On the number of pure Nash equilibria in random two-person games

Klaus Kultti
University of Helsinki

Hannu Salonen
University of Turku

Hannu Vartiainen
University of Turku

March 15, 2010

Abstract

We study pure Nash equilibria in random two player games, when best replies may be multi valued. We first show that, when all best reply correpondences are equally likely, the probability of at least one pure Nash equilibrium approaches one, and the expected number of pure Nash equilibria approaches infinity, when the size of the game becomes large. We then study the case where the utilities of the players are drawn from a finite set of utility indices. This set may depend on the size of the game. We derive an explicit formula for the limit distribution of pure Nash equilibria. The limit distribution is Poisson with mean that depends on the relative size of the set of utility indices against the choice set.

Keywords: random games, pure Nash equilibria
JEL: C71, C72.

1 Introduction

It is of importance to understand how Nash equilibrium behaves as a solution concept, on the "average". A vast literature has analyzed pure Nash equilibria (PNE) in games whose payoffs are drawn randomly from a maximum entropy distribution (see e.g. Stanford 1995a,b, 1996; Powers 1990, Goldberg et al. 1968, Dresher 1970). Of particular interest is the asymptotic distribution of pure Nash equilibria when the size of the game becomes large.[1]

A standard assumption in the literature has been that the payoffs are drawn randomly from a set, typically a continuum, that is much larger than the finite choice set. This assumption guarantees that players are never indifferent, and

[1]Our focus is restricted to pure strategies and finite action sets and independently drawn payoffs. McLennan (1997) allows mixed strategies and Bade et al. (2007) allow infinite action sets. Rinott and Scarsini (2000) study the case where there may be either positive or negative dependence among the players' payoffs.

that best responses are uniquely defined. This assumption simplifies the analysis remarkably.

However, we feel that possibility of multiple best responses should be accounted for. First, it can be argued that best reply correspondences are sufficient descriptions of games, at least as far as one is interested in solutions that depend only on best replies. If all best response correspondences are equally likely, then multiple best responses will materialize with positive probability (as is observed by looking at examples is standard game theory texts). Second, even if there are aspects in strategic interaction that are not captured by best reply correspondences alone, it is (arguably) not natural to think that - in a model where everything else is finite - there are infinitely many possible payoffs. This again leads to multiple best responses with positive probability.

We evaluate the likelihood pure Nash equilibria in random games where multiple best responses are allowed. Our focus is on random two-player matrix games where each player has K choices. We first show that, when all best reply correspondences are equally likely, the probability of at least one pure Nash equilibrium approaches one, and the expected number of pure Nash equilibria approaches infinity, when K becomes large.

The situation is more difficult, however, when randomness concerns the underlying utilities. To model this, we let the payoffs of the players be drawn independently from a finite set of utility indices. Letting the cardinality of the set of utilities, $f(K)$, depend on K in such a way that $f(K)/K$ approaches some real number r as K becomes large, the probability of indifferences does not vanish even in the limit.

The standard result when the utility indices are drawn from a continuum is that the distribution of pure Nash equilibria converges to the Poisson distribution with mean 1 as K becomes large. This, however, does not hold when utilities are drawn from a set of $f(K)$ indices. Our main finding is that the limit distribution of pure Nash equilibria converges to Poisson with mean

$$\left(\frac{1/r}{1-e^{-1/r}}\right)^2.$$

Since this number converges to 1 as r tends to infinity, our result can be interpreted as a generalization of the previous findings. Our result requires new combinatorial argument that takes into account the dependency between possibly of multiple best responses; the existence of a pure Nash equilibrium in one row (column) affects the probability - but does not rule out the possibility - that there is another pure Nash equilibrium in the same row (column).

The paper is organized as follows. The notation is given in Section 2. In the rest of the paper, we characterize the distribution of the number of equilibria and give the expected number of pure Nash equilibria, as the number of pure strategies goes to infinity.

2 Preliminaries

There are two players, the row player and the column player, playing a $K \times K$ matrix game. Now we define the concepts of a *payoff matrix* and a *best response matrix* of the column player. By swapping the roles of rows and columns, the definitions below apply also to the row player.

