On Nash-equilibria of approximation-stable games

Pranjal Awasthi[1,]*, Maria-Florina Balcan[2], Avrim Blum[1], Or Sheffet[1] and Santosh Vempala[2]

[1]Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3891, Pennsylvania, USA
[2]School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332, USA

One of the goals of algorithmic game theory is to develop efficient methods for predicting the behaviour of self-interested agents in a given scenario, or game. A common approach is to attempt to compute an (approximate) Nash equilibrium, a behaviour such that agents have little incentive to deviate. Yet this computation appears to be quite difficult in general. In this work, we define and study games that are approximation-stable, meaning that all approximate equilibria predict similar behaviour. By analysing their properties, we show that finding approximate equilibria is substantially easier in such stable games.

Keywords: Behaviour prediction, Nash equilibrium, stable games, self-interested agents.

Introduction

GAME theory studies the behaviour of rational agents interacting in a given scenario (known as a game). This long-studied field has vast applications in many branches of science, including economics, philosophy and more recently computer science. As an example, consider several contractors bidding on a job for a client who will select the lowest bidder. This can be modelled as a game where each contractor is a player. Each player in a game has a set of strategies available to him/her; for example, the various bids a contractor could make. Players then individually choose one of their available strategies, and this combined choice then determines the payoffs that each player receives (e.g. which contractor gets the job and how much profit he makes). Players may also play *mixed strategies* which are probabilistic choices, such as 'with probability 1/3 bid 100 and with probability 2/3 bid 200'.

One of the most important concepts in game theory is that of an equilibrium state of a game. This refers to a set of mixed strategies, one for each player, such that no player has any incentive to deviate unilaterally from his own strategy even knowing the mixed strategies used by all the other players. In his seminal work, Nash[1,2] proved that every game has at least one such equilibrium (it could have many), now popularly known as a Nash equilibrium. A closely related notion is that of an approximate Nash equilibrium, where the players have atmost a very small incentive to deviate. Approximate Nash equilibria are especially natural and relevant in games with mixed Nash equilibrium strategies, where the players probabilistically choose their actions. In such cases, players might not be able to perfectly estimate their expected gain for each action. Additionally, in many scenarios, the computation of an approximate Nash equilibrium can reveal interesting structure about the nature of the true Nash equilibria and in fact might be sufficient for the desired tasks at hand such as prediction of behaviour.

With the rise of the Internet, including electronic commerce and social networking, the problem of efficiently computing Nash equilibria and approximate Nash equilibria has become increasingly important, in order to better predict how the participants can be expected to behave. As a first step towards this problem, substantial research has focused on two-player games where each player has n actions to select from[3–6]. One would ideally like to design algorithms whose running time depends polynomially on the complexity of the game. For example, an algorithm whose running time grows as n^2 can be fast even for games with thousands of actions. However, a running time of 2^n, such as would be produced by enumerating all possible mixed-strategy support-sets, is prohibitively large. Unfortunately, recent evidence suggests that substantially improving over 2^n for computing a Nash equilibrium will be difficult even for general two-player games[7,8]. The situation for *approximate* Nash equilibria, however, is somewhat better, and there is an algorithm with running time $n^{O((1/\varepsilon^2)\log n)}$ due to Lipton *et al.*[9] (the $O(\cdot)$ notation is used to hide multiplicative constant factors). This growth rate lies between polynomial and exponential.

The starting point of this work, in aiming to more efficiently compute approximate Nash equilibria, is to notice that it makes sense to focus on games where all approximate equilibria are close to each other. Otherwise, prediction of behaviour may not be possible even if the equilibria are known. If one focuses on such games, does the (approximate) equilibrium computation problem become easier? Here, we answer this question in the affirmative.

*For correspondence. (e-mail: pawasthi@cs.cmu.edu)

We do this by showing that games that are *approximation-stable*, meaning that all approximate equilibria are close to a true equilibrium, satisfy important structural conditions. These conditions then allow us to compute approximate Nash equilibria significantly more efficiently than the best guarantees known for general games.

We note that there has also been significant work on the complementary problem of approximate strategy-proofness in mechanism design, where one would like to *design* the rules of an interaction so that some desired behaviour (such as truth-telling) is an approximate equilibrium[10–12].

