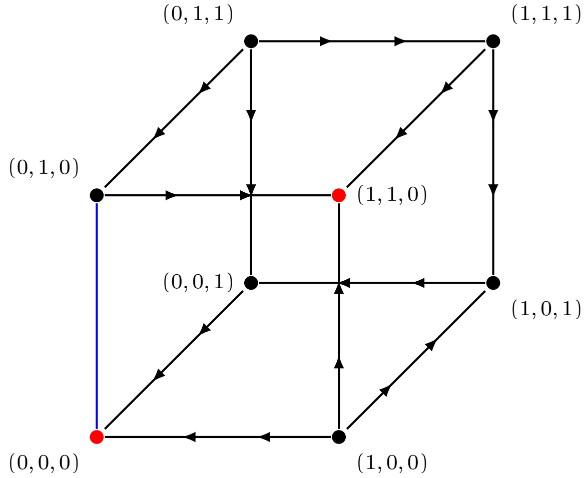

Figure 3: BRD representation

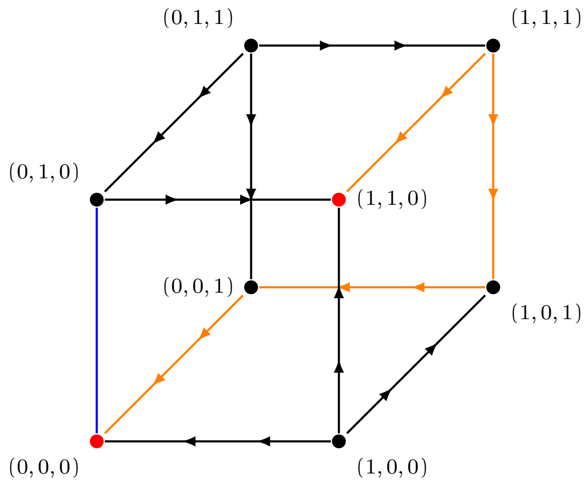

The connection with percolation that we mentioned above can be visualised as follows.

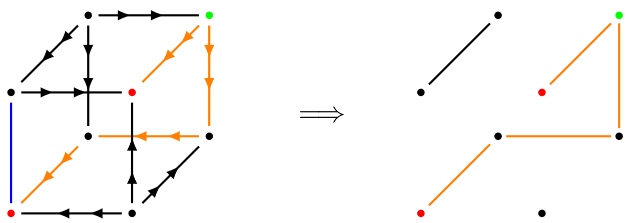

6 Further Study

I started studying games with payoffs that are not identically distributed but still independent. This is the first step in removing the i.i.d. assumption from Amiet et al. I obtained CLT in specific examples. In this case, N players and 2 actions are introduced, if two strategies have distance at least 2 then the events that they are PNE are independent. Hence, we could adapt results of CLT for 2-dependent processes to our context. We aim to obtain a general theory of heterogenous payoffs, and prove that in certain cases PNE localize in a certain region, making it easier for (variant of) BRD to find them.

7 Acknowledgement

I would like to thank the AMSI for providing this research opportunity and my supervisor, Dr Andrea Collevecchio, for his patient directions and meaningful help during the whole subject.

References

Amiet, B., Collevecchio, A., Scarsini, M., and Zhong, Z. (2019). Pure nash equilibria and best-response dynamics in random games. *arXiv preprint arXiv:1905.10758*.

Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.

Blum, A. and Mansour, Y. (2007). Learning, regret minimization, and equilibria. In *Algorithmic game theory*, pages 79–101. Cambridge Univ. Press, Cambridge.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games Econom. Behav.*, 5(3):387–424.

Christodoulou, G., Mirrokni, V. S., and Sidiropoulos, A. (2012). Convergence and approximation in potential games. *Theoret. Comput. Sci.*, 438:13–27.

