ON THE OPTIMAL STRATEGY IN A RANDOM GAME

JOHAN JONASSON [1]

Dept of Mathematics, Chalmers University of Technology, S-412 96 Goteborg, Sweden.
email: `jonasson@math.chalmers.se`

Abstract

Consider a two-person zero-sum game played on a random $n \times n$-matrix where the entries are iid normal random variables. Let Z be the number of rows in the support of the optimal strategy for player I given the realization of the matrix. (The optimal strategy is a.s. unique and Z a.s. coincides with the number of columns of the support of the optimal strategy for player II.) Faris an Maier [4] make simulations that suggest that as n gets large Z has a distribution close to binomial with parameters n and $1/2$ and prove that $P(Z = n) \leq 2^{-(k-1)}$. In this paper a few more theoretically rigorous steps are taken towards the limiting distribution of Z: It is shown that there exists $a < 1/2$ (indeed $a < 0.4$) such that $P\left((\frac{1}{2} - a)n < Z < (\frac{1}{2} + a)n\right) \to 1$ as $n \to \infty$. It is also shown that $\mathbb{E}Z = (\frac{1}{2} + o(1))n$.

We also prove that the value of the game with probability $1 - o(1)$ is at most $Cn^{-1/2}$ for some $C < \infty$ independent of n. The proof suggests that an upper bound is in fact given by $f(n)n^{-1}$, where $f(n)$ is any sequence such that $f(n) \to \infty$, and it is pointed out that if this is true, then the variance of Z is $o(n^2)$ so that any $a > 0$ will do in the bound on Z above.

1   Introduction

A two-person zero-sum game is a game played on an $m \times n$-matrix $A = [a_{ij}]$ (known to both players) where player I and player II simultaneously tell the number of a row, i, and column, j, respectively. Player II then pays a_{ij} dollars to player I (where, of course, a negative a_{ij} is interpreted as player I paying $-a_{ij}$ to player II). The obvious question is how the two players should make their choices. In case A contains a *saddle point*, i.e. an entry which is smallest in its row and greatest in its column, then one quickly realizes that player I should pick the row and player II the column of the saddle point, but in all other cases the question is not meaningful unless one allows *randomization*, i.e. using chance to make your choice. To be precise the two players each use a probability vector, $\mathbf{p} = [p_1 \, p_2 \, \ldots \, p_m]^T$ for player I and $\mathbf{q} = [q_1 \, q_2 \, \ldots \, q_n]^T$ for player II. This is to be interpreted that player I will bet on row i with probability p_i, $i = 1, \ldots, m$, and player II will bet on column j with probability q_j, $j = 1, \ldots, n$. Randomized strategies of this kind are called *mixed strategies*. The goal is now for player I to find a $\mathbf{p}$ that gives him as high expected winnings as possible and for player II to find $\mathbf{q}$ that gives her as low expected losses as possible, or equivalently gives player I as low expected winnings as possible.

According to the well-known Minimax Theorem of von Neumann and Morgenstern (see e.g. [5, Section II.4]) there exist a number V, called the *value* of the game, and mixed strategies $\mathbf{p}_0$ and $\mathbf{q}_0$, called *optimal strategies*, for the two players respectively with the following property: When player I plays $\mathbf{p}_0$ then his expected winnings are at least V whatever player II does and when player II plays $\mathbf{q}_0$ then player I's expected winnings are at most V whatever he does. More formally $V = \mathbf{p}_0^T A \mathbf{q}_0$, $\mathbf{p}_0^T A \mathbf{q} \geq V$ for every $\mathbf{q}$ and $\mathbf{p}^T A \mathbf{q}_0 \leq V$ for every $\mathbf{p}$. Another way to describe $\mathbf{p}_0$ and $\mathbf{q}_0$ are that $\mathbf{p}_0$ maximizes $\min_j C_j(\mathbf{p})$ over all possible mixed strategies $\mathbf{p}$ and $\mathbf{q}_0$

minimizes $\max_i R_i(\mathbf{q})$ over all possible $\mathbf{q}$'s, where $C_j(\mathbf{p}) = \sum_{i=1}^m p_i a_{ij}$, the expected winnings for player I when player II plays column j, and analogously $R_i(\mathbf{q}) = \sum_{j=1}^n q_j a_{ij}$.