Definitions and preliminaries

Games: The focus of this article is two-player n-action games. We will use R and C to denote the payoff matrices for the two players, called the row and column players respectively. The entry $R[i, j]$ denotes the payoff which the row player gets if he chooses action i and the column player chooses action j. Similarly, the entry $C[i, j]$ denotes the corresponding payoff for the column player. An example of a two-player two-action game is 'matching pennies', where the players simultaneously make a binary choice (heads or tails), and R wins if the choices were the same and C wins if the choices were different. For this game, the R and C matrices are as follows

$$R = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

We will use p and q to denote mixed (randomized) strategies for the row and the column players respectively. Notice that any deterministic strategy is also a legal mixed strategy. We will use e_i to denote the probability vector with a 1 in entry i, corresponding to the deterministic strategy of always playing action i. For a given mixed strategy p, the support of p is the set of all actions which have non-zero probability under p. If the row and column players use strategies p and q respectively to play the game, then the expected payoff to the row player is p^TRq and the corresponding payoff to the column player is p^TCq. We will assume that all payoffs are scaled to the range [0, 1]. This scaling does not affect the nature of any equilibrium strategy.

Definition 1 (Nash equilibrium). *Consider a two-player n-action game with payoff matrices R and C. A pair of mixed strategies (p, q) is a Nash equilibrium if for all rows i, we have $e_i^TRq \le p^TRq$, and for all columns j, we have $p^TCe_j \le p^TCq$. In other words, the players cannot improve their expected payoff by unilaterally switching to a different strategy. We will typically use (p^*, q^*) to denote a Nash equilibrium. Note that in a Nash equilibrium (p^*, q^*), all rows i in the support of p^* satisfy $e_i^TRq^* = p^*Rq^*$ and similarly, all columns j in the support of q^* satisfy $p^{*T}Ce_j = p^{*T}Cq^*$.*

Definition 2 (ε-approximate Nash equilibrium). *Consider a two-player n-action game with payoff matrices R and C and payoffs scaled in the range* [0, 1]. *For a given $\varepsilon > 0$, a pair of mixed strategies (p, q) is an ε-equilibrium if for all rows i, we have $e_i^TRq \le p^TRq + \varepsilon$, and for all columns j, we have $p^TCe_j \le p^TCq + \varepsilon$. In other words, the players can increase their payoff by at most ε by unilaterally switching to a different strategy. Note that any Nash equilibrium (p^*, q^*) is an ε-approximate equilibrium for $\varepsilon = 0$.*

To discuss the distance between mixed strategies, we use variation distance, which is the total amount of probability mass one would need to move to make one distribution equal to the other. Specifically we define

$$d(q, q') = \frac{1}{2}\sum_i |q_i - q_i'| = \sum_i \max(q_i - q_i', 0). \tag{1}$$

We then define the distance between two strategy pairs as the maximum of the row-player's and column-player's distances, that is:

$$d((p, q), (p', q')) = \max[d(p, p'), d(q, q')]. \tag{2}$$

We now present our main definition, namely that of a game being *approximation-stable*.

Definition 3 (approximation stability). *A game satisfies (ε, Δ)-approximation stability if there exists a Nash equilibrium (p^*, q^*) such that any (p, q) that is an ε-equilibrium is Δ-close to (p^*, q^*), i.e. $d((p, q), (p^*, q^*)) \le \Delta$.*

We would often be interested viewing Δ as a function of ε. For a fixed ε, a smaller Δ means a stronger condition and a larger Δ means a weaker condition. Every game is $(\varepsilon, 1)$-approximation-stable, and as Δ gets smaller, we might expect for the game to exhibit more useful structure. Many natural games such as matching pennies and rock–paper–scissors satisfy (ε, Δ)-approximation-stability for $\Delta = O(\varepsilon)$; see the next section for analysis of a few simple examples.

All our results also apply to a weaker notion of approximation stability that allows for multiple equilibria, so long as moving distance Δ from any equilibrium produces a solution in which at least one player has ε incentive to deviate.

Some simple examples

A number of natural small games satisfy (ε, Δ)-approximation-stability for every $\varepsilon > 0$ and for $\Delta = O(\varepsilon)$. Here, we give a few simple examples.

Game 1: In this simple game, the row and the column matrices are 2×2 as follows

$$R = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}.$$

Here, the only Nash equilibrium (p^*, q^*) is for the row player to play row 1 and the column player to play column 1, which are dominant strategies. Any strategy that is at a distance Δ from p^* will give the row player Δ incentive to deviate, regardless of the strategy of the column player. Similarly, any deviation of Δ from q^* will give the column player Δ incentive to deviate regardless of the strategy of the row player. Hence, for every $\varepsilon \in [0, 1]$, this game is (ε, Δ)-stable for $\Delta = \varepsilon$.

