Nash Equilibria in Random Games

Imre Bárány,[1,2,*] Santosh Vempala,[3,†] Adrian Vetta[4,‡]

[1]*Rényi Institute of Mathematics, Hungarian Academy of Sciences, Hungary;*
e-mail: barany@renyi.hu
[2]*Department of Mathematics, University College London, UK*
[3]*School of Computer Science, Georgia Tech, Atlanta, GA, USA;*
e-mail: vempala@cc.gatech.edu
[4]*Department of Mathematics and Statistics, School of Computer Science,*
McGill University, Quebec, Canada; e-mail: vetta@math.mcgill.ca

ABSTRACT: We consider Nash equilibria in 2-player random games and analyze a simple Las Vegas algorithm for finding an equilibrium. The algorithm is combinatorial and always finds a Nash equilibrium; on $m \times n$ payoff matrices, it runs in time $O(m^2 n \log \log n + n^2 m \log \log m)$ with high probability. Our result follows from showing that a 2-player random game has a Nash equilibrium with supports of size two with high probability, at least $1 - O(1/\log n)$. Our main tool is a polytope formulation of equilibria. © 2007 Wiley Periodicals, Inc. Random Struct. Alg., 31, 391–405, 2007

Keywords: Nash equilibria; random games; polynomial time; algorithmic game theory

1. INTRODUCTION

In a 2-player game, the first player, namely Alice, has m pure strategies $\mathcal{S} = \{\sigma_1, \sigma_2 \ldots, \sigma_m\}$ while the second player, namely Bob, has n pure strategies $\mathcal{T} = \{\tau_1, \tau_2 \ldots, \tau_n\}$. We are given two payoff matrices A and B for Alice and Bob, respectively. Here, the ijth entry of A is the payoff to Alice when she plays σ_i and Bob plays τ_j. A Nash equilibrium is a pair of mixed strategies (probability distributions) ν, λ such that given ν, the distribution λ on $\mathcal{T}$

---

maximizes $\nu^T B \lambda$, the payoff to Bob, and simultaneously, given λ, the distribution ν on $\mathcal{S}$ maximizes $\nu^T A \lambda$, the payoff to Alice.

The complexity of finding a Nash equilibrium in a game is a pivotal question in algorithmic game theory [22]. For general games this question was answered recently when Chen and Deng [6] proved that the Nash equilibrium problem in 2-player games is PPAD-complete. Determining the complexity of finding equilibria has also lead to much research in a variety of directions, e.g., a quasi-polynomial-time algorithm to find an approximate Nash equilibrium due to Lipton et al. [19]; an investigation into the complexity of finding pure strategy Nash equilibria in succinctly specified games by Fabrikant et al. [12]; a proof that the Lemke-Howson algorithm [18] takes exponential time with all possible initial pivots [26].

In this article, we consider 2-player games where the two payoff matrices are chosen randomly. Our motivation is the question of whether finding Nash equilibria is any easier in random games compared to general games, that is, easier "on average." In a random game, every entry in each of the matrices is drawn independently according to some probability distribution. We consider the uniform distribution on an interval and the standard Normal distribution $N(0, 1)$. In the first case, the distribution of any set of k entries of a payoff matrix is uniform in a k-dimensional cube, while in the second case it is a k-dimensional Normal. In fact, there has been much work in this direction for a special case of 2-player games, the zero-sum case. This case is equivalent to linear programming. Motivated by the question of explaining the success of the simplex algorithm, Borgwardt [5], Smale [28], and Megiddo [21] studied linear programs where the constraints are chosen randomly from spherically symmetric distributions and showed that variants of the simplex algorithm run in polynomial time. Besides simplex, other simple methods (e.g., the perceptron algorithm) also work for random linear programs, demonstrating that they have considerably more structure than arbitrary linear programs.

Here we show random games are indeed much simpler than general games. Specifically, we show that with high probability, there is a Nash equilibrium in which the supports of the mixed strategies of both players have small cardinality.

In a random game, the vectors of the payoff matrices define points in general position with probability one. This implies the well-known property that in a Nash equilibrium the supports of both players have the same cardinality, which we call the "size of the equilibrium support." As a result, the following naive heuristic is a Las Vegas algorithm for finding Nash equilibria: exhaustively check for Nash equilibria with support of cardinality $i = 1, 2, \ldots$ until an equilibrium is found. In fact, see Corollary 12, with high probability only two phases will be required! (Interestingly, this simple heuristic has also recently been shown to perform extremely well experimentally, see [23].)

The key to our result is a reformulation of the problem in terms of random polytopes. For convenience, we will assume that $m = n$. We will see in Section 2 that, given a mixed strategy for Alice (Bob), the supports of best response strategies for Bob (Alice) are precisely those supports that induce facets with non-negative normal vectors in an associated random polytope. Consequently, the algorithmic problem of finding Nash equilibria can be tackled by considering problems relating to the number of points on the convex hull of a set of n random points in d dimensions. In Section 5, we extend the analysis of such random polytopes for our purpose. This allows us to examine the quality of our algorithm in Section 4. In particular, consider the convex hull of n random points in d dimensions. Let N_1 denote the expected number of points that lie on the boundary of the convex hull (this will differ for the Normal and uniform distributions). Then, our main theorem can be stated as follows.

Theorem 1. *The probability that a random $n \times n$ game contains no Nash equilibria with support of size at most d is less than*

$$f(d) \left(\frac{1}{n} + \frac{1}{(N_1)^2} \right)$$

where $f(d)$ is a function of d alone.

It is known that N_1 is unbounded as a function of n for both the distributions we consider and so we get our algorithmic result as a corollary.

Corollary 2. *There is a combinatorial algorithm that finds a Nash equilibrium in a random $n \times n$ game and, with high probability, runs in time $O(n^3 \log \log n)$.*

In related work, Goldberg et al. [13] have studied the expected number of pure strategy Nash equilibria in random games; McLennan and Berg [20] considered the case of (mixed strategy) Nash equilibria in certain random games. Finally, we remark in Section 6 that if our work could be extended to allow matrix entries to have arbitrary means then a polynomial time randomized algorithm for finding approximate Nash equilibria in arbitrary games would follow.