A payoff matrix of the column player defines a utility index for all action pairs $(i, j) \in \{1, ..., K\}^2$

$$U = \begin{bmatrix} u_{11} & \cdots & u_{1K} \\ \vdots & & \vdots \\ u_{K1} & \cdots & u_{KK} \end{bmatrix}.$$

Given a payoff matrix U, denote the induced best response matrix by

$$B(U) = \begin{bmatrix} b_{11}(U) & \cdots & b_{1K}(U) \\ \vdots & & \vdots \\ b_{K1}(U) & \cdots & b_{KK}(U) \end{bmatrix},$$

where

$$b_{ij}(U) = \begin{cases} 1, & \text{if } u_{ij} \geq u_{ij'}, \quad \text{for all } j' = 1, ..., K, \\ 0, & \text{otherwise.} \end{cases}$$

Note that at least one element in a row of a best response matrix must be equal to one. If U is random, then $B(U)$ is random.

Given two best payoff matrices U and U' of the row player and the column player, respectively, an action pair $(i, j) \in \{1, ..., K\}^2$ forms a *pure Nash equilibrium* (PNE) if and only if

$$b_{ij}(U)b_{ij}(U') = 1.$$

If the value of a function ϕ is dependent on a random variable x, we denote by $E_x\phi(x)$ the expected value of the function.

2.1 Random best reply matrices

Two games with the same players and the same strategy sets are *best reply equivalent* if they induce the same best response matrices. Note that all games in the same equivalence class have the same PNE. Given an equivalence class of games, we may take the corresponding best reply matrices (one for each player) as representing this class, since best reply matrices can of course be interpreted as payoff matrices. In this section we draw equivalence classes from a uniform distribution, for each number K of actions. We determine the limit probability that the chosen equivalence class of games has k PNEs as K goes to infinity.

Proposition 1 *Suppose that a best reply equivalence class is drawn from the uniform distribution over all equivalence classes, given K. The probability that a game in this class possesses exactly k PNEs, for $k = 0, 1, ...,$ goes to zero as K goes to infinity.*

Proof. Assume first that $f(K) = 2$ for all K. We may assume w.l.o.g. that 0 and 1 are the possible utility values. Let us call this situation as a $0-1$ game. First we claim that, as K becomes large, there is a row (column) in the payoff matrix of the column (row) player that contains only 0s with probability zero. To verify this, note that the probability of there being no such row (column) is

$$\left(1-\frac{1}{2^K}\right)^K.$$

To see that this goes to 1 as K becomes large, observe that

$$\left(1-\frac{1}{2^K}\right)^K = \left[\left(1-\frac{1}{2^K}\right)^{2^K}\right]^{K/2^K},$$

and that

$$\ln\left(1-\frac{1}{2^K}\right)^K = \frac{K}{2^K}\ln\left(1-\frac{1}{2^K}\right)^{2^K}$$

goes to zero as K grows to infinity. This proves the claim.

As a consequence of the previous claim, when K becomes large, each row (column) of the column (row) player's best response matrix has at least one 1 with probability one. This implies that the payoff matrix of a $0-1$ game coincides with probability one with the best response matrix. Since payoffs are randomly drawn, it follows that the distribution of $0-1$ payoff matrices also reflects almost surely the distribution of random best responses matrices.

Now let $f(K) \geq 2$. By the observation made in the previous paragraph, a game with random best response matrices has k PNEs with the same probability as a game of random $0-1$ payoffs has k PNEs with payoffs $(1,1)$. Since the utility indices of a $0-1$ game are drawn independently from the set $\{0,1\}$, an event "strategy (x,y) induces payoffs $(1,1)$" is a Bernoulli trial that is independent of x and y, and has success probability $1/4$. The probability that there exist exactly $k = 0, 1, ...$ such PNEs is binomially distributed, and equals

$$\begin{pmatrix} K^2 \\ k \end{pmatrix}\left(1-\frac{1}{4}\right)^{K^2-k}\left(\frac{1}{4}\right)^k = \frac{K^2!}{(K^2-k)!k!}\left(\frac{3}{4}\right)^{K^2-k}\left(\frac{1}{4}\right)^k.$$

This number goes to zero as K goes to infinity. ■

An immediate corollary of the previous proposition is that the probability of there being exactly 0 PNEs goes to zero as K becomes large.

Corollary 2 *Suppose that a best reply equivalence class is drawn from the uniform distribution over all equivalence classes, given K. The probability that a game in this class possesses at least one PNE goes to one as K goes to infinity.*

Since there is no cluster point in the distribution of the number PNEs as K becomes large, no sequence of expected number of PNEs can converge to a finite number.