In this paper we will study the optimal strategies and the value when the game matrix A is a random $n \times n$-matrix $\mathcal{X}$ where the entries are iid standard normal. We know from the Minimax Theorem that given the realization of $\mathcal{X}$ there will be optimal strategies for the two players and we shall mainly study the number of rows (columns) that will get a positive probability in an optimal strategy for player I (player II).

References to random games are sparse. The papers [1], [2], [3], [4], [6] and [7] are concerned with situations which are "similar in spirit" to our setting. However most of them treat random general-sum games, a setting which is in fact in a sense is simpler than the zero-sum case because of the often assumed independence between the payoffs for different players. Among the mentioned papers only Cover [3] and Faris-Maier [4] consider the present setting and of these only the latter is directly concerned with the present questions. Faris and Maier show that the probability that an optimal strategy has full support is bounded from above by $2^{-(n-1)}$ and make simulations that suggest that the number of rows/colums in the support is asymptotically close to a binomial distribution with parameters n and $1/2$. In Section 2 we take a few more theoretically rigorous steps towards the limiting distrubtion by showing on one hand that with probability tending to 1 this number is between $01.n$ and $0.9n$ and on the other hand that its expectation is $(1 + o(1)\frac{1}{2}$. We will also provide a high probability upper bound on the value of the game.

2 Random games

From now on we shall play the game on the $n \times n$-matrix $\mathcal{X} = [X_{ij}]_{1 \le i,j \le n}$ where the X_{ij}'s are iid standard normal random variables. (The assumption of normal distribution is mostly for convenience and most of what we do can equally well be done for any continuous distribution symmetric about 0.) Given any realization there will be an optimal strategy $\mathbf{p}$ for player I and an optimal strategy $\mathbf{q}$ for player II. The following is a simple but important observation:

LEMMA 2.1 *With probability 1 the optimal strategies are unique and their supports are of the same size.*

Proof. Note first that since the distribution of the X_{ij}'s is continuous, every sub-matrix of $\mathcal{X}$ is a.s. non-singular. Now pick any subset of the rows, let k denote the size of the subset and let $\mathbf{p}$ be any strategy for player I having that subset as its support. Since the $k \times n$-sub-matrix on these rows is non-singular no more than k of the $C_j(\mathbf{p})$'s can be equal. Consequently $C_j(\mathbf{p}) = V$ for at most k different j's; let $\mathcal{I}(\mathbf{p})$ denote the set of these j's. The reasoning applies in particular when $\mathbf{p}$ is an optimal strategy and the rows picked are those of its support. But then any optimal strategy for player II must have a support which is contained in $\mathcal{I}(\mathbf{p})$; otherwise it would not be optimal. Indeed the support of an optimal strategy for player II must be contained in the intersection of the $\mathcal{I}(\mathbf{p})$'s over all optimal $\mathbf{p}$. Taking the symmetry of the situation into account it now follows that all optimal strategies for player I have the same support, that the same goes for player II and that the two supports are of the same size, k. Finally uniqueness follows from non-singularity of the $k \times k$-matrix on the rows and columns of these supports since if $\mathbf{p}$ and $\mathbf{p}'$ are two optimal strategies for player I then $C_j(\mathbf{p}) = C_j(\mathbf{p}')$ for all j's of this sub-matrix and so $\mathbf{p} = \mathbf{p}'$. □

By Lemma 2.1 both players will have optimal strategies supported on the same number of rows or columns. Denote this random quantity by Z. I.e.