2. A GEOMETRIC INTERPRETATION OF NASH EQUILIBRIA

In this section, we give a geometric characterisation of Nash equilibria in terms of polytopes. For any subset of Alice's strategies, $S \subseteq \mathcal{S}$, let $B^S = \{b_1^S, \ldots, b_n^S\}$ be the set of subcolumns of B induced by S, i.e., whose entries correspond to the rows of S. We will view the elements of B^S as points in $\mathbb{R}^d$, and denote by $\mathsf{conv}(B^S)$ the polytope corresponding to the convex hull of these points. Given a subset of Bob's strategies, $T \subseteq \mathcal{T}$, we denote by $B^S(T)$ the points in B^S corresponding to elements in T. Similarly A^T is the set of subrows of A induced by T. We define $\mathsf{conv}(A^T)$ and $A^T(S)$ in an analogous manner.

Key to the geometric interpretation is the following simple lemma.

Lemma 3. *A pure strategy τ_j of Bob is a best response to some mixed strategy ν with support S of Alice if and only if b_j^S maximizes $\nu \cdot b$ among all points $b \in B^S$.*

Proof. Any mixed strategy λ for Bob corresponds to a point $b(\lambda) \in \mathsf{conv}(B^S)$ where $y(\lambda) = \sum_{i=1}^n \lambda_i b_i^S$. The expected payoff to Bob of the mixed strategy λ is then just $\nu \cdot y$. ∎

Note that only non-negative vectors correspond to feasible mixed strategies. Consequently, our real interest lies in the anti-dominant of the convex hull; the anti-dominant of a polyhedron $K \subseteq \mathbb{R}^d$ is $\mathsf{ant}(K) = \{z \in \mathbb{R}^d : \exists x \in K, z \leq x\}$. For simplicity, we will write $\mathsf{ant}(B^S) = \mathsf{ant}(\mathsf{conv}(B^S))$. We may now give the desired characterization of Nash Equilibria.

Theorem 4. *The pair $S \subseteq \mathcal{S}, T \subseteq \mathcal{T}$ produces a Nash equilibrium if and only if $B^S(T)$ induces a face of $\mathsf{ant}(B^S)$ and $A^T(S)$ induces a face of $\mathsf{ant}(A^T)$.*

Proof. Suppose the points in $B^S(T)$ induce a face in $\mathsf{ant}(B^S)$ contained in a facet with normal $\nu \neq 0$. Then, by Lemma 3, each point in $B^S(T)$ is a best response for Bob when

Alice is using the mixed strategy $\nu/||\nu||_1$ on S. Similarly, suppose the points in $A^T(S)$ induce a face in $\mathsf{ant}(A^T)$ with normal $\lambda \neq 0$. Then, each point in $A^T(S)$ is a best response for Alice is using the mixed strategy $\lambda/||\lambda||_1$ on T. Therefore, the use of the strategies S and T with probability distributions $\nu/||\nu||_1$ and $\lambda/||\lambda||_1$, respectively, gives a Nash equilibrium. Conversely, suppose that the points in $B^S(T)$ do not induce a face in $\mathsf{ant}(B^S)$. Then, it is not possible to find a mixed strategy ν on S such that every point in $B^S(T)$ is a best response for Bob when Alice is using the mixed strategy ν. Thus, no mixed strategy on T for Bob can be a best response against any mixed strategy for Alice on S. ∎

This characterization applies to general games. In our case, it implies that Nash equilibria in random games are closely related to the polytopes produced by random points. In particular, we need to study the faces of random polytopes. Since the entries of A and B are chosen at random, with probability one the corresponding point sets will be in general position and any face of dimension d contains exactly $d + 1$ points. It follows that for any Nash equilibrium in a random game we must have $|S| = |T|$.

3. EXPECTED NUMBER OF NASH EQUILIBRIA

For motivation, let's first consider the simple case of pure strategy Nash equilibrium. Observe that any pure strategy σ_i for Alice is a best response to any pure strategy τ_i for Bob with probability exactly $\frac{1}{n}$ and vice versa. It follows that the expected number of pure strategy Nash equilibria in a random game is exactly one. However, it is easy to show using basic probabilistic techniques that there are no pure strategy Nash equilibria with constant probability (in fact, this probability tends to $\frac{1}{e}$ [13] even for multiplayer games [9]).

Is the probability of success better when we look for mixed strategy Nash equilibria? To answer this question we will use the polyhedral framework developed in Section 2. We will begin by showing how to use this framework to easily count the expected number of mixed strategy Nash equilibria in a random game. Towards this end, let N_i be the expected number of faces of dimension $i - 1$ induced by $\mathsf{conv}(\mathcal{P})$ where $\mathcal{P}$ is a set of n random points (N_1 is the number of vertices, N_d is the number of facets). We will say that a facet of the polytope is useful if its normal vector is positive. Similarly a face of any dimension is called useful if it is contained in some useful facet. Our interest is in the expected number of useful faces of dimension $i - 1$, denoted N_i^+, since such faces are also faces of $\mathsf{ant}(\mathcal{P})$.

We are now ready to calculate the expected number of Nash equilibria. Again, let S and T be strategy subsets for Alice and Bob, respectively, where $|S| = |T| = d$. We let $\mathcal{F}(T) = \{S_1, S_2, \ldots, S_p\}$ be the set of facets of the polyhedron $\mathsf{ant}(A^T)$ (observe that $\mathsf{E}(p)$, the expectation of p, is just N_d^+); each S_i corresponds to the set of rows that induce the facet. We also define $\bar{\mathcal{F}}(T)$ to be the set of all faces contained in facets of $\mathcal{F}(T)$. To avoid confusion between row and column vectors, the sets $\mathcal{G}(S)$ and $\bar{\mathcal{G}}(S)$ are defined similarly w.r.t. the polyhedron $\mathsf{ant}(B^S)$. Then $S \in \mathcal{F}(T)$ iff Alice's strategies induced by S are all best responses to some mixed strategy by Bob on the strategy set induced by T. Note that S and T induce a Nash equilibrium, denoted by $S \leftrightarrow T$, if and only if $S \in \mathcal{F}(T)$ and $T \in \mathcal{G}(S)$. We denote by $\mathcal{E}_{ST}$ the event that $S \leftrightarrow T$, and by χ_{ST} the indicator variable for this event. By the independence of the payoff matrices for Alice and Bob, the probability that $S \leftrightarrow T$ is exactly the product of the probabilities that $S \in \mathcal{F}(T)$ and $T \in \mathcal{G}(S)$.

