THEORY OF COMPUTING, Volume 13 (13), 2017, pp. 1–31
www.theoryofcomputing.org

Nash Equilibria in Perturbation-Stable Games

Maria-Florina Balcan[*] Mark Braverman[†]

Received November 23, 2012; Revised August 30, 2017; Published November 13, 2017

Abstract: Motivated by the fact that in many game-theoretic settings, the game analyzed is only an approximation to the game being played, in this work we analyze equilibrium computation for the broad and natural class of bimatrix games that are stable under perturbations. We specifically focus on games with the property that small changes in the payoff matrices do not cause the Nash equilibria of the game to fluctuate wildly. For such games we show how one can compute approximate Nash equilibria more efficiently than the general result of Lipton et al. (EC'03), by an amount that depends on the degree of stability of the game and that reduces to their bound in the worst case. Additionally, we show that for stable games, the approximate equilibria found will be close in variation distance to true equilibria, and moreover this holds even if we are given as input only a perturbation of the actual underlying stable game.

For uniformly stable games, where the equilibria fluctuate at most quasi-linearly in the extent of the perturbation, we get a particularly dramatic improvement. Here, we

[*]Supported in part by NSF grants CCF-0953192 and CCF-1101283, ONR grant N00014-09-1-0751, AFOSR grant FA9550-09-1-0538, a Google Research Award, and a Microsoft Faculty Fellowship. This work was done in part while the author was visiting Microsoft Research NE.

[†]This work was done in part while the author was a member of Microsoft Research NE. Research supported in part by NSF Awards, DMS-1128155, CCF-1525342, and CCF-1149888, a Packard Fellowship in Science and Engineering, and the Simons Collaboration on Algorithms and Geometry.

ACM Classification: F.1.3, F.2.1, J.4

AMS Classification: 68Q17, 68W25, 91A05, 91A10

Key words and phrases: algorithmic game theory, approximation algorithms, Nash equilibrium, beyond worst-case analysis, perturbation stability

 DOI: 10.4086/toc.2017.v013a013

achieve a fully quasi-polynomial-time approximation scheme, that is, we can find $1/\text{poly}(n)$-approximate equilibria in quasi-polynomial time. This is in marked contrast to the general class of bimatrix games for which finding such approximate equilibria is PPAD-hard. In particular, under the (widely believed) assumption that PPAD is not contained in quasi-polynomial time, our results imply that such uniformly stable games are inherently easier for computation of approximate equilibria than general bimatrix games.

1 Introduction

The Nash equilibrium solution concept has a long history in economics and game theory as a description for the natural result of self-interested behavior [31, 35]. Its importance has led to significant effort in the computer science literature in recent years towards understanding their computational structure, and in particular on the complexity of finding both Nash and approximate Nash equilibria. A series of results culminating in the work by Daskalakis, Goldberg, and Papadimitriou [18] and Chen, Deng, and Teng [15, 16] showed that finding a Nash equilibrium or even a $1/\text{poly}(n)$-approximate equilibrium, is PPAD-complete even for 2-player bimatrix games. For general values of ε, the best known algorithm for finding ε-approximate equilibria runs in time $n^{O((\log n)/\varepsilon^2)}$, based on a structural result of Lipton et al. [28] showing that there always exist ε-approximate equilibria with support over at most $O((\log n)/\varepsilon^2)$ strategies. This structural result has been shown to be existentially tight [23]. Even for large values of ε, despite considerable effort [20, 19, 33, 23, 14, 26], polynomial-time algorithms for computing ε-approximate equilibria are known only for $\varepsilon \geq 0.3393$ [33]. Recently, Rubinstein [32] showed that unless surprisingly fast algorithms exist for PPAD-complete problems (running in time $2^{\tilde{O}(\sqrt{n})}$), no general poly-time algorithm for computing ε-approximate equilibria, for a sufficiently small constant ε, exists. These results suggest a difficult computational landscape for equilibrium and approximate equilibrium computation on worst-case instances.