Game 2: We now consider the matching pennies game described earlier. Denoting the indicator vectors as e_1 and e_2, the Nash equilibrium (p^*, q^*) is equal to $((e_1 + e_2)/2, (e_1 + e_2)/2)$. We now show that for any strategy which is Δ far from (p^*, q^*), at least one player must have ε incentive to deviate for $\varepsilon = \Delta((1 + 2\Delta)/(1 + 4\Delta))$.

Specifically, let (p, q) be Δ-far from (p^*, q^*), and without loss of generality assume $d(p, p^*) = \Delta$. We may further assume without loss of generality (by symmetry) that $p = (1/2 + \Delta)e_1 + (1/2 - \Delta)e_2$. Let $q = (1/2 - \Delta')e_1 + (1/2 + \Delta')e_2$ for $\Delta' \in [-\Delta, \Delta]$. In this case the row player is getting a payoff $p^TRq = (1/2 - 2\Delta\Delta')$. Furthermore, he can move to row 2 and get payoff $e_2^TRq = (1/2 + \Delta')$. Hence, the incentive to deviate is $(e_2 - p)^TRq \geq \Delta'(1 + 2\Delta)$. Similarly, the column player has payoff $p^TCq = (1/2 + 2\Delta\Delta')$, whereas $p^TCe_2 = (1/2 + \Delta)$, and hence has at least $\Delta(1 - 2\Delta')$ incentive to deviate. The maximum of these two is at least $\Delta((1 + 2\Delta)/(1 + 4\Delta))$ (with this value occurring at $\Delta' = \Delta/(1 + 4\Delta)$). Therefore, the incentive to deviate in any (p, q) that is Δ-far from (p^*, q^*) is at least this large. Solving for Δ as a function of ε, this game is (ε, Δ)-approximation-stable for $\Delta = \varepsilon + 2\varepsilon^2$.

Game 3: Rock–paper–scissors.

$$R = \begin{bmatrix} 0.5 & 0 & 1 \\ 1 & 0.5 & 0 \\ 0 & 1 & 0.5 \end{bmatrix}, \quad C = \begin{bmatrix} 0.5 & 1 & 0 \\ 0 & 0.5 & 1 \\ 1 & 0 & 0.5 \end{bmatrix}.$$

A case analysis (omitted) shows that this game is (ε, Δ)-approximation-stable for $\Delta = 4\varepsilon$, for any $\varepsilon \leq 1/6$.

Preliminaries

Before presenting the main technical result we present a few preliminary facts that apply to any two-player general-sum game.

Claim 1. *If* (p, q) *is* α*-close to a Nash equilibrium* (p^*, q^*) (*i.e. if* $d((p, q), (p^*, q^*)) \leq \alpha$), *then* ($p$, q) *is a* 3α*-Nash equilibrium.*

Proof. (see Appendix.) □

Claim 1 is useful because while it may be hard to determine how close some pair (p, q) is to a true equilibrium, it is easy to check how much incentive players have to deviate. Say that a Nash equilibrium (p^*, q^*) is *nontrivial* if at least one of p^* or q^* does not have full support over all the rows or columns. Trivial Nash equilibria, if they exist, can be computed in polynomial-time using linear programming. We then have

Claim 2. *For any nontrivial Nash equilibrium* (p^*, q^*) and any $\alpha > 0$, there exists (p, q) such that $d((p, q), (p^*, q^*)) \geq \alpha$ and (p, q) is an α-approximate equilibrium.

Proof. Without loss of generality, assume that p^* does not have full support. Let e_i be a row not in the support of p^*. Consider a pair of distributions (p, q^*) where $p = (1 - \alpha)p^* + \alpha e_i$. Since i was not in the support of p^*, (p, q^*) has variation distance α from (p^*, q^*). Yet, in (p, q^*), with probability $(1 - \alpha)$ both the players are playing their best responses to each other. Hence, no player has more than α incentive to deviate. □

Corollary 1. *Assume that the game* $\mathcal{G}$ *satisfies* (ε, Δ)*-approximation-stability and has a non-trivial Nash equilibrium. Then we must have* $\Delta \geq \varepsilon$.