Lemma 5. *The expected number of $d \times d$ Nash equilibria in a random $n \times n$ game is*

$$\mu = \mathsf{E}\left(\sum_{S,T:|S|=|T|=d} \chi_{ST} \right) = \left(N_d^+ \right)^2.$$

Proof. Given S let us first evaluate the probability that $T \in \mathcal{G}(S)$ for some T. We have seen that this is the case if and only if $B^S(T)$ induces a useful facet in B^S. Thus, this probability is exactly

$$\frac{N_d^+}{\binom{n}{d}}.$$

Similarly, we also have that

$$\mathsf{P}(S \in \mathcal{F}(T)) = \frac{N_d^+}{\binom{n}{d}}.$$

By the independence of the payoff matrices, we obtain $\mathsf{P}(S \leftrightarrow T) = \left(\frac{N_d^+}{\binom{n}{d}} \right)^2$. Summing up over all pairs S and T, the lemma follows. ∎

We will use the following notation. We write $\alpha \lesssim \beta$ if there is a fixed function $f(d)$, of the dimension d alone, such that $\alpha \leq f(d)\beta$. We write $\alpha \approx \beta$ if $\alpha \lesssim \beta$ and $\alpha \gtrsim \beta$. Observe that, for any $1 \leq s \leq d$, $N_s^+ \gtrsim N_1^+ \gtrsim N_1/2^d$, where the second inequality follows from the observation that there are 2^d sign patterns for the normal of any face. Thus, by Lemma 5, we obtain a lower bound on the expected number of Nash equilibria by examining the behaviour of N_1. There has already been a large amount of work studying N_i for various distributions (for example, [1, 11, 14, 15, 30]); for a rather comprehensive survey see [31]. We begin with two basic results regarding the uniform and Normal distributions, respectively.

Theorem 6. [10] *Given n points drawn independently and uniformly from a d-dimensional unit cube, the expected number of points N_1 on the convex hull is $\gtrsim (\log n)^{d-1}$.*

Theorem 7. [24] *Given n points in $\mathbb{R}^d$ with coordinates drawn independently from $N(0,1)$, the expected number of points N_1 on the convex hull is $\gtrsim (\log n)^{\frac{1}{2}(d-1)}$.*

For completeness and to illustrate some of the techniques applied later, we will present a proof of Theorem 6. Our proof is based on the approach of Devroye [7] for counting maxima (undominated points) in random polytopes. We will also require the notion of a cap. Given a convex body $K \subseteq \mathbb{R}^d$, the intersection of K with a closed halfspace H is called a cap C. Intuitively, caps are useful in counting convex hull points in that a randomly chosen point is more likely to be on the convex hull if there is some cap of small volume that contains it.

Proof of Theorem 6. It suffices to give a lower bound on N_1^+. By symmetry, we just need to count the expected number of vertices on the dominant (rather than the anti-dominant) of the convex hull. Take a point $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_d) \in \mathbb{R}^d$. Note that α lies on the hyperplane

$$H = \left(x : \frac{x_1}{\alpha_1} + \frac{x_2}{\alpha_2} + \cdots + \frac{x_3}{\alpha_d} = d \right).$$

Now consider the cap C formed by intersecting the negative halfspace induced by H with the unit cube. Observe that C is contained within the simplex formed by the vertices $(0, \ldots, 0)$ and $(0, \ldots, d\alpha_i, \ldots, 0)$ for $1 \leq i \leq d$. This simplex has volume

$$\frac{\prod_{i=1}^{d}(d\alpha_i)}{d!} = \frac{d^d}{d!} \prod_{i=1}^{d} \alpha_i.$$

Denote the distribution of the ith coordinate by U_i. This is the uniform distribution on the unit interval. Since each coordinate is drawn independently we have that

$$N_1^+ \geq n \, \mathsf{E}\left(\left(1 - \frac{d^d}{d!}\prod_{i=1}^{d} U_i\right)^{n-1}\right)$$

$$= n \int_{x=0}^{\frac{d!}{d^d}} \left(1 - \frac{d^d}{d!}x\right)^{n-1} f(x)dx.$$

Here f is density function for the product of d independent uniform random variables. It is well known that that

$$f(x) = \frac{\log^{n-1}\left(\frac{1}{x}\right)}{\Gamma(d)}, \quad 0 < x < 1$$

where Γ is the gamma function. Hence, for small ϵ we have

$$N_1^+ \geq n \left(1 - \frac{d^d}{d!}\epsilon\right)^{n-1} \int_{x=0}^{\epsilon} f(x)dx$$

$$= n \left(1 - \frac{d^d}{d!}\epsilon\right)^{n-1} \epsilon \sum_{j=0}^{d-1} \frac{\left(\log\frac{1}{\epsilon}\right)^j}{j!}$$

$$\geq n \left(1 - \frac{d^d}{d!}\epsilon\right)^{n-1} \epsilon \frac{\left(\log\frac{1}{\epsilon}\right)^{d-1}}{(d-1)!}.$$

Setting $\epsilon = \frac{1}{n}$, we obtain

$$N_1^+ \geq \left(1 - \frac{d^d}{nd!}\right)^{n-1} \frac{(\log n)^{d-1}}{(d-1)!}$$

$$\gtrsim (\log n)^{d-1}. \qquad \blacksquare$$

We remark that similar techniques can be applied to give a matching upper bound on the expected number of vertices on the convex hull.

4. PERFORMANCE ANALYSIS OF THE LAS VEGAS ALGORITHM

Here we analyze the performance of aforementioned simple algorithm. To do this an expectation on the number of Nash equilibria is not sufficient; we also need to obtain concentration bounds. Recall that the algorithm exhaustively checks for Nash equilibria with support of size 1, then for Nash equilibria with support of size 2, etc., until we find a Nash equilibrium.

