A COMBINATORIAL AND EXPERIMENTAL ANALYSIS OF NASH EQUILIBRIA FOR FINITE COMMON PAYOFF GAMES

RICHARD VOEPEL

Abstract. In this paper we build upon the work of Stanford [1] in analyzing the distribution of the number of pure Nash equilibria in "random" two player finite common payoff games. By utilizing a combinatorial framework, complete distribution information for the cases of $2 \times n$ games are found. Larger games can be analyzed in the same fashion, though the complexity is best handled through the use of computer algebra systems, such as Maple$^{\text{TM}}$. We also show the lack of need for this information for practical purposes by casting the problem into an experimental framework, collecting data on two player games with larger strategy sets and three player games on small strategy sets.

1. Preamble

In William Stanford's "On the number of pure strategy Nash equilibria in finite common payoffs games" [1], a healthy amount of time is spent discussing the link between political philosophy, identical preferences, and common payoff games. Like Stanford, we are primarily concerned only with probabilistic features of "random" two player finite common payoff games, and as such will not repeat a summary of those topics here. We instead focus on extending the analysis Stanford began by tackling (in part) the open question of complete information on the probability distribution of the number of pure Nash equilibria (PNEs) in these games.

Following our exact analysis, we turn to discuss the role of experimental mathematics in studying problems where data can be simulated, which is the case here. For practical purposes, there is very little reason beyond innate curiosity for providing rigorous proofs and exact formulae for the probability distributions in question; not only are the requirements for being "random" exceedingly strict and therefore unlikely to occur naturally, but it takes only a day on a modest personal computer with unoptimized code to simulate millions of data points for small sized games, which ultimately can provide highly accurate estimates of the distributions for a fraction of the human effort. In a situation where this information could be useful to an individual working in the field, tabulating statistics once and for all is likely to be faster, cheaper, and just as effective when compared to a "purely theoretical" approach.

Author addresses: Richard Z. Voepel, Department of Mathematics, Rutgers University (New Brunswick), Hill Center-Busch Campus, 110 Frelinghuysen Rd., Piscataway, NJ 08854-8019, USA
http://www.math.rutgers.edu/∼rzv2/ : rzv2@math.rutgers.edu .

2. The Combinatorial Approach

Let m and n be integers with $n \geq m \geq 2$. We define a random two player finite common payoff game in bimatrix form (A, B) to be a game satisfying the following two properties:

(1) The payoff matrices $A = (a_{ij})$ and $B = (b_{ij})$ are both of size $m \times n$, with $A = B$.
(2) The matrix A is drawn uniformly at random from the set of $m \times n$ matrices with distinct entries in $\{1, 2, \ldots, mn\}$.

As Stanford notes in [1], this second property is, for the purposes of finding probabilistic information on the number of PNEs, equivalent to saying that the individual payoffs a_{ij} are IID continuous random variables, since there is a natural reduction from such payoff matrices to weak ordinal preference matrices, and the case of non-distinct payoffs form a probability-zero event.

However, there is another natural equivalence that we can leverage – probability statements about uniform random variables on finite sets are the same as counting arguments on those sets. Thus, rather than utilizing analytic or probabilistic tools to pin down X, the random variable representing the number of PNEs in our games, we may instead simply count the total number of payoff matrices defined above that have a specified number of PNEs.

To that end, we first recall the definition of a PNE: a strategy (i^*, j^*) for a bimatrix game (A, B) is said to be a PNE if $a_{i^*j^*} \geq a_{ij^*}$ and $b_{i^*j^*} \geq b_{i^*j}$ for all appropriate i and j. As Stanford notes, this implies that the range of X in the common payoff games we are considering is from 1 to $\min(m, n)$, inclusive, since there can be at most one PNE per row and per column. This range informs us about the number of cases to consider when building our counting argument. We also introduce the set

$$\Lambda_{m,n} = \{(\lambda_2, \lambda_3, \ldots, \lambda_m) \in \mathbb{N}^{m-1} : m + n - 2 < \lambda_m < \cdots < \lambda_2 < mn\}$$

to shorten the appearance of formulae below.

2.1. The $m = 2$ Case.

We claim the following:

Fact 2.1. Let $m = 2$ and let $n \geq 2$. If X is the random variable counting the number of PNEs in a random two player common payoff game of size $2 \times n$, then the probability generating function for X is given by

$$G_{2,n}(z) = \frac{2}{n+1} z^1 + \frac{n-1}{n+1} z^2$$

Proof. By relabeling the strategy sets of the two players, we may assume that the entries along the diagonal of the payoff matrix A are maximal in their "lower" truncated submatrices, i.e. $a_{ii} \geq a_{jk}$ whenever $j, k \geq i$. In particular, this means the first entry $a_{11} = 2n = \lambda_1$, and is always a PNE. Since there can only be two possible values for X in this case, and our final probabilities must sum to 1, it is sufficient to count only the case where $X = 2$.

So how many $2 \times n$ strict ordinal payoff matrices are there with exactly 2 PNEs in this canonical form? Consider the matrix below:

$$\begin{bmatrix} \lambda_1 & | & \\ \text{—} & \lambda_2 & \text{———} \end{bmatrix}$$

The requirements for how to populate the entries of this payoff matrix are simple:

(1) $\lambda_2 > 2 + n - 2$
(2) $\lambda_2 < 2n$
(3) All entries sharing a row or column with λ_2 must be less than λ_2.

