Complex Systems **1** (1987) 235–244

The Value of a Random Game: The Advantage of Rationality

William G. Faris
Robert S. Maier
Department of Mathematics, University of Arizona, Tucson, Arizona 85721, USA

Abstract. Two players play against each other in a game with payoffs given by a random n by n matrix with mean zero. If one player adopts a uniform, purely random strategy, then his loss is limited by the law of averages to a quantity proportional to $\sqrt{\log n}/\sqrt{n}$. On the other hand, if he plays an optimal strategy his losses will typically be considerably less. Numerical evidence is presented for the following conjecture: the standard deviation of the value of the game is asymptotically proportional to $1/n$. This smaller loss exhibits the advantage of rationality over randomness. The rational player, moreover, tends as $n \to \infty$ to employ a strategy vector that has half its components zero.

1. Introduction

The asymptotic properties of random two-person zero-sum games with a large number of strategies have been considered in work of J. E. Cohen and C. M. Newman [1]. They were interested in the stabilizing effect of strategy diversification in evolutionary biology. In the course of their work they bounded the value of the game, *i.e.*, the expected payoff of one player to the other if both adopt optimal strategies, in terms of the expected payoff obtaining if one of the players uses a uniform (purely random) strategy. In part 3 of this paper we present numerical results giving evidence for the actual asymptotic behavior of the value. We also report on the asymptotics of the optimal strategies.

Random game theory is not yet a subject amenable to extended analytic investigation. To a given game or game matrix, a variety of formulae and algorithms may be applied to yield the value, the optimal strategies, *etc.* But if the game matrix is random, these will be random variables. Their distribution is not easy to determine analytically for any fixed matrix size, still less their asymptotics as the matrix size tends to infinity. Part 2 contains such analytic results as we have.

The games we consider are of the simplest kind, where there are two players in strict competition. Each of the two players has a finite number of choices. The consequence of a pair of choices is given by a payoff matrix. Each player makes a choice in secret. The two choices are then revealed, and the matrix determines the payment from one player to the other. Since the sum of the gains is zero, this type of game is known as a two-person zero-sum game.

We usually interpret the payoff matrix as a matrix giving the payment from the player who chooses the columns to the player who chooses the rows. The column player thus prefers small payoffs, while the row player prefers large payoffs. Thus in the game

$$A = \begin{pmatrix} 1 & 0 \\ 2 & 3 \end{pmatrix} \tag{1.1}$$

it is apparent that the row player will choose the second row. The column player may yearn after the 0 entry in the (1, 2) position, but it does no good; his only choice now is the first column. Thus the result of the game is that the row player wins the amount 2 given in the (2, 1) entry.

The players are allowed to choose mixed strategies, that is, strategies based on probabilities. In the game

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 0 \end{pmatrix} \tag{1.2}$$

the column player may choose to employ the uniform strategy, that is, to play the columns with equal probabilities. Then the expected winning for the row player is 3/2, no matter what choice he makes. One can also look at the game from the point of view of the row player. He has a low stakes choice (the first row) and a high stakes choice (the second row). If the row player tries to play a uniform strategy employing the two rows with equal probability, then the column player may play the second column and limit his expected losses to 1. The optimal strategy for the row player is more complicated; he should play the first row with probability 3/4 and the second row with probability 1/4. Then the expected loss by the column player is 3/2, no matter what choice he makes. Clearly if both players play optimally the expected payment from the column player to the row player is 3/2. The value of this game to the row player is thus 3/2.

The situation we are interested in is a situation in which the games are generated by nature, so that the game matrix itself is random. The conceptual experiment is conducted as follows. A number of game matrices are independently generated. When a game matrix is generated, it is made public. The two players each choose a (mixed) strategy for this matrix. They then play this game, perhaps a number of times, so that the outcome of the game for the row player is considered to be the expected winnings with these strategies.