Corollary 3 *Suppose that a best reply equivalence class is drawn from the uniform distribution over all equivalence classes, given K. The expected number of PNEs goes to infinity as K goes to infinity.*

3 Random payoff matrices

In this section, we let the utility indices be the primitive of the model. We assume that the players payoffs are drawn independently and uniformly from the set $\{1/f(K), ..., 1\}$, where $f(K)$ is a natural number. We also assume that there is a nonnegative real number r such that

$$\lim_{K\to\infty} \frac{f(K)}{K} = r.$$

First we observe the following lower bound on the number of pure PNEs in the limit game. When payoffs for a $K \times K$ game are taken from the set $\{1/f(K), \ldots, 1\}$, the best possible PNE is the one with payoffs $(1,1)$. Our result says that the distribution of number of such best equilibria is approximately Poisson with mean $1/r^2$ as K is large.

Proposition 4 *The number of PNE with payoffs* $(1,1)$ *is Poisson distributed with mean* $1/r^2$ *as* K *goes to infinity.*

Proof. Fix $r > 0$. The probability that an action profile results in payoffs $(1,1)$ gets arbitrarily close to $1/r^2K^2$ as K grows. The probability of payoffs $(1,1)$ for a given action profile is independent of the realization of the payoffs for other action profiles. For a large fixed K, the number of action profiles with payoffs $(1,1)$ is approximately binomially distributed with success probability $1/r^2K^2$. The number of trials is K^2 and so $K^2 \cdot (1/r^2K^2) = 1/r^2$. By the well-known approximation theorem, the limit distribution is Poisson with mean $1/r^2$. ■

As a corollary of the previous proposition it follows that the probability of at least one PNE with payoffs $(1,1)$ converges to $1 - e^{-1/r^2}$ as K becomes large. However, for all K there is also a positive probability that a PNE materializes with payoffs strictly lower than 1. As long as $r > 0$, this probability does not vanish when K becomes large, and it needs to be taken into account when evaluating the distribution of PNEs.

Let U be a payoff matrix of, say, the *column player* in a $K-$game. Denote the proportion of *rows* in which he has n best responses of all rows by

$$\alpha^K(n:U) = \frac{\sum_i I\left\{\sum_j b_{ij}(U) = n\right\}}{K}, \quad \text{for all } n = 1, ..., K.$$

Also denote the proportion of *columns* in which the column player has m best responses of all columns by

$$\beta^K(m:U) = \frac{\sum_j I\left\{\sum_i b_{ij}(U) = m\right\}}{K}, \quad \text{for all } m = 1, ..., K.$$

Since the payoff matrix U is a random variable, $\alpha^K(m:U)$ and $\beta^K(m:U)$ are random variables.

Similarly, by changing the roles of columns and rows, if U is a payoff matrix of the *row player* in a $K-$game, then $\alpha^K(n:U)$ reflects the proportion of *columns* in which the row player has n best responses and $\beta^K(m:U)$ the proportion of *rows* in which the row player has m best responses.

Since the total number of best responses of a player is independent on whether one counts them on the basis of columns or of rows, the average number of best responses in a row or in a column is the same. Given U, denote this average by

$$\mu^K(U) := \sum_{m=1}^{K} \alpha^K(n:U)n = \sum_{m=1}^{K} \beta^K(m:U)m. \tag{1}$$

Define

$$\bar{\mu} = \frac{1/r}{1-e^{-1/r}}.$$

Lemma 5 *$\mu^K(U)$ converges to $\bar{\mu}$ as K goes to infinity.*

Proof. The probability that the number of, say, the column player's best responses in a row is n is the probability that n actions generate the same payoff v times the probability that all other actions generate lower payoffs, given v. Since the distribution over the set $\{1/f(K), 2/f(K), ..., 1\}$ is uniform, we have, under given K,

$$\begin{aligned} E_U\alpha^K(n:U) &= \sum_{v=1}^{f(K)} \begin{pmatrix} K \\ n \end{pmatrix} \left(\frac{1}{f(K)}\right)^n \left(\frac{v-1}{f(K)}\right)^{K-n} \\ &= \begin{pmatrix} K \\ n \end{pmatrix} \left(\frac{1}{f(K)}\right)^n \sum_{x=1}^{f(K)} \left(1-\frac{v}{f(K)}\right)^{K-n}, \end{aligned}$$