There are $\binom{n}{d}^2$ pairs of supports of size d and determining whether a pair of supports induce a Nash equilibrium can be done in polynomial time using the solution to a system of linear equations. Thus, provided that our game has a Nash equilibrium with a support of constant size, we obtain a polynomial time algorithm. We will show that with high probability, a random game has a 2×2 equilibrium. More generally, here we consider the probability that there is no $d \times d$ Nash equilibrium in a random game. So for a fixed d, we are interested in the probability that the random variable $Z = \sum_{S,T} \chi_{ST} > 0$. We will use the notation

$$\Delta = \sum_{(S,T),(S',T'):(S,T)\bowtie(S',T')} \mathsf{P}(\mathcal{E}_{ST} \wedge \mathcal{E}_{S'T'}),$$

where $\bowtie$ signifies that the events $\mathcal{E}_{ST}$ and $\mathcal{E}_{S'T'}$ are dependent. We remark that such a dependency arises if and only if $S \cap S' \neq \emptyset$ or $T \cap T' \neq \emptyset$. Our interest in Δ arises because, applying standard techniques [2], we obtain:

$$\mathsf{P}(Z = 0) \leq \frac{Var(Z)}{\mathsf{E}(Z)^2} \leq \frac{\mathsf{E}(Z) + \Delta}{\mathsf{E}(Z)^2}. \tag{1}$$

The next lemma bounds Δ.

Lemma 8. *In a random game, with μ as in Lemma 5, we have*

$$\Delta \lesssim \mu^2 \left(\frac{1}{n} + \frac{1}{\left(N_1^+\right)^2} \right).$$

Proof. Let $X = S^* \cap S$ and $Y = T^* \cap T$, $s = |X|$ and $t = |Y|$. We remark that only the cardinalities of these intersections will be of consequence. Recall that $X \in \bar{\mathcal{F}}(T^*)$ means that the vertices corresponding to X form a useful face (i.e., are all best responses) in the game induced by the strategies of T^*. So

$$\Delta = \sum_{s,t} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s,|Y|=t}} \mathsf{P}(S \leftrightarrow T \text{ and } S^* \leftrightarrow T^*)$$

$$= \sum_{s,t} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s,|Y|=t}} \mathsf{P}(S \in \mathcal{F}(T) \text{ and } S^* \in \mathcal{F}(T^*)) \, \mathsf{P}(T \in \mathcal{G}(S) \text{ and } T^* \in G(S^*))$$

$$= 2\sum_{s\geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s,|Y|=0}} \mathsf{P}(S \in \mathcal{F}(T) \text{ and } S^* \in \mathcal{F}(T^*)) \, \mathsf{P}(T \in \mathcal{G}(S) \text{ and } T^* \in \mathcal{G}(S^*))$$

$$+ \sum_{s,t\geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s,|Y|=t}} \mathsf{P}(S \in \mathcal{F}(T) \text{ and } S^* \in \mathcal{F}(T^*)) \, \mathsf{P}(T \in \mathcal{G}(S) \text{ and } T^* \in \mathcal{G}(S^*))$$

$$= 2\Delta_1 + \Delta_2.$$

For clarity, we will defer the proofs of the following two claims until the next section. We will need these claims to bound the terms Δ_1 and Δ_2. ∎

Claim 9. *For n sufficiently larger than d, if $T \cap T^* = 0$ then*

$$\mathsf{P}(T \in \mathcal{G}(S) \wedge T^* \in \mathcal{G}(S^*)) \lesssim \left(\frac{N_d^+}{\binom{n}{d}}\right)^2.$$

Claim 10. *For n sufficiently larger than d,*

$$\mathsf{P}(T \in \mathcal{G}(S) \text{ and } T^* \in \mathcal{G}(S^*)) \lesssim \frac{\binom{n}{t}(N_d^+)^2}{\binom{n}{d}^2 N_t^+}.$$

Now we bound Δ_1 from above. Observe that since $t = 0$ the events $S \in \mathcal{F}(T)$ and $S^* \in \mathcal{G}(T^*)$ are independent. Therefore,

$$\mathsf{P}(S \in \mathcal{F}(T) \text{ and } S^* \in \mathcal{F}(T^*)) = \mathsf{P}(S \in \mathcal{F}(T))^2 = \left(\frac{N_d^+}{\binom{n}{d}}\right)^2.$$

Applying Claim 9, we have

$$\Delta_1 = \sum_{s \geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s, |Y|=0}} \mathsf{P}(S \in \mathcal{F}(T) \text{ and } S^* \in \mathcal{F}(T^*)) \, \mathsf{P}(T \in \mathcal{G}(S) \text{ and } T^* \in \mathcal{G}(S^*))$$

$$\lesssim \sum_{s \geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s, |Y|=0}} \frac{(N_d^+)^4}{\binom{n}{d}^4}.$$

However,

$$\frac{(N_d^+)^4}{\binom{n}{d}^4} \sum_{s:d \geq s \geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s, |Y|=0}} 1 = \frac{\mu^2}{\binom{n}{d}^4} \sum_{s:d \geq s \geq 1} \binom{n}{d}\binom{d}{s}\binom{n-d}{d-s}\binom{n}{d}\binom{n-d}{d} \lesssim \frac{\mu^2}{n}.$$

Thus,

$$\Delta_1 \lesssim \frac{\mu^2}{n}.$$

Next, we bound Δ_2. Applying Claim 2, we obtain

$$\Delta_2 = \sum_{s,t \geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s, |Y|=t}} \mathsf{P}(S \in \mathcal{F}(T) \text{ and } S^* \in \mathcal{F}(T^*)) \, \mathsf{P}(T \in \mathcal{G}(S) \text{ and } T^* \in \mathcal{G}(S^*))$$

$$\lesssim \sum_{s,t \geq 1} \sum_{\substack{(S,T),(S^*,T^*): \\ |X|=s, |Y|=t}} \frac{\binom{n}{s}\binom{n}{t}(N_d^+)^4}{\binom{n}{d}^4 N_s^+ N_t^+}.$$

To finish the proof, observe that we have the following simple relationship between the number of faces of a given dimension and the number of points on the convex hull:

Observation 11. *The number of faces of conv($\mathcal{P}$) of dimension s is at least $\frac{1}{s+1}\binom{d-1}{s}$ times the number of points on the convex hull.*

Proof.   Every vertex of the convex hull is contained in at least d facets. Take a vertex x and consider any d facets containing x. Each of these is a simplex of d vertices and induces a set of $\binom{d-1}{s}$ faces of dimension s that contain x. The union of these sets may contain at most d copies of any face. Moreover, as all the points are in general position, every s-dimensional face contains exactly $s + 1$ points. Thus, summing over all vertices, we obtain the result.