We are actually interested in large matrices. Consider an m by n matrix. As above, the payoff from the row player to the opposing column player is defined to be the (i,j)th matrix element if they choose the ith and jth alternatives respectively. The entries of the matrix are independent, identically distributed random variables with a distribution that is symmetric about zero. When $m = n$ all probabilities are symmetric between the row player and the column player. In part 2 we nonetheless show that if one player adopts a uniform (that is, completely random) strategy, the other may adopt a strategy that ensures gains that are asymptotically proportional to $\sqrt{\log n}/\sqrt{n}$ as $n \to \infty$.

But according to the numerical results of part 3, if the players use their optimal strategies the average return to the players (*i.e.*, the value of the game) appears to have standard deviation asymptotically proportional to $1/n$. These contrasting asymptotics indicate that rational play, which takes advantage of the structure of the particular matrix, produces a more even outcome.

A striking feature of rational play, as revealed by our investigations, is the asymptotic behavior of the number of zero components of the optimal strategy vectors. The numerics suggest that the proportion of zero components is asymptotically equal to $1/2$, with the variance of the proportion tending to zero as $1/n$. This is consistent with a rigorous result of part 2; namely, that the probability of the extreme cases, *i.e.*, of the optimal strategy vector having $n-1$ zeroes or no zeroes, decays at least exponentially in n.

2. Analytical results

In the following we shall have occasion to refer to an n-component vector x as a *probability vector* if all its entries are greater than or equal to zero and the sum of the entries is one. Thus we have

$$\mathbf{x} \geq 0 \qquad \sum \mathbf{x} = 1. \tag{2.1}$$

The probability vectors range over a simplex of dimension $n-1$. They represent strategies for a player with n alternatives.

The game theory interpretation applies to an arbitrary real matrix. Let A be an m by n matrix. Player 1 (the row player) has m alternatives and player 2 (the column player) has n alternatives. This matrix represents the payoff from player 2 to player 1 corresponding to a choice of alternatives by each of the players. The gain of player 1 is the loss of player 2. Thus player 1 wants to maximize the payoff, and player 2 wants to minimize it.

Let $\mathbf{x}$ be an m-component column probability vector representing a strategy for player 1, and let $\mathbf{y}$ be an n-component column probability vector representing a strategy for player 2. Write $\mathbf{x}'$ for the m-component row vector that is the transpose of $\mathbf{x}$. Then if player 1 plays strategy $\mathbf{x}$, the average payoffs for each choice by player 2 are represented by the n-component row vector $\mathbf{x}'A$. The number $\min \mathbf{x}'A$ is a lower bound on the average gain of player 1 that is independent of the choice of the other player. Similarly, if player 2 plays strategy $\mathbf{y}$, the average payoffs for each choice by player 1 are represented by the m-component column vector $A\mathbf{y}$. Again, $\max A\mathbf{y}$ is an upper bound on the average loss by player 2.

It is easy to see that

$$\min \mathbf{x}'A \leq \max A\mathbf{y}. \tag{2.2}$$

In other words, the least gain to 1 due to a choice of strategy by player 1 involves a smaller payoff from 2 to 1 than the greatest loss to 2 due to a choice of strategy by player 2. In fact, we have

$$\min \mathbf{x}'A \leq \mathbf{x}'A\mathbf{y} \leq \max A\mathbf{y}, \tag{2.3}$$

which says that the actual average payoff $\mathbf{x}'A\mathbf{y}$ due to choices by both players is between the two extremes.

The minimax theorem of game theory states that there exists optimal strategy vectors such that equality holds [4]. In this case the common value v is the *value* of the game. For every pair of probability vectors $\mathbf{x}$ and $\mathbf{y}$ we have the bound

$$\min \mathbf{x}'A \leq v \leq \max A\mathbf{y}. \tag{2.4}$$

In our case, the matrix A is random, and its value is a random variable. Since the optimal strategy vectors depend on the matrix, they are also random. In the following we assume that the entries of the matrix are independent and have the same distribution, with mean zero and standard deviation σ.