where the second equality follows by reversing the order of summation. Letting K become large,

$$\begin{aligned} \lim_K E_U\alpha^K(n:U) &= \lim_K \frac{1}{n!}\left(\frac{K}{f(K)}\right)^n \sum_{v=1}^{f(K)} \left(1-\frac{v}{f(K)}\right)^{K-n} \\ &= \frac{\sum_{v=1}^{\infty} e^{-v/r}}{r^n n!} \\ &= \frac{e^{-1/r}}{(1-e^{-1/r})r^n n!}, \end{aligned}$$

where the second equality follows by taking a component wise limit of the summation. Since best responses in distinct rows of the column player are independently distributed, it follows, by the law of large numbers, that

$$\lim_K \alpha^K(n:U) = \lim_K E_U\alpha^K(n:U), \quad \text{for all } n.$$

Thus

$$\begin{aligned}
\lim_K \mu^K(U) &= \lim_K \sum_{n=1}^{f(K)} n\alpha^K(n:U) \\
&= \sum_{n=1}^{\infty} \frac{ne^{-1/r}}{(1-e^{-1/r})r^n n!} \\
&= \frac{e^{-1/r}/r}{1-e^{-1/r}} \sum_{n=1}^{\infty} \frac{1}{r^{n-1}(n-1)!} \\
&= \frac{1/r}{1-e^{-1/r}},
\end{aligned}$$

where the final equality follows from noting that $\sum_{n=1}^{\infty}\left[r^{n-1}(n-1)!\right]^{-1}$ is a Taylor expansion of $e^{1/r}$. ■

Lemma 6 *Let $\theta_K(k)$ be the probability that a randomly selected column of a K-matrix contains k PNEs. Then*

$$\lim_{K\to\infty} K \cdot \theta_K(k) = \begin{cases} \bar{\mu}^2, & \text{if } k = 1, \\ 0, & \text{if } k > 1. \end{cases} \tag{2}$$

Proof. Fix an arbitrary column. Let there be m column player's best responses and n row player's best responses in this column. Since each allocation of the given m and n best responses in the column is equally likely, the number $k \le \max\{m,n\}$ of PNEs in the column is hypergeometrically distributed. The probability $\eta_K(k:m,n)$ of k PNEs is

$$\begin{aligned}
\eta_K(k:m,n) &= \frac{\binom{m}{k}\binom{K-m}{n-k}}{\binom{K}{n}} \\
&= \frac{m!n!}{(m-k)!(n-k)!}\frac{(K-n)!(K-m)!}{K!(K-m-n+k)!}.
\end{aligned}$$

As K becomes large,

$$\lim_K K \cdot \eta_K(k:m,n) = \begin{cases} mn, & \text{if } k = 1, \\ 0, & \text{if } k > 1. \end{cases} \tag{3}$$

By definition,

$$\theta_K(k) = E_{m,n}\eta_K(k:m,n).$$

Since the column player's and the row player's best responses are independently distributed, n and m are independent random variables. Hence, by (3),

$$\lim_K K \cdot \theta_K(k) = \begin{cases} \lim_K \left(E_U \sum_n \alpha^K(n:U)n \right) \left(E_U \sum_m \beta^K(m:U)m \right), & \text{if } k=1, \\ 0, & \text{if } k>1. \end{cases} \tag{4}$$

By (1) and Lemma 5, (4) implies (2). ■

Proposition 7 *The number of PNEs is Poisson distributed with mean* $\bar{\mu}^2$ *as* K *goes to infinity*

Proof. Let K be large so that, by the law of large numbers, $\beta^K(\cdot : U)$ can be interpreted as the empirical distribution of a K-sequence of independent trials drawn from the distribution $\beta^K(\cdot : U)$ itself. Because of this independency, $K \cdot \theta_K(k)$ approximates the expected number of columns that contain exactly k PNE. By Lemma 6, this expectation is close to zero for all $k > 1$. Hence any column contains almost surely at most one Nash equilibrium, and, counting on the basis of columns, the existence of a PNE in a column can be interpreted as an independent Bernoulli trial with success rate $\theta_K(1)$. This means that the total number of PNEs is binomially distributed with success rate $\theta_K(1)$. By the standard approximation result, Lemma 6 implies that the number of PNE is Poisson distributed with mean equal to $\lim_K K \cdot \theta_K(1) = \bar{\mu}^2$. ■