The support of equilibria in stable games

A structural result of Lipton *et al.*[9] shows that for a two-player game with n actions available to each player, there always exists, for each player, a multiset of actions of size at most $O((1/\varepsilon^2) \log n)$ with the following property: if the strategy of each player is to choose an action uniformly at random from his/her respective sets and play it, then the two strategies form an ε-approximate Nash equilibrium. This gives an immediate $n^{O((1/\varepsilon^2) \log n)}$-time algorithm for computing ε-approximate equilibria and has also been shown to be existentially tight[13]. We now show that approximation-stable games have a structure that can be used to improve the efficiency of such algorithms for computing approximate equilibria.

Theorem 1. *For any game satisfying* (ε, Δ)*-approximation-stability, there exists an* ε*-equilibrium where each player's strategy has support* $O((\Delta/\varepsilon)^2 \log(1 + 1/\Delta) \log n)$.

Corollary 2. *There is an algorithm to find* ε*-equilibria in games satisfying* (ε, Δ)*-approximation-stability, running in time* $n^{O((\Delta/\varepsilon)^2 \log(1 + 1/\Delta) \log n)}$.

Let $S = c(\Delta/\varepsilon)^2 \log n$ for some absolute constant c, and let (p^*, q^*) denote the Nash equilibrium such that all ε-equilibria lie within distance Δ of (p^*, q^*). Theorem 1 is proven in stages. First, in Lemma 1 we show that given a pair of distributions (p, q), if p is near-uniform over a large support, then p can be written as a convex combination $p = xp_1 + (1 - x)p_2$, where p_1 and p_2 have disjoint supports, and for every column j, j's performance against p_1 is close to its performance against p_2. This implies p^* itself cannot be near-uniform over a large-sized support, since otherwise we could write it in this way and then shift Δ probability mass from p_2 to p_1, producing a new distribution p' such that under (p', q^*), the column player has less than ε incentive to deviate (and the row player has zero incentive to deviate since $\mathrm{supp}(p') \subseteq \mathrm{supp}(p^*)$). This contradicts the fact that the game is (ε, Δ)-approximation-stable. We then build on this to show that if p^* is not near-uniform and does have a large support, it must be well-approximated by a distribution of small support (roughly $O(S \log(1/\Delta))$). This analysis combines Lemma 1 together with the sampling idea of Lipton *et al.*[9]. The same, of course, applies to q^*. For the rest of this section we assume that $\Delta \le 1/4$.

Lemma 1. *For any distributions* p *and* q, *if* p *satisfies* $||p||_2^2 \le 1/S$, *where* $S = c(\Delta/\varepsilon)^2 \log n$ *for sufficiently large constant* c, *then* p *can be written as a convex combination* $p = xp_1 + (1 - x)p_2$ *of two distributions* p_1 *and* p_2 *over disjoint supports such that:*

(i) $x \le 3/4 \le 1 - \Delta$.
(ii) $\forall j, (p_1 - p)^T C(e_j - q) < \varepsilon/(4\Delta)$.

The point of Lemma 1 is that by (i) and (ii), modifying p by moving Δ probability mass from p_2 to p_1 can improve the performance of e_j relative to q for the column player by at most ε. The proof of Lemma 1 makes extensive use of the Hoeffding bound.

Theorem 2 (Hoeffding bound). *Let* X_i, $i = 1, 2, \ldots, n$, *be* n *random variables, such that* $\forall i$, $X_i \in [a_i, b_i]$. *Let* $\mu_i = \mathbf{E}[X_i]$. *Then for every* $t > 0$ *we have that*

$$\mathbf{Pr}\left[\sum_i X_i > t + \sum_i \mu_i\right] \le \exp\left(-\frac{t^2}{\sum_i (b_i - a_i)^2}\right). \quad (3)$$

Proof [Lemma 1]. Let r be a random subset of the support of p; that is, for every element in $\mathrm{supp}(p)$, add it to r with probability 1/2. Also, let C_i denote the ith entry of Cq. The idea of the proof is just to argue that for any column j, by the Hoeffding bound, with high probability over the choice of r, the distribution p_1 induced by p restricted to r satisfies the desired condition that $p_1^T C(e_j - q)$ is within $\varepsilon/(4\Delta)$ of $p^T C(e_j - q)$. We then simply perform a union bound over j.