An alternative to the optimal strategy is the adoption of a uniform strategy by one or the other of the players. We take $\mathbf{x}$ and $\mathbf{y}$ to be the constant vectors $1/m$ and $1/n$. If both players adopt this uniform strategy, the average payoff is

$$\frac{1}{mn}\sum_i \sum_j A_{ij}. \tag{2.5}$$

If the entries of the matrix are independent with mean zero and standard deviation σ^2, then the average payoff has mean zero and variance $\sigma^2/\sqrt{mn}$. If $m = n$ this is proportional to $1/n$. Thus this rather dull and stupid play by both players results in a rather close outcome, by the law of averages.

If one of the players adopts the uniform strategy, the worst case analysis gives a bound on the value of the game. The value v is bounded below and above by the minimum column average and maximum row average of A:

$$\min_j \frac{1}{m}\sum_i A_{ij} \leq v \leq \max_i \frac{1}{n}\sum_j A_{ij}. \tag{2.6}$$

The column and row averages $\frac{1}{m}\sum_i A_{ij}$ and $\frac{1}{n}\sum_j A_{ij}$ have mean zero and standard deviations $\sigma/\sqrt{m}$ and $\sigma/\sqrt{n}$. What remains is to take into account the effect of the maximum and minimum.

Take the case of the maximum loss by player 2 after his adoption of a uniform strategy. Assume for simplicity that the distributions of the entries of the matrix is Gaussian. Then the distribution of the row averages is also Gaussian.

The theory of maxima of independent, identically distributed random variables is well known [2]. If $Z_1, \ldots, Z_n$ are independent random variables with the same distribution, and $U_m = \max(Z_1, \ldots, Z_m)$, then

$$\mathcal{P}[U_m \leq t] = \mathcal{P}[Z \leq t]^m = (1 - \mathcal{P}[Z > t])^m. \tag{2.7}$$

From this we see that the interesting asymptotics are when $\mathcal{P}[Z > t_m] \approx 1/m$. For the Gaussian distribution with mean zero and standard deviation σ this is when $t_m \approx \sigma\sqrt{2 \log m}$. In fact, it is known that in this circumstance

$$\mathcal{P}[|U_m - \sigma\sqrt{2 \log m}| < \epsilon] \to 1 \tag{2.8}$$

as $m \to \infty$.

In the game theory situation, this says that for the uniform strategy $\mathbf{y}$ we have

$$\mathcal{P}[|\max A\mathbf{y} - \sigma\sqrt{2 \log m}/\sqrt{n}| < \epsilon] \to 1 \tag{2.9}$$

as $m \to \infty$, uniformly in n. When $m = n$, this gives the ratio $\sqrt{\log n/n}$, which may be regarded as roughly proportional to $1/\sqrt{n}$ for most practical purposes.

We now turn to the case when the entries are not assumed to be Gaussian. However the row and column averages will still be approximately Gaussian, so we expect similar results. The following result is obtained by a refinement of the method of Cohen and Newman [1].

Theorem 1. *Consider an m by n matrix A with independent, identically distributed entries A_{ij}, having mean zero and standard deviation σ. Assume that the moment generating function $\mathcal{E}[\exp(tA_{ij})]$ exists for all real t sufficiently near to zero. Let $\mathbf{y}$ be the strategy that is uniformly $1/n$. Then for every $\sigma_1 > \sigma$ and $\epsilon > 0$,*

$$\lim_{m\to\infty} \mathcal{P}[\max A\mathbf{y} \leq \frac{\sigma_1(\sqrt{2 \log m} + \epsilon)}{\sqrt{n}}] = 1. \tag{2.10}$$

Lemma 1. *Let $\sigma_1 > \sigma$. Then the row average will satisfy a probability estimate of the form*

$$\mathcal{P}[\frac{1}{n}\sum_j A_{ij} \geq s] \leq \exp(-\frac{ns^2}{2\sigma_1^2}) \tag{2.11}$$

for small $s > 0$.