$$Z = |\{i \in [n] : p_i > 0\}| = |\{j \in [n] : q_j > 0\}|$$

where $\mathbf{p}$ and $\mathbf{q}$ are the unique optimal strategies for player I and player II respectively. Also given that $Z = k$ the whole game will take place on the corresponding $k \times k$-sub-matrix on which $C_j(\mathbf{p}) = R_i(\mathbf{q}) = V$ for all rows i and columns j. In particular the $C_j(\mathbf{p})$'s are equal for all j on this sub-matrix as are the $R_i(\mathbf{q})$'s for all i. Denote the random set of rows on which $\mathbf{p}$ has its support by $\mathcal{G}_I$ and the set of columns on which $\mathbf{q}$ has its support by $\mathcal{G}_{II}$ and put $\mathcal{G} = \mathcal{G}_I \times \mathcal{G}_{II}$ We will work on bounding the probability that $\mathcal{G}$ equals a beforehand specified $k \times k$ sub-array $B = B_I \times B_{II}$ of $[n] \times [n]$. A consequence of what we just stated is that a necessary condition on B for the event $\{\mathcal{G} = B\}$ is

(E) There are mixed strategies $\mathbf{p}$ and $\mathbf{q}$ for the two-person zero-sum game played on $\mathcal{X}(B)$ such that $C_i(\mathbf{p}) = C_j(\mathbf{p})$ and $R_i(\mathbf{q}) = R_j(\mathbf{q})$ for all $(i,j) \in B$.

The two strategies in this condition are called *equalizing strategies* (for B and for player I and player II respectively). It takes more than (E) for B to be $\mathcal{G}$, namely that the two equalizing strategies are optimal on the whole of $\mathcal{X}$, i.e. that the C_j's and R_i's are all not just equal but actually equal to V. Now saying that $\mathbf{p}$ and $\mathbf{q}$ are optimal on $\mathcal{X}$ is the same as saying on one hand that they are optimal on $\mathcal{X}(B)$ and on the other hand that there is no other square sub-matrix of $\mathcal{X}$ for which any of the players can find a better strategy. The first part is covered by (E) so consider now the latter part. Assume without loss of generality that B is the array $[k] \times [k]$, i.e. corresponding to the first k rows and columns. Denote by V_B the value of the game on $\mathcal{X}(B)$ and let $\mathbf{p}_B$ and $\mathbf{q}_B$ denote the corresponding optimal strategies. We want a simple criterion for determining when $V_B = V$. Create an $(n-k+1) \times (n-k+1)$-matrix $S_B = [s_{ij}]$ by putting $s_{ij} = X_{k-1+i,k-1+j}$ when $i \geq 2$ and $j \geq 2$, $s_{11} = V_B$, $s_{1j} = \sum_{i=1}^{k} p_B(i) X_{i,k-1+j}$ when $j \geq 2$ and $s_{i1} = \sum_{j=1}^{k} q_B(j) X_{k-1+i,j}$ when $i \geq 2$.

LEMMA 2.2 *One has $V_B = V$ if and only if s_{11} is a saddle point of S_B.*

Proof. Assume first that $V_B = V$. Then $\mathbf{p}_B$ and $\mathbf{q}_B$ are optimal strategies on the whole of $\mathcal{X}$ (with $\mathbf{p}_B(i) = \mathbf{q}_B(i)$ regarded as 0 for $i > k$). Thus since the entries $2, 3, \ldots, n-k+1$ of the first row of S_B are $C_{k+1}(\mathbf{p}_B), C_{k+2}(\mathbf{p}_B), \ldots, C_n(\mathbf{p}_B)$ and the entries $2, 3, \ldots, n-k+1$ of the first column are $R_{k+1}(\mathbf{q}_B), R_{k+2}(\mathbf{q}_B), \ldots, R_n(\mathbf{q}_B)$ one has that $s_{11} = V_B$ must be smallest in the first row and greatest in the first column, i.e. a saddle point. On the other hand if s_{11} is a saddle point none of the players can do better than using $\mathbf{p}_B$ and $\mathbf{q}_B$ respectively and therefore $V_B = V$. $\square$

Thus in order for a square sub-array B of $[n] \times [n]$ to be $\mathcal{G}$ we can add the following condition:

(S) The entry s_{11} of S_B is a saddle point.