Cohen, J. E. (1998). Cooperation and self-interest: Pareto-inefficiency of Nash equilibria in finite random games. *Proc. Natl. Acad. Sci. USA*, 95(17):9724–9731.

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

Daskalakis, C., Dimakis, A. G., and Mossel, E. (2011). Connectivity and equilibrium in random games. *Ann. Appl. Probab.*, 21(3):987–1016.

Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259.

Dresher, M. (1970). Probability of a pure equilibrium point in n-person games. *J. Combinatorial Theory*, 8:134–145.

Durand, S., Garin, F., and Gaujal, B. (2019). Distributed best response dynamics with high playing rates in potential games. *Performance Evaluation*, 129:40–59.

Durand, S. and Gaujal, B. (2016). Complexity and optimality of the best response algorithm in random potential games. In *Algorithmic Game Theory*, volume 9928 of *Lecture Notes in Comput. Sci.*, pages 40–51. Springer, Berlin.

Dütting, P. and Kesselheim, T. (2017). Best-response dynamics in combinatorial auctions with item bidding. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 521–533. SIAM, Philadelphia, PA.

Fabrikant, A., Jaggard, A. D., and Schapira, M. (2013). On the structure of weakly acyclic games. *Theory Comput. Syst.*, 53(1):107–122.

Friedman, J. W. and Mezzetti, C. (2001). Learning in games by random sampling. *J. Econom. Theory*, 98(1):55–84.

Galla, T. and Farmer, J. D. (2013). Complex dynamics in learning complicated games. *Proc. Natl. Acad. Sci. USA*, 110(4):1232–1236.

Goemans, M., Mirrokni, V., and Vetta, A. (2005). Sink equilibria and convergence. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–151.

Goldberg, K., Goldman, A. J., and Newman, M. (1968). The probability of an equilibrium point. *J. Res. Nat. Bur. Standards Sect. B*, 72B:93–101.

Goldman, A. J. (1957). The probability of a saddlepoint. *Amer. Math. Monthly*, 64:729–730.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games Econom. Behav.*, 14(1):124–143.

Pangallo, M., Heinrich, T., and Doyne Farmer, J. (2019). Best reply structure and equilibrium convergence in generic games. *Science Advances*, 5(2).

Pei, T. and Takahashi, S. (2019). Rationalizable strategies in random games. *Games Econom. Behav.*, 118:110–125.

Powers, I. Y. (1990). Limiting distributions of the number of pure strategy Nash equilibria in N-person games. *Internat. J. Game Theory*, 19(3):277–286.

Raič, M. (2003). Normal approximation by Stein's method. In Mrvar, A., editor, *Proceedings of the Seventh Young Statisticians Meeting*, pages 71–97.

Rinott, Y. and Scarsini, M. (2000). On the number of pure strategy Nash equilibria in random games. *Games Econom. Behav.*, 33(2):274–293.

Stanford, W. (1995). A note on the probability of k pure Nash equilibria in matrix games. *Games Econom. Behav.*, 9(2):238–246.

Stanford, W. (1996). The limit distribution of pure strategy Nash equilibria in symmetric bimatrix games. *Math. Oper. Res.*, 21(3):726–733.

Stanford, W. (1997). On the distribution of pure strategy equilibria in finite games with vector payoffs. *Math. Social Sci.*, 33(2):115–127.

Stanford, W. (1999). On the number of pure strategy Nash equilibria in finite common payoffs games. *Econom. Lett.*, 62(1):29–34.

Takahashi, S. (2008). The number of pure Nash equilibria in a random game with nondecreasing best responses. *Games Econom. Behav.*, 63(1):328–340.

Takahashi, S. and Yamamori, T. (2002). The pure Nash equilibrium property and the quasi-acyclic condition. *Econ. Bull.*

Tardos, E. and Vazirani, V. V. (2007). Basic solution concepts and computational issues. In *Algorithmic game theory*, pages 3–28. Cambridge Univ. Press, Cambridge.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.