Proof: By Chebyshev's inequality, independence, and Taylor's theorem

$$\mathcal{P}[\frac{1}{n}\sum_j A_{ij} \geq s] \leq \mathcal{E}[\exp(tA_{ij})]^n \exp(-nst) \leq \exp(\frac{n\sigma_1^2}{2}t^2)\exp(-nst) \tag{2.12}$$

for small t. Take $t = s/\sigma_1^2$. ∎
Proof of theorem: For every s we have

$$\mathcal{P}[\max Ay \le s] = \prod_i \mathcal{P}[(Ay)_i \le s] = \prod_i (1 - \mathcal{P}[(Ay)_i > s]). \tag{2.13}$$

Now take the s of the lemma to be $s = \sigma_1(\sqrt{2\log m} + \epsilon)/\sqrt{n}$ as in the statement of the theorem. Then

$$\mathcal{P}[\max Ay \le a] \ge (1 - \frac{1}{m} e^{-(\epsilon\sqrt{2\log m} + \epsilon^2/2)})^m. \tag{2.14}$$

This approaches 1 as $m \to \infty$. ∎

We have thus obtained a bound on the probability of a positive fluctuation of the value of the random game. A very similar argument applied to the column averages produces a bound on the probability of a negative fluctuation.

This argument proves that the value of the n by n game should decrease at a rate proportional to $\sqrt{\log n}/\sqrt{n}$. Of course this is the bound obtained by having a player adopt a uniform (completely random) strategy, and it is plausible that there should be room for a faster decrease, if each player plays rationally.

Now we examine the strategies. The number of zero components in an n-component strategy vector may range from 0 to $n-1$. The extreme cases correspond respectively to interior points and corners of the simplex of probability vectors. In the game context, generically the optimal strategy for each player will be a corner when the game matrix has a saddle point. (A saddle point is an entry that is maximal in its column and minimal in its row.) A. J. Goldman [3] has shown that for independent matrix entries with a common continuous distribution, the probability of a saddle point is $m!n!/(m+n-1)!$. When $m = n$ this is asymptotically $(n!)^2/(2n-1)! \sim 2n\sqrt{\pi n}4^{-n}$, which decays exponentially as $n \to \infty$.

The opposite extreme, that of a strategy vector lying in the interior of the simplex, also has exponentially decreasing probability as $n \to \infty$. Our result is given by the following theorem.

Theorem 2. *Assume that the n by n matrix A has independent, identically distributed Gaussian random variables as matrix elements. Then the probability that the optimal strategy is in the interior of the simplex is bounded by 2^{1-n}.*

Proof: Denote by $\mathcal{O}_\pm$ the union of the positive and negative orthants in $\mathbf{R}^n$, *viz.*, $\{\mathrm{x} \in \mathbf{R}^n \mid \forall i\, \mathrm{x}_i > 0 \text{ or } \forall i\, \mathrm{x}_i < 0\}$. Then a standard result [4] has it that an interior point obtains only if the two events E1 and E2 given respectively by $A^{-1}\mathbf{1} \in \mathcal{O}_\pm$ and $A'^{-1}\mathbf{1} \in \mathcal{O}_\pm$ both occur. Here $\mathbf{1}$ denotes the n-component vector with unit entries, and A' is the transpose of A. The standard result assumes that the value of A is non-zero, but this occurs with probability one.

Both E1 and E2 have probability 2^{1-n}. To verify this, note that the probability measure on the space of matrices is invariant under right multiplication by every orthogonal matrix Q. The matrix Q may be taken random, so long as it is taken independent of A. So,

$$\begin{aligned} \mathcal{P}[\mathrm{E1}] &= \mathcal{P}[A^{-1}\mathbf{1} \in \mathcal{O}_\pm] \\ &= \mathcal{P}[(AQ)^{-1}\mathbf{1} \in \mathcal{O}_\pm] \\ &= \mathcal{P}[Q^{-1}A^{-1}\mathbf{1} \in \mathcal{O}_\pm] \\ &= \mathcal{P}[A^{-1}\mathbf{1} \in Q\mathcal{O}_\pm]. \end{aligned} \tag{2.15}$$

If the measure on the orthogonal matrices Q is taken to be Haar measure (the group invariant measure), this is simply twice the probability of an n-component vector lying in a random orthant, i.e., 2^{1-n}. A similar argument suffices for E2. The bound in the statement of the theorem is an immediate consequence of this probability calculation. ∎

The conclusion of this discussion is that for large matrices neither extreme, corner or interior, is likely to occur. The optimal strategy vector will not be concentrated on one component, but it will still have some zeros.