Fix column e_j. Let Y_{ij} be the random variable defined as $2p_i(C_{ij} - C_i)$ if element i was added to r, and 0 otherwise. Observe that $\mathbf{E}[\sum_i Y_{ij}] = 1/2 \sum_i 2p_i(C_{ij} - C_i) = p^T C(e_j - q)$. Let Z_i be the random variable defined as $2p_i$ with probability 1/2 (if element i was added to r), and 0 otherwise. Observe $\mathbf{E}[\sum_i Z_i] = 1$. Observe also that for every i we have $Z_i, Y_{ij} \in [-2p_i, 2p_i]$.

The obvious reason for defining Y_{ij} and Z_i is that by denoting the distribution p restricted to r (renormalized to have L_1 norm equal to 1) as p_r, we have:

$$p_r^T C(e_j - q) = \frac{\sum_{i \in r} p_i (C_{ij} - C_i)}{\sum_{i \in r} p_i} = \frac{\sum_i Y_{ij}}{\sum_i Z_i}, \quad (4)$$

so by bounding the numerator from above and the denominator from below, we can hope to find r for which $p_r^T C(e_j - q) < \mathbf{E}[\sum_i Y_{ij}] + (\varepsilon/4\Delta)$, thus decomposing p into the desired $p_1 = p_r$ and $p_2 = p_{\bar{r}}$. We can do this using the Hoeffding bound and plugging the value of S.

$$\mathbf{Pr}\left[\sum_i Y_{ij} > p^T C(e_j - q) + \frac{\varepsilon}{10\Delta}\right]$$

$$< \exp\left(\frac{-(\varepsilon/10\Delta)^2}{\sum_i (4p_i)^2}\right) \le \exp\left(\frac{-S\varepsilon^2}{(40\Delta)^2}\right) < \frac{1}{2n},$$

where the last inequality is by definition of S. Thus, $\mathbf{Pr}[\exists j, \sum_i Y_{ij} > p^T C(e_j - q) + (\varepsilon/10\Delta)] < 1/2$. Similarly (and even simpler), we have $\mathbf{Pr}[\sum_i Z_i < 1 - (\varepsilon/(10\Delta))] < 1/2$, and so the existence of r for which both events do not hold is proven. Observe that for this specific r we have

$$\frac{\sum_i Y_{ij}}{\sum_i Z_i} \le \frac{p^T C(e_j - q) + \varepsilon/10\Delta}{1 - \varepsilon/10\Delta} \le p^T C(e_j - q)$$

$$+ \frac{\varepsilon/5\Delta}{1 - \varepsilon/10\Delta} \le p^T C(e_j - q) + \frac{\varepsilon}{4\Delta},$$

using the fact that $p^T C(e_j - q) \le 1$. Thus, we have the desired decomposition of p. □

Proof [Theorem 1]. We begin by partitioning p^* into its *heavy* and *light* parts. Specifically, greedily remove the largest entries of p^* and place them into a set H (the heavy elements) until either (a) $\mathbf{Pr}[H] \ge 1 - 4\Delta$, or (b) the remaining entries L (the light elements) satisfy the condition that $\forall i \in L$, $\mathbf{Pr}[i] \le 1/S\,\mathbf{Pr}[L]$ for S as in Lemma 1, whichever comes first. We analyse each case in turn.

If case (a) occurs first, then clearly H has at most $S \log(1/4\Delta)$ elements. We now simply apply the sampling argument of Lipton *et al.*[9] to L and combine the result with H. Specifically, decompose p^* as $p^* = \beta p_H + (1 - \beta)p_L$, where β denotes the total probability mass over H. The sampling argument of Lipton *et al.*[9] now implies by sampling a multiset $\mathcal{X}$ of S elements from $\text{supp}(p_L) = L$, we are guaranteed, by definition of S, that for any column e_j, $|(U_{\mathcal{X}})^T Ce_j - p_L^T Ce_j| \le (\varepsilon/8\Delta)$, where $U_{\mathcal{X}}$ is the uniform distribution over $\mathcal{X}$. This means that for $\tilde{p} = \beta p_H + (1-\beta)U_{\mathcal{X}}$, all columns e_j satisfy $|p^{*T}Ce_j - \tilde{p}^{\,T}Ce_j| \le \varepsilon/2$. We have thus found (the row portion of) an ε-equilibrium with support of size $S(1 + \log(1/4\Delta))$ as desired, and now simply apply the same argument to q^*.