Hence $N_{s+1} \geq \frac{1}{s+1}\binom{d-1}{s}N_1$, and so we have

$$\mu^2 \sum_{s,t \geq 1} \frac{\binom{d}{s}\binom{d}{t}}{N_s^+ N_t^+} \lesssim \frac{\mu^2}{\left(N_1^+\right)^2}.$$

Thus

$$\Delta_2 \lesssim \frac{\mu^2}{\left(N_1^+\right)^2}.$$

and this completes the proof of Lemma 8.   ∎

Theorem 1 now follows from (1) and Lemma 8. Then, by Theorems 6 and 7, which imply that N_1^+ becomes much larger than $f(d)$ as n increases, we obtain our result for random games in which the payoff entries are either uniformly or Normally distributed. In fact, it is enough to look for 2×2 equilibria.

Corollary 12.   *A random $n \times n$ game with normally (respectively, uniformly) distributed payoffs contains no 2×2 Nash equilibrium with probability $\lesssim \frac{1}{\log n}$ (respectively, $\lesssim \frac{1}{\log^2 n}$).*

It follows from Corollary 12 that the run time of the algorithm is $O(m^2 n \log \log n + n^2 m \log \log m)$, with high probability. To see this, observe that we need to calculate the convex hull of n points for each pair of strategies of Alice, and calculate the convex hull of m points for each pair of strategies of Bob. Since we can find the convex hull of k points in 2 dimensions in time $O(k \log h)$, where h is the number of points on the convex hull [17], this takes time $O(m^2 n \log \log n + n^2 m \log \log m)$. If a pair of strategies for Alice and a pair of strategies for Bob mutually induce facets with non-negative normals in the convex hull associated with the other pair then we have a Nash equilibrium. The normals to these facets also give the probability distributions on the strategy supports (at a Nash equilibrium). Thus our algorithm is entirely combinatorial.

5. RANDOM POLYTOPES UNDER GAUSSIAN AND UNIFORM DISTRIBUTIONS

In this section, we present proofs of the main technical results used in analysing the Las Vegas algorithm. Specifically, we present proofs of Claim 9 and Claim 10. To do so, we will be interested in the following general question. Let $\mathcal{P} = \{x_1, \ldots, x_n\}$ be a set of i.i.d. random points from a distribution with density function f. Let F be the set of subsets of points that induces facets of $\mathsf{conv}(\mathcal{P})$, and let Y_1 and Y_2 be two subsets of $\mathcal{P}$. The probability that one subset Y_1 induces a facet is well-understood for many important distributions, including the Gaussian and the cube. What, though, is the probability that both Y_1 and Y_2 induce facets, i.e., $\mathsf{P}(Y_1, Y_2 \in F)$? We prove the following results (which may be of independent interest).

Lemma 13. *Suppose f is the normal density or the uniform density over a cube. If Y_1, Y_2 are disjoint subsets of $\mathcal{P}$, then*

$$\mathsf{P}(Y_1, Y_2 \in F) \lesssim \mathsf{P}(Y_1 \in F)^2.$$

Lemma 14. *Suppose f is the normal density or the uniform density over a cube. If Y_1, Y_2 are subsets of $\mathcal{P}$ with $|Y_1 \cap Y_2| > 0$, then*

$$\mathsf{P}(Y_1, Y_2 \in F) \lesssim \frac{\mathsf{P}(Y_1 \in F)^2}{\mathsf{P}(Y_1 \cap Y_2 \text{ is a face of } \mathsf{conv}(\mathcal{P}))}.$$

The proof of these lemmas for the normal distribution can be carried out directly using the density function. We give such a proof of Lemma 13 at the end of this section (the proof of Lemma 14 is similar). For a cube, however, things are more complicated and we will use an economic cap-covering [4]. The next theorem is from [3].

Theorem 15. *Assume K is a convex body in $\mathbb{R}^d$, $\mathsf{vol}(K) = 1$ and $0 < \varepsilon < \varepsilon_0(d)$ where $\varepsilon_0(d)$ is a constant depending only on d. Let $K(\varepsilon)$ denote the set of points from K that can be cut off by a cap of volume ε. Then there are caps $C_1, \ldots, C_m$ and pairwise disjoint convex sets $C'_1, \ldots, C'_m$ such that $C'_i \subset C_i$ for every i such that every cap of volume ε is contained in one of the C_i and, further,*

$$\bigcup_1^m C'_i \subset K(\varepsilon) \subset \bigcup_1^m C_i$$
$$\mathsf{vol}(C'_i) \gtrsim \varepsilon \text{ and } \mathsf{vol}(C_i) \lesssim \varepsilon.$$

Simply put, Theorem 15 states that $K(\varepsilon)$ can be covered by caps of volume $\approx \varepsilon$ without much overlapping, that is, we have an economic cap-covering. Hence, the theorem implies that $m \approx \frac{1}{\varepsilon} \mathsf{vol}(K(\varepsilon))$. Moreover, if K_n is the random polytope on n points from K then $K - K(\frac{1}{n})$ is a good approximation to K_n [4]; of course, this is why the study of $K(\varepsilon)$ and, consequently, of cap coverings is of use here. It follows that, for m corresponding to $\varepsilon = 1/n$, the expected number of $k - 1$-dimensional faces of K_n is $N_k(K_n) \approx m$.

For the case in which K is the unit cube, we may take the sets C_i to be axis-parallel dyadic boxes instead of caps. In fact, the proof of Theorem 15 first finds an economic covering using Macbeath regions (boxes in the case of the unit cube) and then replaces them by closely fitting caps. In particular, the box-covering for $K(\varepsilon)$ consists of *all* boxes of the form

$$\prod_{j=1}^d [0, 2^{-f_j}]$$

where the f_j are non-negative integers summing to f, and $2^{-f} \approx \varepsilon$. (We remark that these boxes only form a covering near a single vertex, the origin, of the unit cube. As before, by symmetry, this is all that is required for the useful faces.)

One can define cap coverings for other distributions as well, replacing the volume of each set C_i by its measure. Such a covering exists for the normal distribution if the complement of a suitably large ball is deleted.