3. Numerical results

The method used for our numerical results is standard, and we give only the briefest description. Let A be the game matrix, and let $\mathbf{y}$ be a probability vector representing the strategy of player 2. Assume that λ is a constant such that the vector $A\mathbf{y} \leq \lambda$. (Thus player 2 can be sure that his losses are limited by λ.) Since the value $v \leq \max A\mathbf{y} \leq \lambda$, the minimum value of λ is obtained when $\mathbf{y}$ is an optimal strategy vector and $v = \lambda$.

Thus the problem of finding the value v of the game is the problem of finding the minimum value of λ subject to

$$A\mathbf{y} \leq \lambda \qquad \mathbf{y} \geq 0 \qquad \sum \mathbf{y} = 1. \tag{3.1}$$

Now assume that the value of the game $v > 0$. (This may always be arranged by adding a constant to the matrix A.) Make the change of variables $\mathbf{z} = \mathbf{y}/\lambda$. Then the problem is to find the maximum value of $1/\lambda = \sum \mathbf{z}$ subject to

$$A\mathbf{z} \leq 1 \qquad \mathbf{z} \geq 0. \tag{3.2}$$

This is thus the linear programming problem of finding the maximum of a linear function on a convex polyhedron. This maximum is to be assumed at a vertex. In the present context the simplex method is to start at the known vertex $\mathbf{z} = 0.$ and move at each step to an adjacent vertex with a larger value of $1/\lambda = \sum \mathbf{z}$. After a finite number of steps the maximum is achieved.

This algorithm has been implemented in Fortran and run on a mini-supercomputer manufactured by Scientific Computing Systems, consuming about 200 hours of computer time. We present our results in Tables 1 and

n	$\hat{P}[z=n-1]$	$\hat{P}[z=0]$	$\bar{d}$	$s_z/\sqrt{n}$	ns_v
2	0.67	0.33	-0.085	0.3337	1.248
3	0.30	0.099	-0.048	0.3457	1.359
4	0.11	0.028	-0.033	0.3482	1.425
5	0.041	0.0076	-0.021	0.3514	1.457
6	0.013	0.0022	-0.015	0.3521	1.483
7	0.0042	0.00045	-0.011	0.3527	1.494
8	0.0013	0.00014	-0.009	0.3530	1.511
9	0.00041	0.00001	-0.000	0.3522	1.516
10	0.00012	0.00002	-0.005	0.3537	1.521
11	0.00002		-0.008	0.3529	1.529
12			0.000	0.3537	1.533
13			0.002	0.3525	1.541
14			0.003	0.3532	1.536
15			-0.002	0.3536	1.540
16			0.002	0.3531	1.549
30			0.002	0.3542	1.557
31			0.007	0.3538	1.561

Table 1: Statistics for Gaussian entries

2, which correspond to the two cases of Gaussian and uniformly distributed A_{ij}. In the former case $A_{ij} = Z_{ij}(0,1)$, a Gaussian with mean zero and unit variance. In the latter the A_{ij} have the same mean and variance, but are uniformly distributed over $[-\sqrt{3}, \sqrt{3}]$. In both cases we generated 100,000 n-square pseudo-random game matrices A for each value of n between 2 and 16, as well as for 30 and 31, and computed v for each A. The number of zero components of the optimum strategy vector was determined as well.