If case (b) occurs first, we show that the game cannot satisfy (ε, Δ)-approximation-stability. Specifically, let p_L denote the induced distribution produced by restricting p^* to L and renormalizing so that $\sum_i (p_L)_i = 1$, then $\sum_i (p_L)_i^2 \le 1/S \sum_i (p_L)_i = 1/S$. Using Lemma 1, we can write p_L as a convex combination: $p_L = xp_1 + (1-x)p_2$ of p_1 and p_2 satisfying the properties of Lemma 1. Again, by denoting β as the total probability mass over H, we have

$$p^* = \beta p_H + (1-\beta)xp_1 + (1-\beta)(1-x)p_2, \tag{5}$$

where p_H is the induced distribution over H. We now consider the transition from p^* to p' defined as

$$p' = \beta p_H + ((1-\beta)x + \Delta)p_1 + ((1-\beta)(1-x) - \Delta)p_2. \tag{6}$$

Notice that by Lemma 1, $x \le 3/4$ and hence $(1-\beta) \times (1-x) - \Delta \ge (1-\beta)/4 - \Delta \ge 0$, so p' is a valid probability distribution. Also, since p_1 and p_2 are distributions over disjoint support, p' is Δ far from p^*. Note that since p' is obtained from an internal deviation within the support of p^*, the row player has no incentive to deviate when playing p' against q^*. So, if the game is (ε, Δ)-approximation-stable, then playing p' against q^* must cause the column player to have more than ε incentive to deviate. However, by transitioning from p^* to p', the expected gain of switching from q^* to any e_j is

$$\begin{aligned} p'^T C(e_j - q) &= (p^* + \Delta(p_1 - p_2))^T C(e_j - q^*) \\ &\le \Delta(p_1 - p_2)^T C(e_j - q^*) \\ &\quad (\text{since } p^{*T} Cq^* \ge p^{*T} Ce_j). \end{aligned}$$

From Lemma 1 we know that for every column j, $(p_1 - p_L)^T C(e_j - q^*) < \varepsilon/4\Delta$. Also we have that $p_2 = (p_L - xp_1)/(1-x)$. Using this we can write $\Delta(p_1 - p_2)^T C(e_j - q^*) = (\Delta/(1-x))(p_1 - p_L)^T C(e_j - q^*) < (\Delta/(1-x))(\varepsilon/4\Delta) \le \varepsilon$, where the last step follows from $x \le 3/4$. So the column player has less than ε incentive to deviate, which contradicts the fact that the game is (ε, Δ)-approximation-stable.

Polynomial-time algorithms when Δ and ε are close

While in the previous section we show that one can get some advantage over the previously best known results in the case of approximation stable games, we now show how to get a dramatic advantage for $\Delta \le 2\varepsilon$. More formally, we prove that if $\Delta \le 2\varepsilon - 6\varepsilon^2$, then there must exist an $O(\varepsilon)$-equilibrium where each player's strategy has support $O(1/\varepsilon)$. Thus, in this case, for constant ε, we have a polynomial-time algorithm for computing $O(\varepsilon)$-equilibria.

Theorem 3. *For any game satisfying (ε, Δ)-approximation-stability for $\Delta \le 2\varepsilon - 6\varepsilon^2$, there exists an $O(\varepsilon)$-equilibrium where each player's strategy has support $O(1/\varepsilon)$. Thus, $O(\varepsilon)$-equilibria can be computed in time $n^{O(1/\varepsilon)}$.*

Proof. Let (p^*, q^*) be a Nash equilibrium of the game. First, if there is no set S of rows having a combined total probability mass $x \in [\Delta, \Delta + \varepsilon]$ in p^*, then this implies that except for rows of total probability mass less than Δ, all rows in the support of p^* have probability greater than ε. Therefore, p^* is Δ-close to a distribution of support at most $1/\varepsilon$. If this is true for q^* as well, then this implies (p^*, q^*) is Δ-close to a pair of strategies (p, q) each of support $\le 1/\varepsilon$, which by Claim 1 and the assumption $\Delta < 2\varepsilon$, is an $O(\varepsilon)$-equilibrium as desired. So, to prove the theorem, it suffices to show that if such a set S exists, then the game cannot satisfy (ε, Δ)-approximation-stability for $\Delta \le 2\varepsilon - 6\varepsilon^2$.