Let $\mathcal{E}_B$ be the event that a given $k \times k$ sub-array B satisfies (E) and denote by $\mathcal{S}_B$ the event that B satisfies (S). We have argued that $\{\mathcal{G} = B\} = \mathcal{E}_B \cap \mathcal{S}_B$ and so

$$P(\mathcal{G} = B) = P(\mathcal{E}_B \cap \mathcal{S}_B).$$

Let us now bound the right hand side. Begin by conditioning on $\mathcal{X}(B)$:

$$P(\mathcal{E}_B \cap \mathcal{S}_B) = \mathbb{E}[P(\mathcal{E}_B \cap \mathcal{S}_B | \mathcal{X}(B))] = \mathbb{E}[1_{\mathcal{E}_B} P(\mathcal{S}_B | \mathcal{X}(B))],$$

the second equality following from the fact that $\mathcal{E}_B$ is $\mathcal{X}(B)$-measurable. Now $P(\mathcal{S}_B | \mathcal{X}(B))$ is the probability that $s_{12}, \ldots, s_{1,n-k+1}$ are all greater that V_B and $s_{21}, \ldots, s_{n-k+1,1}$ are all less than V_B. Since given any outcome of $\mathcal{X}(B)$ these entries constitute two sets of iid normal random variables with expectation 0 (but with different variances; $\|\mathbf{p}_B\|^2$ and $\|\mathbf{q}_B\|^2$ respectively) we see that, since V_B is either positive or negative, this conditional probability is at most $2^{-(n-k)}$. Therefore

$$P(\mathcal{E}_B \cap \mathcal{S}_B) \leq \frac{1}{2^{n-k}} P(\mathcal{E}_B) \tag{2.1}$$

and it remains to bound $P(\mathcal{E}_B)$. Doing this will take slightly more work: For any $k \times k$-matrix M with iid standard normal entries, put $\mathcal{E} = \mathcal{E}(M)$ for the event that M satisfies (E). Write $\mathcal{E} = \mathcal{R} \cap \mathcal{C}$ where $\mathcal{R}$ is the event that there exists a strategy $\mathbf{p}$ on M for player I such that $C_i(\mathbf{p}) = C_j(\mathbf{p})$ for all $i, j \in [k]$ and $\mathcal{C}$ is the event that there exists a strategy $\mathbf{q}$ on B for player II such that $R_i(\mathbf{q}) = R_j(\mathbf{q})$ for all i and j. Another way of expressing $\mathcal{C}$ is that the (almost surely) unique vector $\mathbf{x}$ such that $M\mathbf{x} = \alpha \mathbf{1}$ for some $\alpha \neq 0$ and $\mathbf{x}^T \mathbf{1} = 1$ (unique because M is a.s. non-singular) has only positive entries. This in turn is the same event as that the a.s. unique solution to $M\mathbf{x} = \mathbf{1}$ has either all positive or all negative entries. Since $\mathbf{x} = M^{-1}\mathbf{1}$, this is the same as saying that the row sums of M^{-1} are either all positive or all negative. Now since the distribution of M is invariant under (right or left) multiplication by any orthogonal matrix, then so is M^{-1}. More formally, if Q is a $k \times k$ orthogonal matrix, then

$$QM^{-1} = (MQ^T)^{-1} =_d M^{-1}$$

and
$$M^{-1}Q = (Q^T M)^{-1} =_d M^{-1}.$$

In particular the distribution of M^{-1} is invariant under the operation of changing the sign of all the entries in a given set of rows; this corresponds to taking Q to be diagonal and orthogonal (i.e. with all diagonal entries being 1 or -1) and multiplying M^{-1} from the left by Q. Thus

$$P(\mathcal{C}) = P(\mathcal{C}(M)) = P(\mathcal{C}(MQ^T))$$

for all such Q's. However there are always exactly two Q's for which $\mathcal{C}(MQ^T)$ occurs. Since there are exactly 2^k diagonal orthogonal matrices of dimension $k \times k$ we get that

$$P(\mathcal{C}) = \frac{1}{2^{k-1}}.$$

An analogous argument also tells us that $\mathcal{R}$ has the same probability. An immediate consequence of these two results is that $P(\mathcal{E}) \leq 2^{-(k-1)}$. Putting $M = \mathcal{X}(B)$ and taking (2.1) into account we have established:

PROPOSITION 2.3 *Let B be any $k \times k$ sub-matrix of $\mathcal{X}$. Then*

(a) $P(\mathcal{C}_B) = P(\mathcal{R}_B) = \frac{1}{2^{k-1}}$,

(b) $P(\mathcal{E}_B) \leq \frac{1}{2^{k-1}}$,

(c) $P(\mathcal{E}_B \cap \mathcal{S}_B) \leq \frac{1}{2^{n-1}}$.

Note that Faris and Maier's result follows from (b) with $B = \mathcal{X}$.
We now arrive at our main result:

THEOREM 2.4 *Let Z be defined as above. Then there exists $a < 1/2$ such that*

$$\lim_{n \to \infty} P\left(\left(\frac{1}{2} - a\right)n < Z < \left(\frac{1}{2} + a\right)n\right) = 1.$$

Proof. To be specific we will prove that one can take $a = 0.4$. By Proposition 2.3

$$P(\mathcal{G} = B) \leq \frac{1}{2^{n-1}}$$

for any $k \times k$ sub-array B of $[n] \times [n]$. The number of $k \times k$ sub-arrays is $\binom{n}{k}^2$ and so

$$P(Z = k) \leq \frac{\binom{n}{k}^2}{2^{2n-k-1}}.$$

By Stirling's formula

$$\binom{n}{k} \leq \frac{n^n}{(n-k)^{n-k} k^k}$$

and writing $k = bn$ thus yields

$$P(Z = k) \leq 2\left(2(b^b(1-b)^{1-b})^2\right)^{-n}.$$

Put $r(b) = 2(b^b(1-b)^{1-b})^2$, i.e. the expression in brackets on the right hand side. The function $r(b)$ is symmetric about $b = 1/2$, increasing on $[0, 1/2)$ and decreasing on $(1/2, 1]$. Since $r(0.1) = r(0.9) < 1$ we get as $n \to \infty$ that

$$P\left(0.1n \leq Z \leq 0.9n\right) \leq 0.4n(r(0.1n))^n \to 0$$

which completes the proof. $\square$

Remark. As indicated earlier, nothing of what we have done relies on the fact that the X_{ij}'s are normal. All that matters is that the X_{ij}'s are iid with some continuous distribution symmetric

about 0. In this more general case the distribution of $\mathcal{X}$ is no longer invariant under arbitrary orthogonal transformations. However invariance remains for the very special cases of orthogonal transformations that we use.

So the bounds of Proposition 2.3 allow us to conclude that with high probability the support of the optimal strategy will be between $0.1n$ and $0.9n$. However it is difficult from reading the arguments not to note that we "wasted" a lot of information: E.g. when bounding $P(\mathcal{E})$ we only used the obvious bound $P(\mathcal{E}) = P(\mathcal{R} \cap \mathcal{C}) \leq P(\mathcal{C})$ and it seems likely that $P(\mathcal{E})$ should be of much smaller order than $P(\mathcal{R})$ and $P(\mathcal{C})$. In fact we believe that $\mathcal{C}$ and $\mathcal{R}$ are slightly positively correlated but asymptotically not far from independent. (Independence in a strict sense does not hold, for example in the case $k = 2$ one has $P(\mathcal{C}) = P(\mathcal{R}) = 1/2$ whereas $P(\mathcal{E}) = 1/3$.) Simulations and heuristics support the following conjecture:

CONJECTURE 2.5 *Let M be a $k \times k$ matrix whose entries are iid standard normal. Then*

$$\frac{1}{4^{k-1}} \leq P(\mathcal{E}_M) \leq \frac{2}{4^{k-1}}.$$

Despite considerable effort we have not been able to prove Conjecture 2.5 which we think is interesting in its own right. However the following result is very similar and may serve as strong evidence for the conjecture: Recall that $\mathcal{R}$ and $\mathcal{C}$ are the events that the entries of $(M^{-1})^T \mathbf{1}$ and $M^{-1}\mathbf{1}$ respectively are either all positive or all negative. Replace the vector $\mathbf{1}$ by $\mathbf{e}_i$, the i'th base vector of $\mathbb{R}^k$ and denote the corresponding new events by $\mathcal{R}^i$ and $\mathcal{C}^i$. Then:

THEOREM 2.6 *The events $\mathcal{R}^i$ and $\mathcal{C}^j$ are independent for all $i, j \in [k]$.*

Proof. The event $\mathcal{R}^i$ is the event that the entries of the i'th column of M^{-1} all have the same sign and $\mathcal{C}^j$ is the event that entries of the j'th row of M^{-1} all have the same sign. Thus $\mathcal{R}^i \cap \mathcal{C}^j$ is the event that all the entries of the $2k - 1$-dimensional vector

$$\mathbf{b} := (\beta_{1i}, \ldots, \beta_{j-1,i}, \beta_{j+1,i}, \ldots, \beta_{ki}, \beta_{ji}, \beta_{j1}, \ldots, \beta_{j,i-1}, \beta_{j,i+1}, \ldots, \beta_{jk}),$$

where the β's are the entries of M^{-1}, have the same sign. Since the distribution of M^{-1} is invariant under orthogonal transformations, the distribution of $\mathbf{b}$ is invariant under orthogonal transformations that leave either the first k entries or the last k entries fixed. This implies that the distribution of $\mathbf{b}$ is invariant under changing the sign of any given subset of its entries. Since there are exactly 2^{2k-1} ways of doing this and since all these correspond to disjoint transformations of $\mathcal{R}^i \cap \mathcal{C}^j$ of which exactly two always occur (as in the proof of $P(\mathcal{C}) = 2^{-(k-1)}$), we get that

$$P(\mathcal{R}^i \cap \mathcal{C}^j) = \frac{2}{2^{2k-1}} = \frac{1}{4^{k-1}}$$

as desired. $\square$

Assuming that Conjecture 2.5 holds would produce a considerably smaller a in Theorem 2.4, but it would not bring a arbitrarily close to 0. However what we believe, and what is supported by Faris and Maier's simulation as well as our own, is that in fact any $a > 0$ will do:

CONJECTURE 2.7 *For any $a > 0$*

$$\lim_{n \to \infty} P\left((\frac{1}{2} - a)n < Z < (\frac{1}{2} + a)n\right) = 1.$$

To shed further light on the distribution of Z and give furher evidence of Conjecture 2.7, we next calculate $\mathbb{E}Z$:

THEOREM 2.8

$$\mathbb{E}Z = (1 + o(1))\frac{n}{2}.$$

Proof. Lemma 2.2 easily generalizes to cases when B is not square, so the event that a particular row, say row 1, is in the support of player I's optimal strategy is the event that $V_1 := \sum_{j=1}^n q'_j X_{1j} > V'$ where V' is the value of the game on the matrix $\mathcal{X}'$ obtained from $\mathcal{X}$ by removing row 1 and $\mathbf{q}'$ is the optimal strategy for player II on $\mathcal{X}'$. Clearly $V' \leq V$ and the distribution of V_1 given $\mathcal{X}'$ is ymmetric about 0, so the probability that row 1 is in the support is at least $1/2$. We must show that this probability is also not essentially bigger than $1/2$. We do this by showing that the expected loss caused to player I from the removal of row 1 is small compared to the variance $\|\mathbf{q}'\|^2$ of V_1 given $\mathcal{X}'$. Let player I use a strategy $\mathbf{p}'$ on $\mathcal{X}'$ given by putting $p'_1 = 0$ (obviously) and $p'_i = p_i + p_1/(n-1)$ where $\mathbf{p}$ is the optimal strategy on the whole of $\mathcal{X}$. (To be correct one must set p'_i to something else in cases when $p_i + p_1/(n-1) > 1$. However it is easily seen that this will with probability tending to 1 not happen: Playing on at most two rows would cause player I a loss of at least order $O(1)$.) Then