One can define similar cap or box coverings for other distributions as well, replacing the volume of each set C_i by its measure. Such a covering does exist for instance for the Normal distribution, and could be used to give alternative proofs of Lemma 13 and Lemma 14.

We will use Theorem 15 to prove the following generalization of Lemma 14. Again, we focus on the case when K is the unit cube, but the proof applies whenever a cap covering exists.

Lemma 16. *Let S, S^*, T, T^* be sets of cardinality d with $S \cap S^* = X$, $T \cap T^* = Y$, $|X| = s > 0$, $|Y| = t > 0$. We denote by $K = Q^{|S \cup S^*|} \subset \mathbb{R}^{2d-s}$ the unit cube in $\mathbb{R}^{2d-s}$. Assume $x_1, \ldots, x_n$ are uniform random points from K. Assume $T = \{x_1, \ldots, x_d\}$ and $T^* = \{x_1, \ldots, x_t, x_{d+1}, \ldots, x_{2d-t}\}$. Let $\mathcal{F}(S)$ denote the facets of the convex hull of the x_is projected onto $\mathbb{R}^S$, and similarly for $\mathcal{F}(S^*)$. Then,*

$$\mathsf{P}[T \in \mathcal{F}(S) \text{ and } T^* \in \mathcal{F}(S^*)] \lesssim \frac{(\log n)^{2d-s-1}}{n^{2d-t}}.$$

Proof. We take the box covering M_f for every large enough f for the unit cube. We write $L|S$ and $L|S^*$ for the projection of a set $L \subset \mathbb{R}^{|S \cup S^*|}$ onto $R^{|S|}$ and $R^{|S^*|}$, respectively. We need a special minimal cap $C(T)$ of K that contains T: namely, writing $C(T|S)$ for the minimal cap containing $T|S$ in $Q^{|S|}$ we let

$$C(T) = C(T|S) \times Q^{|S^* \setminus S|}.$$

Similarly,

$$C(T^*) = C(T^*|S^*) \times Q^{|S \setminus S^*|},$$

note that both $C(T)$ and $C(T^*)$ are caps of K. Now,

$$\mathsf{P}[T \in \mathcal{F}(S) \text{ and } T^* \in \mathcal{F}(S^*)] = \int_K \ldots \int_K \chi[T \in \mathcal{F}(S)]\, \chi[T^* \in \mathcal{F}(S^*)]\, \mathrm{d}x_1 \cdots \mathrm{d}x_n,$$

where $\chi[E]$ is the indicator of the event E. We can restrict integration to the subset where $\mathsf{vol}(C(T)) \leq \frac{c \log n}{n}$ for a fixed constant c and the same for $C(T^*)$, since, with probability $1 - n^{-5d}$, K_n contains all points of K that cannot be cut off from K by a cap of volume $\frac{c \log n}{n}$, if c is large enough. Define f_0 by $2^{-f_0} = \frac{c \log n}{n}$. We replace the above integral by a double sum in the following way. Given $x_1, \ldots, x_d$ and $C(T)$, let f be the largest integer with $C(T) \subset C_f$ for some $C_f \in M_f$ where M_f is the aforementioned box covering for $\varepsilon = 2^{-f}$. Of course, the C_f that matters is of the form

$$\prod_{j \in S \cup S^*} [0, 2^{-f_j}]$$

with $f_j = 0$ when $j \in S^* \setminus S$. Let M_f^1 be the set of these elements of M_f. Obviously, $\mathsf{vol}(C(T)) \geq c(d)2^{-f}$ with a suitable small $c(d) > 0$, and similarly for $C(T^*)$. Analogously, we get a C_g^* from M_g for each $C(T^*)$, and the ones that matter are collected in $M_g^2 \subset M_g$.

We integrate then on each $C_f \in M_f^1$ and $C_g^* \in M_g^2$ for $f, g \geq f_0$. Assuming $g \geq f$, the integrand is at most

$$(1 - c(d)2^{-f})^{n-(2d-t)}(\mathsf{vol}(C_f))^{d-t}\big(\mathsf{vol}(C_g^*)\big)^{d-t}\big(\mathsf{vol}(C_f \cap C_g^*)\big)^t$$

since $x_1, \ldots, x_t$ come from $C_f \cap C_g^*, x_{t+1}, \ldots, x_d$ from C_f, and $x_{d+1}, \ldots, x_{2d-t}$ from C_g^*, while the rest of the x_i come from $K \setminus C(T)$. Summing this for all $g \geq f \geq f_0$, and all $C_f \in M_f^1$, all $C_g^* \in M_g^2$ we are lead to the sum

$$\sum_{f \geq f_0} \sum_{g \geq f} \exp\left\{-\frac{1}{2}c(d)2^{-f}n\right\} 2^{-f(d-t)} 2^{-g(d-t)} S_{f,g}$$

where,

$$S_{f,g} = \sum_{M_f^1} \sum_{M_g^2} \left(\text{vol}\left(C_f \cap C_g^*\right)\right)^t.$$

Now $C_f \cap C_g^*$ is an axis-parallel dyadic box in K with side lengths $2^{-h_j} = \min(2^{-f_j}, 2^{-g_j^*})$ for $1 \leq j \leq d$. Thus, it is equal to a unique box $C_h \in M_h$ where $h = \sum h_j$. Consequently,

$$S_{f,g} \leq \sum_{h \geq g} 2^{-ht} \sum_{C_h \in M_h} W(f,g)$$

where $W(f,g)$ is the number of pairs C_f and C_g^* with $C_h = C_f \cap C_g^*$. For a fixed C_h, it is not hard to see that $W(f,g) \lesssim (2h - f - g + 1)^s$. Further, $|M_h| \lesssim h^{2d-s-1}$. Thus

$$S_{f,g} \lesssim \sum_{h \geq g} 2^{-ht} h^{2d-s-1} (2h - g - f + 1)^s.$$

The rest is a computation: one shows first that $S_{f,g}$ is dominated by the term when $h = g$. Then the next sum

$$\begin{aligned} &\sum_{g \geq f} \exp\left\{-\tfrac{1}{2} c(d) 2^{-f} n\right\} 2^{-f(d-t)} 2^{-g(d-t)} S_{f,g} \\ \lesssim\ &\sum_{g \geq f} \exp\left\{-\tfrac{1}{2} c(d) 2^{-f} n\right\} 2^{-f(d-t)} 2^{-g(d-t)} 2^{-gt} g^{2d-s-1} (g - f + 1)^s \end{aligned}$$

is again dominated by the term when $g = f$, implying that

$$\mathsf{P}[T \in \mathcal{F}(S) \text{ and } T^* \in \mathcal{F}(S^*)] \lesssim \sum_{f \geq f_0} \exp\left\{-\frac{1}{2} c(d) 2^{-f} n\right\} 2^{-f(2d-t)} f^{2d-s-1}.$$