Therefore, assume for contradiction that p^* can be written as a convex combination

$$p^* = xp_1 + (1-x)p_2, \tag{7}$$

where p_1, p_2 have disjoint supports and $x \in [\Delta, \Delta + \varepsilon]$. Let $\gamma = p_1^T Cq^* - p_2^T Cq^*$ and let $V_C = p^{*T} Cq^*$. We now consider two methods for moving distance Δ from p^*: moving probability from p_1 to p_2, and moving probability from p_2 to p_1. Let

$$p' = (x - \Delta)p_1 + (1 - x + \Delta)p_2 \tag{8}$$

$$= \left(1 + \frac{\Delta}{1-x}\right)p^* - \left(\frac{\Delta}{1-x}\right)p_1. \tag{9}$$

Since p' has distance Δ from p^* and its support is contained in the support of p^*, by approximation-stability, there must exist some column e_j such that $p'^T Ce_j \ge p'^T Cq^* + \varepsilon$. From eq. (8) we have $p'^T Cq^* = V_C - \Delta(p_1 - p_2)^T \times Cq^* = V_C - \Delta\gamma$. From eq. (9) and the fact that $p^{*T} Ce_j \le V_C$, we have $p'^T Ce_j \le V_C(1 + (\Delta/(1-x)))$. Therefore we have the constraint

$$V_C\left(1 + \frac{\Delta}{1-x}\right) \ge V_C - \Delta\gamma + \varepsilon. \tag{10}$$

Now, consider moving Δ probability mass from p_2 to p_1. Specifically, let

$$p'' = (x + \Delta)p_1 + (1 - x - \Delta)p_2 \quad (11)$$

$$= \left(1 - \frac{\Delta}{1-x}\right) p^* + \left(\frac{\Delta}{1-x}\right) p_1. \quad (12)$$

Again, there must exist some column e_k such that $p''^T Ce_k \geq p''^T Cq^* + \varepsilon$. From eq. (11) we have $p''^T Cq^* = V_C + \Delta(p_1 - p_2)^T Cq^* = V_C + \Delta\gamma$. From eq. (12) and the fact that $p^{*T} Ce_k \leq V_C$, we have

$$p''^T Ce_k \leq V_C \left(1 - \frac{\Delta}{1-x}\right) + \frac{\Delta}{1-x}.$$

Therefore, we have the constraint

$$V_C \left(1 - \frac{\Delta}{1-x}\right) + \frac{\Delta}{1-x} \geq V_C + \Delta\gamma + \varepsilon. \quad (13)$$

From constraint (10) we have

$$V_C \left(\frac{\Delta}{1-x}\right) \geq \varepsilon - \Delta\gamma$$

and from constraint (13) we have

$$V_C \left(\frac{\Delta}{1-x}\right) \leq \frac{\Delta}{1-x} - \Delta\gamma - \varepsilon.$$

Therefore, $\Delta/(1 - x) \geq 2\varepsilon$, contradicting $\Delta \leq 2\varepsilon - 6\varepsilon^2$. □

We would like to point out that one can give a near-matching lower bound to the results mentioned in the section on support of equilibria in stable games, showing that there exist stable games in which the Nash equilibrium and all approximate equilibria have support $\Omega(\log n)$. Formally, we have the following theorem whose proof is omitted.

Theorem 4. *For any $\Delta \leq 1/2$, there exists an (ε, Δ)-approximation-stable game $\mathcal{G}$ for some $\varepsilon > 0$, such that all ε-equilibria have support $\Omega((\Delta^4/\varepsilon^2) \log n)$.*

Summary

In this work we define and analyse a natural notion of stability for two-player games, motivated by the goal of finding approximate equilibria for predictive purposes. We show that many natural games satisfy our notion of approximation-stability. We then design an algorithm for computing an approximate Nash equilibrium in such stable games. We show that our algorithm runs faster on stable games and is no worse than previously proposed algorithms for general games. We also design polynomial time algorithms for computing approximate equilibria when the stability parameters Δ and ε are close to each other. Our future goal is to better understand for what values of Δ and ε can one find approximate equilibria efficiently. Recently, Balcan and Braverman[14] have shown that this may be intrinsically hard for $\Delta \geq 8\varepsilon^{1/4}$. More broadly, perhaps our notion of stability can shed light on the important problem of obtaining an efficient algorithm for finding approximate equilibria in general games.

Appendix. Additional proofs

Proof (*Claim 1*). Define $p_i^{\min} = \min(p_i, p_i^*)$ and $p^{\min} = (p_1^{\min}, \ldots, p_n^{\min})$. So, $p = p^{\min} + p'$, $p^* = p^{\min} + p''$, where $\sum_i p_i' \leq \alpha$ and $\sum_i p_i'' \leq \alpha$. Similarly, for q, q^*, define $q^{q\min}$, q', q''.