Thus, in the limit, the expected number of PNE is

$$\bar{\mu}^2 = \left(\frac{1/r}{1 - e^{-1/r}} \right)^2$$

and the probability of *at least* one PNE is

$$1 - e^{-\bar{\mu}^2} = 1 - e^{-[r(1-e^{-1/r})]^{-2}}.$$

For example, when the number of utility indices grows with the same speed as the number of choices, i.e. $r = 1$, the expected number of Nash equilibria in the limit is $\left(1 - e^{-1}\right)^{-2} \approx 2.502$ and the probability of at least one PNE is $1 - e^{-\left(1-e^{-1}\right)^{-2}} \approx 0.329$.

Since

$$\lim_{r \to \infty} \frac{1/r}{1 - e^{-1/r}} = 1,$$

Proposition 7 implies the following classic result (see e.g. Goldberg et al., 1968; Powers, 1990; Stanford, 1995a,b): when payoffs are drawn from a set that is much (infinitely times) larger than the set of choices, the number of pure PNE is Poisson distributed with mean 1 as the set of choices becomes large. Conversely,

$$\lim_{r \to 0} \frac{1/r}{1 - e^{-1/r}} = \infty$$

implies that when payoffs are drawn from a set that is small relative to the size of the game the number of PNE approaches infinity, a result parallel to Proposition 4.

Note also that the ratio between Poisson mean r^{-2} in Proposition 4 - the lower bound of the expected number of equilibria - and the Poisson mean $\bar{\mu}^2$ in Proposition 7, i.e.

$$\left(\frac{1}{1-e^{-1/r}}\right)^2$$

tends to one as r goes to 0, reflecting the fact that when the set of utility indices is small relative the size of the game, most of the PNE are with maximal payoffs.

4 A note on the limit game

The natural limit game when K increases without limit is the one in which both players have $\mathbb{N} = \{0, 1, \ldots\}$ as their strategy sets. If $f(K)$ increases without limit as well, then the uniform distribution over $\{1/f(K), \ldots, 1\}$ weakly converges to the uniform distribution over $[0, 1]$. Assume indeed that the strategy sets are $\mathbb{N}$ and payoffs to both players and to each strategy pair are *i.i.d.* draws from the uniform distribution over $[0, 1]$. In this game there are *no* pure Nash equilibria with probability 1. To see this, note that player $i = 1, 2$ gets utility strictly less than 1 from every strategy pair with probability 1. Hence a Nash equilibrium (b, b) should be such that player $i = 1$, say, gets equilibrium payoff $y < 1$. But with probability one he gets payoff $x > y$ from some other action $b' \neq b$. This is one reason why the limit results are of interest: if there were pure Nash equilibria in the limit game, then such an equilibrium might qualify as an approximate solution to a large but finite matrix game.

References

[1] Bade, S., Haeringer, G., and L. Renou (2007), More strategies, more Nash equilibria, Journal of Economic Theory 135, pp 551-557

[2] Dresher, M. (1970), Probability of a pure Nash equilibrium point in n-person games. J. Combinat. Theory 8, pp. 134–145.

[3] Goldberg, K., Goldman, A. and M. Newman (1968), The probability of an equilibrium point. J. Res. Nat. Bur. Stand. U.S.A. 72B, pp. 93–101.

[4] McLennan, A. (2005), The Expected Number of Nash Equilibria of a Normal Form Game, Econometrica 73, pp 141-174.

[5] Nash, J. (1951), Noncooperative games. Ann. Math. 54, pp. 286–295.

[6] Powers, I. (1990), Limiting distributions of the number of pure strategy Nash equilibria in N-person games. Int. J. Game Theory 19, pp. 277–286.

[7] Rinott, Y. and M. Scarsini (2000), On the Number of Pure Strategy Nash Equilibria in Random Games, Games and Econom. Behav. 33, pp 274-293

[8] Stanford, W. (1995a), A note on the probability of k pure Nash equilibria in matrix games. Games and Econom. Behav. 9, pp. 238–246.

[9] Stanford, W. (1995b), On the distribution of pure strategy equilibria in finite games with vector payoffs. Mathematical Social Sciences. 2, pp. 115-127.

[10] Stanford, W. (1996), The Limit Distribution of Pure Strategy Nash Equilibria in Symmetric Bimatrix Games. Math. Oper. Res. 21, pp. 726–733