$$
\begin{aligned}
V' &\geq \min_j \sum_{i=2}^n p'_i X_{ij} = \min_j \sum_{i=2}^n (p_i + \frac{1}{n-1}) X_{ij} \\
&= \min_j \left(\sum_{i=1}^n p_i X_{ij} - p_1 X_{1j} + \frac{p_1}{n-1} \sum_{i=2}^n X_{ij} \right) \\
&\geq V - \frac{p_1}{n-1} \max_j |\sum_{i=2}^n X_{ij}| - p_1 \max_j X_{1j}.
\end{aligned}
$$

The sums in the second term are normal with mean 0 and variance $n-1$ and so the expectation of the maximum does not exceed $(n \log n)^{1/2}$. Since the maximum is not correlated with p_1 and $\mathbb{E}p_1$ obviously equals n^{-1}, the expectation of the second term does not exceed $n^{-3/2}(\log n)^{1/2}$. Similarly the expectation of the third term is bounded by $n^{-1}(\log n)^{1/2}$ and so we can write $V' \geq V - Y$ where Y is a positive random variable with expectation $O(n^{-1}(\log n)^{1/2})$. By Markov's inequality $P(Y \leq n^{-1} \log n) \to 1$ and so with probability tending to 1, $V' \geq V - n^{-1} \log n$. Since $\|\mathbf{q}'\|^2$ is bounded from below by n^{-1}, the standard deviation of V_1 given $\mathcal{X}'$ is at least $n^{-1/2}$ and so

$$
P(V_1 \leq V - n^{-1} \log n | V_1 \leq V) \to 1
$$

from which it follows that

$$
P(V_1 \geq V') = (1 + o(1))\frac{1}{2}
$$

as desired. $\square$

Note that Theorem 2.8 together would together with Chebyshev's inequality imply Conjecture 2.7 if it could also be established that the variance of Z is $o(n^2)$. Unfortunately we have not been able to bound the variance of Z in any useful way. To do that we would e.g. need:

CONJECTURE 2.9 *Let $f(n)$ be any sequence such that $f(n) \to \infty$. Then*

$$
P(|V| \leq f(n)n^{-1}) \to 1
$$

as $n \to \infty$.

Since it is very natural to ask for results on the value of a random game, results in the spirit of Conjecture 2.9 are interesting in their own right and not only as instruments for proving results on the support of the optimal strategy. In order to motivate our belief in Conjecture 2.9 and to prove a weaker statement, consider the equalizing strategy for, say, player II on an $n \times n$-matrix M when such exists. In other words consider the a.s. unique solution to $M\mathbf{q} = \alpha\mathbf{1}$ with $\sum_i q_i = 1$ on the event $\mathcal{C}$ that $\mathbf{q} \in \mathbb{R}^n_+$. The solution is

$$
\mathbf{q} = \frac{M^{-1}\mathbf{1}}{\mathbf{1}^T M^{-1}\mathbf{1}}.
$$

Denote the denominator by W^{-1}. Note that on the event $\mathcal{R} \cap \mathcal{C}$ that there are equalizing strategies for both players, W is the value of the game on M. We will show that with probability tending to 1 given $\mathcal{C}$, $|W|$ will be of order $O(f(n)n^{-1})$ for any $f(n) \to \infty$.

To make things a little easier we recall from above that the distribution of M^{-1} is invariant under orthogonal transformations. Thus multiplying M^{-1} from the right by an orthogonal matrix whose first column is $n^{-1/2}\mathbf{1}$ does not change its distribution. Therefore W^{-1} is distributed as $n^{1/2}$ times the sum of the elements of the first column of M^{-1} and the event $\mathbf{q} \in \mathbb{R}_+^n$ transforms into the event that the first column of M^{-1} contains only positive or only negative elements.