Let $v_R = p^{*T} Rq^*$ be the value to the row player in (p^*, q^*). Since (p^*, q^*) is Nash, we know $e_i^T Rq^* \leq v_R$ for all rows i. We now show that the best response to q is at most $v_R + \alpha$. To see this, consider some row e_i. The expected payoff is

$$\sum_j q_j R_{i,j} = \sum_j (q_j^{q\min} + q_j') R_{i,j}$$

$$\leq (q_j^{q\min} + q_j'') R_{i,j} + \alpha \leq v_R + \alpha. \quad (14)$$

The middle inequality holds because $\sum_j q_j' R_{i,j} \leq \alpha$ (since $R_{i,j} \in [0, 1]$), and the last because $\sum_j q_j'' R_{i,j} \geq 0$.

We now show that the expected payoff of p against q is at least $v_R - 2\alpha$.

$$p^T Rq = \sum_{i,j} p_i q_j R_{i,j} = \sum_{i,j} (p_i^* - p_i'' + p_i')(q_j^* - q_j'' + q_j') R_{i,j} \quad (15)$$

$$= v_R + \left(\sum_{i,j} p_i^* (-q_j'' + q_j') R_{i,j} + \left(\sum_{i,j} (-p_i'' + p_i') q_j R_{i,j} \right) \right) \quad (16)$$

$$\geq v_R + \left(\sum_i p_i^* (-\alpha) \right) + \left(\sum_j (-\alpha) q_j \right) = v_R - 2\alpha. \quad (17)$$

Thus, under (p, q), the row player has no more than 3α incentive to deviate, and we have the analogous argument for the column player. So, (p, q) is a 3α-equilibrium. □

1. Nash, J. F., Equilibrium points in *n*-person games. In Proceedings of the National Academy of Sciences, USA, 1950, vol. 36, pp. 48–49.
2. Nash, J. F., Non-cooperative games. *Ann. Math.*, 1951, **54**, 286–295.

3. Daskalakis, C., Mehta, A. and Papadimitriou, C. H., Progress in approximate Nash equilibria. In Proceedings of the 8th ACM Conference on Electronic Commerce, San Diego, CA, USA, 2007.
4. Tsaknakis, H. and Spirakis, P. G., An optimization approach for approximate Nash equilibria. In *WINE*, 2007, pp. 42–56.
5. Bosse, H., Byrka, J. and Markakis, E., New algorithms for approximate Nash equilibria in bimatrix games. *Theor. Comput. Sci*., 2010, **411**, 164–173.
6. Bárány, I., Vempala, S. and Vetta, A., Nash equilibria in random games. *Random Struct. Algorithms*, 2007, **31**, 391–405.
7. Chen, X. and Deng, X., Settling the complexity of two-player Nash equilibrium. In Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, Berkeley, CA, USA, 2006, pp. 261–272.
8. Daskalakis, C., Goldberg, P. W. and Papadimitriou, C. H., The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 2009, **39**, 195–259.
9. Lipton, R. J., Markakis, E. and Mehta, A., Playing large games using simple strategies. In Proceedings of the 4th ACM Conference on Electronic Commerce, 2003, pp. 36–41.
10. Procaccia, A. D. and Tennenholtz, M., Approximate mechanism design without money. In Proceedings of the 10th ACM Conference on Electronic Commerce, Stanford, CA, USA, 2009.
11. Lu, P., Sun, X., Wang, Y. and Zhu, Z. A., Asymptotically optimal strategy-proof mechanisms for two-facility games. In Proceedings of the 11th ACM Conference on Electronic Commerce, 2010.
12. Birrell, E. and Pass, R., Approximately strategy-proof voting. In Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, Barcelona, Spain, 2011, vol. 1.
13. Feder, T., Nazerzadeh, H. and Saberi, A., Approximating Nash equilibria using small-support strategies. In Proceedings of the 8th ACM Conference on Electronic Commerce, San Diego, 2007, pp. 352–354.
14. Balcan, M.-F. and Braverman, M., Nash equilibria in perturbation resilient games. Manuscript, 2012.

ACKNOWLEDGEMENTS. This work was supported in part by NSF grants CCF-0830540 and CCF-0953192, ONR grant N00014-09-1-0751,and AFOSR grant FA9550-09-1-0538.