Finally, the last sum is dominated by the term when $2^{-f} = \frac{1}{n}$ which gives

$$\mathsf{P}[T \in \mathcal{F}(S) \text{ and } T^* \in \mathcal{F}(S^*)] \lesssim \frac{(\log n)^{2d-s-1}}{n^{2d-t}}. \qquad \blacksquare$$

Remark 1. *This method works when $T \cap T^* = \emptyset$. Then $(\text{vol}(C_f \cap C_g^*))^t = 1$ since $t = 0$ and M_h does not appear at all. With a similar computation one could prove Claim 1 in the following form: $\mathsf{P}[T \in \mathcal{F}(S) \text{ and } T^* \in \mathcal{F}(S^*)] \lesssim \mathsf{P}[T \in \mathcal{F}(S)]\mathsf{P}[T^* \in \mathcal{F}(S^*)]$.*

Remark 2. *For general convex bodies the outcome depends on $S_{f,g}$ and then on $W(f,g)$. For smooth convex bodies $W(f,g)$ is a constant, and the computation is simpler.*

We now give a more direct proof for the case of the normal distribution.

Proof of Lemma 13 (Normal distribution). We prove the lemma directly for the Normal density. We write

$$\mathsf{P}(Y_1, Y_2 \in F) = \int_{x_1,\ldots x_n \in \mathbb{R}^d} \chi[Y_1 \in F]\, \chi[Y_2 \in F]\, \mathrm{d}f(x_1) \cdots \mathrm{d}f(x_n).$$

Assume that $Y_1 = \{x_1, \ldots, x_d\}, Y_2 = \{x_{d+1}, \ldots, x_{2d}\}$. For a subset Y, let $H(Y)$ be the hyperplane spanning Y and $V(Y)$ be measure of the distribution in the halfspace bounded by this hyperplane not containing the origin. Then the above probability can be bounded as

$$\mathsf{P}(Y_1, Y_2 \in F) \leq 2 \int_{x_1, \ldots, x_{2d}, V(Y_1) \geq V(Y_2)} (1 - V(Y_1))^{n-2d} \, \mathrm{d}f(x_1) \cdots \mathrm{d}f(x_{2d}). \tag{2}$$

We now estimate this when f is the standard Normal density in $\mathbb{R}^d$. Clearly, $V(Y_1)$ depends only on the distance of $H(Y_1)$ from the origin. So, it will be convenient to parametrize in terms of hyperplanes, and positions on them. This is achieved by the Blaschke-Petkantschin formula (for example, see [25]). For a spherically symmetric density function f we have,

$$\int_{x_1, \ldots, x_d \in \mathbb{R}^d} g(x_1, \ldots, x_d) \, \mathrm{d}f(x_1) \cdots \mathrm{d}f(x_d)$$
$$= \psi(d) \int_{H:(d-1)\text{-flat}} \int_{x_1, \ldots, x_d \in H} g(x_1, \ldots, x_d) \, \mathrm{d}f_H(x_1) \cdots \mathrm{d}f_H(x_d) d\mu(H).$$

Here $\psi(d)$ is a function of d alone and $\mathrm{d}\mu(H)$ is the measure induced on $(d-1)$-flats determined by picking d points from f. By spherical symmetry, $\mathrm{d}\mu(H) = \mathrm{d}h\mathrm{d}u$ where u is a unit vector (normal to H) and h is the distance of H from the origin. For the normal density, the measure of a random point on a flat is determined by its position on the flat and the distance of the flat to the origin, i.e., $df_H(x_1) = df_h(x_1)$. The integrand on the RHS of (2) depends only on the distance of $H(Y_1)$ from the origin. Thus, applying the formula twice, once for $x_{d+1}, \ldots, x_{2d}$ and then for $x_1, \ldots, x_d$, we have

$$\int_{x_1, \ldots, x_{2d}, V(Y_1) \geq V(Y_2)} (1 - V(Y_1))^{n-2d} \, \mathrm{d}f(x_1) \cdots \mathrm{d}f(x_{2d})$$
$$= \psi(d)^2 \int_{t_1 \leq t_2} (1 - V(t_1))^{n-2d} \, \mathrm{d}\nu(t_1)\mathrm{d}\nu(t_2)$$
$$\leq \frac{\psi(d)^2 \mathsf{vol}(S_d)^2}{(2\pi)^d} \int_{t_1} (1 - V(t_1))^{n-2d} e^{-\mathrm{d}t_1^2} (t_1)^{2(d-1)} \, \mathrm{d}t_1$$
$$\leq \frac{\psi(d)^2 \mathsf{vol}(S_d)^2}{(2\pi)^d} \int_t \left(1 - \frac{e^{-t^2/2}}{2\sqrt{2\pi}t}\right)^{n-2d} e^{-\mathrm{d}t^2} t^{2(d-1)} \, \mathrm{d}t$$
$$\lesssim n^{-2d} (\log n)^{d-1}.$$

where S_d is the d-dimensional unit sphere. Since

$$\mathsf{P}(Y_1 \in F) \gtrsim \frac{(\log n)^{(d-1)/2}}{\binom{n}{d}},$$

the lemma follows for the normal density. ∎

Proof of Claim 9. Let $\mathcal{G}^{-T}(S^*)$ be the set of useful facets of the polytope induced by the rows of S^* if we ignore the d points corresponding to the columns of T; define $\mathcal{G}^{-T^*}(S)$ similarly. Then clearly

$$\mathsf{P}(T \in \mathcal{G}(S) \wedge T^* \in \mathcal{G}(S^*)) \leq \mathsf{P}(T \in \mathcal{G}^{-T^*}(S) \wedge T^* \in \mathcal{G}^{-T}(S^*))$$
$$= \mathsf{P}(T \in \mathcal{G}^{-T^*}(S) | T^* \in \mathcal{G}^{-T}(S^*))\mathsf{P}(T^* \in \mathcal{G}^{-T}(S^*)).$$

But $\mathsf{P}(T \in \mathcal{G}^{-T^*}(S)|T^* \in \mathcal{G}^{-T}(S^*))$ is maximised when $S = S^*$. The result then follows from Lemma 13. ∎

Similarly Claim 10 follows by applying Lemma 14 or Lemma 16.