If z is the random variable equal to the number of zero components of this vector, Goldman's formula gives in both cases the exact expression $\mathcal{P}[z=n-1] = (n!)^2/(2n-1)!$. Our program counted strategy vectors with $n-1$ or no zeroes, and we present in the tables their frequencies to two significant figures as estimates of $\mathcal{P}[z=n-1]$ and $\mathcal{P}[z=0]$. The former affords a check on the adequacy of our data run. Notice that the frequencies of $n-1$ zeros in the two tables are very close to each other; in fact they are also very close to the exact probabilities, which decay exponentially in n. The frequencies of no zeros also appear to decay exponentially with n. The falloff is much more rapid, however, than that present in the rigorous upper bound $\mathcal{P}[z=0] \leq 2^{1-n}$ of Theorem 2. For $n \geq 12$ we have no data on $\mathcal{P}[z=n-1]$ or $\mathcal{P}[z=0]$, as none of our 100,000 matrices had $z=n-1$ or $z=0$.

The distribution of the number of zeros z is reminiscent of but not identical to a binomial distribution with $p=1/2$. The numerics suggest a mean $E[z] \approx n/2 - 0.25$ and a standard deviation $\sigma_z \approx 0.35n^{1/2}$. The evidence for this is presented in the tables. The deviations $d = z - n/2 - 0.25$

n	$\hat{P}[z=n-1]$	$\hat{P}[z=0]$	$\bar{d}$	$s_z/\sqrt{n}$	ns_v
2	0.67	0.33	-0.082	0.3329	1.433
3	0.30	0.097	-0.048	0.3441	1.604
4	0.11	0.028	-0.031	0.3482	1.661
5	0.040	0.0076	-0.016	0.3512	1.685
6	0.012	0.0020	-0.013	0.3531	1.678
7	0.0039	0.00055	-0.002	0.3525	1.667
8	0.0011	0.00020	-0.000	0.3539	1.662
9	0.00031	0.00005	-0.003	0.3554	1.649
10	0.00013	0.00002	-0.004	0.3547	1.636
11	0.00004		0.008	0.3539	1.639
12			0.008	0.3550	1.619
13			0.008	0.3545	1.630
14			0.018	0.3541	1.621
15			0.019	0.3546	1.618
16			0.008	0.3531	1.610
30			0.029	0.3536	1.598
31			0.018	0.3535	1.586

Table 2: Statistics for uniform entries

in the number of zeros were computed. The sample mean estimates $\bar{d} = \bar{z} - n/2 - 0.25$ for $E[d] = E[z] - n/2 - 0.25$ and the normalized sample standard deviation estimates $s_z/\sqrt{n}$ for $\sigma_z/\sqrt{n}$ are tabulated, the latter to four significant figures. The asymptotic independence of the choice of distribution for A_{ij} is striking.

The asymptotic behavior of the value v is apparently manifested less rapidly as $n \to \infty$ than are the asymptotics of z. We use the usual sample standard deviation estimate s_v of the standard deviation σ_v of v. Our normalized estimates ns_v of $n\sigma_v$ are given to four significant figures. They suggest that in both the Gaussian and uniform cases the standard deviation of v displays leading behavior $\sim 1.6n^{-1}$. Our data do not rule out the possibility that the constant might differ between Tables 1 and 2, *i.e.*, depend on more than the first two moments of the A_{ij}. But the evidence for a $1/n$ falloff of the standard deviation is strong. Rational play is superior to random.

Acknowledgements

We thank Charles Newman for advice and the Arizona Center for Mathematical Sciences (supported by the Air Force) for support of computing.

References

[1] Joel E. Cohen and Charles M. Newman, "Host-parasite relations and random zero-sum games: The stabilizing effect of strategy diversification", *American Naturalist*, to appear.

[2] B. V. Gnedenko, "Sur la distribution limité du terme maximum d'une série aléatoire", *Annals of Mathematics*, **44** (1943) 423–453.

[3] A. J. Goldman, "The probability of a saddlepoint", *American Mathematical Monthly*, **64** (1957) 729–730.

[4] J. C. C. McKinsey, *Introduction to the Theory of Games*, (McGraw-Hill, 1952).