Now the first column of M^{-1} can be written as $\mathbf{n}/Y$ where $\mathbf{n}$ is a unit vector orthogonal to all but the first row of M and Y is the scalar product of the first row of M with $\mathbf{n}$. (To be specific, choose $\mathbf{n}$ by a fair coin flip from the two possible vectors.) Note that Y is standard normal and independent of the direction of $\mathbf{n}$. We have that $|W|^{-1}$ is distributed as $n^{1/2}|\mathbf{1}^T\mathbf{n}|/|Y|$ and on the event that all entries of $\mathbf{n}$ have the same sign we claim that for any $b > 0$ with probability tending to 1

$$(\frac{1}{\sqrt{2}} - b)n^{1/2} \le |\mathbf{1}^T\mathbf{n}| \le (\frac{1}{\sqrt{2}} + b)n^{1/2}$$

which entails that $|W|$ is of order $f(n)n^{-1}$. To show the claim we use that $\mathbf{n}$ has the distribution of $X/\|X\|$ where $X = (X_1, X_2, \ldots, X_n)$ with iid standard normal entries. The expectation of $|X_i|$ is $\sqrt{2/\pi}$ and by standard large deviation theory $\sum_{i=1}^n |X_i|$ does with high probability not significantly deviate from $n\sqrt{2/\pi}$. Similarly, the expectation of X_i^2 is 1 and so $\|X\|^2$ is with high probability close in the same sense to n so that $\|X\|$ is close to $n^{1/2}$. Since conditioning in that $X \in \mathbb{R}_+^n$ does not change the distribution of the $|X_i|$'s, the claim is proved. We have shown:

THEOREM 2.10 *For any $b > 0$ and any $f(n) \to \infty$,*

$$P(|W| \le f(n)n^{-1}|\mathcal{C}) \to 1$$

as $n \to \infty$.

In order to prove that Conjecture 2.9, we would need to prove that the statement of Theorem 2.10 holds when M is a sub-matrix of $\mathcal{X}$ and when we also condition on $\mathcal{R}(M)$ and $\mathcal{S}(M)$. Heuristically both of these events tend to decrease $|W|$ even further, but we have not been able to find a way to handle the intricate dependencies between the entries of $\mathcal{X}$ that arise from conditioning on all three events. Simulations seem to indicate that $|V|$ is of even smaller order than the one given in Conjecture 2.9. In any case, by using brute force Theorem 2.10 at least gives a nontrivial upper bound (which to the best of our knowledge is the best known bound) on the value of the game: The random variable $Y =: Y(M)$ that appears in the proof is standard normal and since there are less than 4^n square sub-matrices M of $\mathcal{X}$, there is a constant D such that $\max_M Y(M)$ does with probability tending to 1 not exceed $Dn^{1/2}$. Thus $Y_\mathcal{G} \le Dn^{1/2}$. Also there exists $d > 0$ such that $|\mathbf{1}^T\mathbf{n}| \ge dn$ for all B and therefore also for $\mathcal{G}$. Therefore, putting $C = D/d$,

THEOREM 2.11 *There exists $C < \infty$ such that with probability tending to 1*

$$|V| \le Cn^{-1/2}.$$

REFERENCES

1. E. BALDER, On the existence of maximum likelihood Nash equilibria. Stochastic equilibrium problems in economics and game theory, *Ann. Oper. Res.* **114** (2002), 57-70.

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

3. T. COVER, The probability that a random game is unfair, *Ann. Math. Statist.* **37** (1966), 1796-1799.

4. W. G. FARIS AND R. S. MAIER, The value of a random game; the advantage of rationality, *Complex Systems* **1** (1987), 235-244.

5. G. OWEN, "Game Theory," 3rd edition, Academic Press, San Diego, 2001.

6. Y, RINOTT AND M. SCARSINI, On the number of pure strategy Nash equilibria in random games, *Games Econom. Behav.* **33** (2000), 274-293.

7. W. STANFORD, A note on the probability of k pure Nash equilibria in matrix games, *Games Econom. Behav.* **9** (1995), 238-246.

Johan Jonasson
Dept. of Mathematics
Chalmers University of Technology
S-412 96 Sweden
Phone: +46 31 772 35 46
`jonasson@math.chalmers.se`