6. CONCLUDING REMARKS

We have shown that finding equilibria on average is easy. This raises several questions: (i) Can we extend the analysis to more general distributions? Our result is unaffected by linear transformations of the payoff matrices. So, for example, each matrix can be chosen from an arbitrary Gaussian. (ii) Our analysis shows that random games do not have small support Nash equilibria with polylogarithmically small probability. Can this bound be improved further? Even more importantly, does this simple algorithm have polynomial expected running time? (iii) We crucially use the fact that the mean of each entry is the same. Is this necessary? An algorithm that works for Gaussian entries with arbitrary means (and time polynomial in the largest variance), akin to smoothed analysis [29], would give a polynomial-time randomized algorithm for finding approximate Nash equilibria in arbitrary games [16]: add random Gaussians to the entries of the given payoff matrices; an equilibrium of the perturbed game will be an approximate equilibrium of the original game with high probability, given that the variance of the Gaussians is small enough. The current best algorithm for finding approximate equilibria has quasi-polynomial complexity [19].

Finally, we observe that finding approximate equilibria in random games is quite easy: for both the distributions we consider, with high probability there will be many pure strategy approximate equilibria and hence, by sampling, one of them can be found in time sublinear in the input size.

ACKNOWLEDGMENTS

REFERENCES

[1] F. Affentranger and J. Wieacker, On the convex hull of uniform random points in a simple d-polytope, Discrete Comput Geometry 6 (1991), 291–305.

[2] N. Alon and J. Spencer, The Probabilistic method, 2nd edition, Wiley, New York, 2000.

[3] I. Bárány, Intrinsic volumes and f-vectors of random polytopes, Math Ann 285 (1989), 671–699.

[4] I. Bárány and D. Larman, Convex bodies, economic cap coverings, random polytopes, Mathematika 35 (1988), 274–291.

[5] K. Borgwardt, Algorithms and Combinatorics Vol. 1, Springer-Verlag, 1980.

[6] X. Chen and X. Deng, Settling the complexity of 2-player Nash-equilibrium, Proceedings of 47th FOCS, Berkeley, October 2006, pp. 267–272.

[7] L. Devroye, A note on finding convex hulls via maximal vectors, Inform Process Lett 11 (1980), 53–56.

[8] J. Dickhaut and T. Kaplan, A program for finding Nash equilibria, The Math J 1 (1991), 87–93.

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

[10] R. Dwyer, On the convex hull of random points in a polytope, J Appl Probab 25 (1988), 688–699.

[11] R. Dwyer, Convex hulls of samples from spherically symmetric distributions, Discrete Appl Math 31 (1991), 113–132.

[12] A. Fabrikant, C. Papadimitriou, and K. Talwar, The complexity of pure-strategy Nash equilibria, Proceedings of 36th STOC, Chicago, June 2004.

[13] K. Goldberg, A. Goldman, and M. Newman, The probability of an equilibrium point, J Res Natl Bur Stand, Section B, 72 (1968), 93–101.

[14] H. Groemer, Limit theorems for the convex hull of random points in higher dimensions, Pac J Math 45 (1973), 525–533.

[15] I. Heuter, Limit theorems for the convex hull of random points in higher dimensions, Trans Am Math Soc 351 (1999), 4337–4363.

[16] A. Kalai, personal communication, 2005.

[17] D. Kirkpatrick and R. Seidel, The ultimate planar convex hull algorithm, SIAM J Comput 15 (1988), 287–299.

[18] C. Lemke and J. Howson, Equilibrium points of bimatrix games, J Soc Ind Appl Math 12 (1964), 413–423.

[19] R. Lipton, E. Markakis, and A. Mehta, Playing large games using simple strategies, Proceedings of E-Commerce, San Diego, June 2003, pp. 36–41.

[20] A. McLennan and J. Berg, The asymptotic expected number of Nash equilibria of two player normal form games, Game Econ Behav 51 (2005), 264–295.

[21] N. Megiddo, Improved asymptotic analysis of the average number of steps performed by the self-dual simplex algorithm, Math Program 35 (1986), 140–172.

[22] C. Papadimitriou, Algorithms, games and the internet, Proceedings of 33rd STOC, Heraklion, Crete, July 2001, pp. 749–753.

[23] R. Porter, E. Nudelman, and Y. Shoham, Simple search methods for finding a Nash equilibrium, Proceedings of 19th AAAI, San Jose, July 2004, pp. 664–669.

[24] H. Raynaud, Sur l'enveloppe convexe des nuages de points aleatiores dans $\mathbb{R}^n$. I., J Appl Probab 7 (1970), 35–48.

[25] L. Santaló, Integral geometry and geometric probability, Addison-Wesley, 1976.

[26] R. Savani and B. von Stengel, Exponentially many steps for finding a Nash equilibrium in a bimatrix game, Proceedings of 45th FOCS, Rome, October 2004, pp. 258–267.

[27] R. Schneider, "Discrete aspects of stochastic geometry," In Eds. J. Goodman and J. O'Rourke, Handbook of discrete and computational geometry, Boca Raton, CRC Press, 2004, pp. 255–278.

[28] S. Smale, On the average number of steps in the simplex method of linear programming, Math Program 27 (1983), 241–262.

[29] D. Spielman and S. Teng, Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time, J ACM 51 (2004), 385–463.

[30] B. van Wel, The convex hull of a uniform sample from the interior of a simple d-polytope, J Appl Probab 26 (1989), 259–273.

[31] W. Weil and J. Wieacker, "Stochastic geometry," In Eds. P. Gruber and J. Wills, Handbook of convex geometry, North-Holland, Amsterdam, 1993, pp. 1391–1438.