Thus, for each choice of valid λ_2, we need to choose $2+n-2$ payoffs less than λ_2 and permute them amongst the shared row and column positions. We then permute the remaining payoffs amongst the last entries, and finally take into account the number of ways our relabeling can be undone and the number of ways we could arrive at the same matrix. The final count then has the following formula:

$$\sum_{\bar{\lambda} \in \Lambda_{2,n}} \frac{2! \cdot n!}{(n-2)!} \binom{\lambda_2 - 1}{2 + n - 2} \cdot (2 + n - 2)! \cdot (n-2)!$$

After dividing by $(2n)!$ to recover the probability statement and simplifying (which can be aided by a computer algebra system such as Maple™), we have:

$$\Pr(X = 2) = \frac{n-1}{n+1}$$

Thus, the probability generating function for X on $2 \times n$ games is as claimed. □

This fact completely answers Stanford's open questions relating to common payoff games for the $m = 2$ case; the probability generating function $G_{2,n}(z)$ encodes information about all moments of X, so the distribution is completely understood (at least in principle).

2.2. **The General Case.**

For $m \geq 3$ our pattern of analysis will be exactly the same; use relabeling in order to organize the payoff matrix's diagonals such that they are maximal in their "lower" truncated submatrices, consider the various cases for the value of X, and construct counting arguments as appropriate. However, the total number of cases to consider grows quite rapidly, and each case becomes quite intricate. When $m = 3$ there are 19 different counting arguments in order to find the probability generating function $G_3(z)$!

Clearly a complete, exact analysis for large values of m would require a great deal of effort by hand. Luckily the counting arguments do not change dramatically in structure, and so it is likely that generation of the appropriate formulae could be automated with a comparatively small investment. We do not concern ourselves with that problem here, as we will elaborate on later. The exception to this tedium is the case $X = m$, where there is a full number of PNEs in our payoff matrix. Consider the matrix below:

$$\begin{bmatrix} \lambda_1 & \vert & \vert & \vert & \vert & \\ \text{---} & \lambda_2 & \vert & \vert & \vert & \text{------} \\ \text{---} & \text{---} & \lambda_3 & \vert & \vert & \text{------} \\ \text{---} & \text{---} & \text{---} & \ddots & \vert & \text{------} \\ \text{---} & \text{---} & \text{---} & \text{---} & \lambda_m & \text{------} \end{bmatrix}$$

By analogy with our reasoning for the $m = 2$ case, the formula for the final count can be written as follows:

$$\sum_{\bar{\lambda} \in \Lambda_m} \frac{m! \cdot n!}{(n-m)!} \prod_{l=0}^{m-1} \binom{\lambda_{m-l} - l(m+n) + (l+1)^2 - 2(l+1)}{m+n-2(l+1)} \cdot (m+n-2(l+1))!$$

While Maple™ will not simplify this expression with symbolic m, we may still pick specific values and simplify after dividing by $(mn)!$. When we do this for $m = 3$ we find that

$$\Pr(X = 3) = \frac{(n-1)(n-2)}{(n+1)(n+2)}$$

and for $m = 4$ we find that

$$\Pr(X = 4) = \frac{(n-1)(n-2)(n-3)}{(n+1)(n+2)(n+3)}$$

and for $m = 5$ we find that

$$\Pr(X = 5) = \frac{(n-1)(n-2)(n-3)(n-4)}{(n+1)(n+2)(n+3)(n+4)}$$

which leads to an obvious conjecture:

Conjecture 2.2. For a random two player $m \times n$ finite common payoff game

$$\Pr(X = m) = \prod_{i=1}^{m-1} \frac{n-i}{n+i}$$

3. Experimental Results on Larger Games

Given the immense human effort required to pin down the probability generating functions for the cases $m \geq 3$, clearly we should focus our efforts on automatically generating the counting arguments, shouldn't we? We claim that no, we should in fact not. While having complete information about the probability distributions in question is certainly appealing, for applications in the real world more likely scenarios would be that the game being studied is "small" in size, the hypotheses for random two player finite common payoff games won't actually be satisfied, or both.

If we are in such a scenario, then how might we proceed more efficiently with our analysis? The answer lies within the experimental mathematics paradigm; we will proceed by computing extensive statistics for the type of game we are studying through simulation of many instances of the game. These statistics can be as accurate as we need them to be within the constraints of measurement error and the quality of our random number generators.

In the tables below, we have provided some statistics regarding two player random common payoff games of small sizes and three player games of very small sizes. For the two player games each experiment ran 3 million trials, and for the three player games each experiment ran 1 million trials. At least for probability, expectation, and variance calculations, the propagated error from using these statistics can be made small with reasonably sized experiments. This data can be found in raw form at http://sites.math.rutgers.edu/~rzv2/Nash, along with Maple™packages and scripts for generating new experiments.

Probabilities for Two Player Random Common Payoff Games

$X = k$	4×4 Probabilities	5×5 Probabilities	6×6 Probabilities	7×7 Probabilities
1	.11428	.03987	.01286	.00406
2	.51451	.31712	.16239	.07322
3	.34280	.47604	.43321	.30576
4	.02841	.15900	.32421	.40815
5		.00798	.06510	.18373
6			.00223	.02449
7				.00058

Probabilities for Three Player Random Common Payoff Games

$X = k$	$2 \times 2 \times 2$ Probabilities	$3 \times 3 \times 3$ Probabilities
1	.2145090	.0044700
2	.6064810	.0804030
3	.1432210	.2933770
4	.0357890	.3607990
5		.2037790
6		.0503130
7		.0065140
8		.0003320
9		.0000130

References

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