In this paper we go beyond worst-case analysis and investigate the equilibrium computation problem in a natural class of bimatrix games that are stable under perturbations. As we argue, on one hand, such games can be used to model many realistic situations. On the other hand, we show that they have additional structure which can be exploited to provide better algorithmic guarantees than believed to be possible on worst-case instances. The starting point of our work is the realization that games are typically only abstractions of reality and except in the most controlled settings, payoffs listed in a game that represents an interaction between self-interested agents are only approximations to the agents' exact utilities.[1] As a result, for problems such as equilibrium computation, it is natural to focus attention on games that are robust to the exact payoff values, in the sense that small changes to the entries in the game matrices do not cause the Nash equilibria to fluctuate wildly; otherwise, even if equilibria can be computed, they may not actually be meaningful for understanding behavior in the game that is played. In this work, we focus on such games and analyze their structural properties as well as their implications to the equilibrium computation problem. We show how their structure can be leveraged to obtain better algorithms for computing approximate Nash equilibria, as well as strategies close to true Nash equilibria.

[1] For example, if agents are two corporations with various possible actions in some proposed market, the precise payoffs to the corporations may depend on specific quantities such as demand for electricity or the price of oil, which cannot be fully known in advance but only estimated.

Furthermore, we provide such algorithmic guarantees even if we are given only a perturbation of the actual stable game being played.

To formalize such settings we consider bimatrix games G that satisfy what we call (ε,Δ) *perturbation stability*, meaning that for any game G' within L_∞ distance ε of G (each entry changed by at most ε), each Nash equilibrium (p',q') in G' is Δ-close to some Nash equilibrium (p,q) in G, where closeness is given by variation distance. Clearly, any game is $(\varepsilon,1)$ perturbation stable for any ε, and the smaller the Δ, the more structure the (ε,Δ) perturbation-stable games have. In this paper we study the meaningful range of parameters, several structural properties, and the algorithmic behavior of these games.

Our first main result shows that for an interesting and general range of parameters, the structure of perturbation stable games can be leveraged to obtain better algorithms for equilibrium computation. Specifically, we show that for any $0 \le \varepsilon \le \Delta \le 1$, all n-action (ε,Δ) perturbation-stable games with at most $n^{O((\Delta/\varepsilon)^2)}$ Nash equilibria must have a well-supported ε-equilibrium of support

$$O\left(\frac{\Delta^2\log(1+\Delta^{-1})}{\varepsilon^2}\log n\right).$$

This yields an $n^{O((\Delta^2/\varepsilon^2)\log(1+\Delta^{-1})\log n)}$-time algorithm for finding an ε-equilibrium, improving by a factor $O(1/(\Delta^2\log(1+\Delta^{-1})))$ in the exponent over the bound of [28] for games satisfying this condition (and reducing to the bound of [28] in the worst-case when $\Delta = 1$).[2] Moreover, the stability condition can be further used to show that the approximate equilibrium found will be Δ-close to a true equilibrium, and this holds even if the algorithm is given as input only a perturbation to the true underlying stable game.

A particularly interesting class of games for which our results provide a dramatic improvement are those that satisfy what we call *t-uniform stability to perturbations*. These are games that for some $\varepsilon_0 = 1/\mathrm{poly}(n)$ and some $t \ge 1$ satisfy the $(\varepsilon, t\varepsilon)$ perturbation-stability condition for all $\varepsilon < \varepsilon_0$. For games satisfying t-uniform perturbation-stability with $t = \mathrm{poly}(\log(n))$, our results imply that we can find $1/\mathrm{poly}(n)$-approximate equilibria in $n^{\mathrm{poly}(\log(n))}$ time, i. e., achieve a fully quasi-polynomial-time approximation scheme (FQPTAS). This is especially interesting because the results of [16] prove that it is PPAD-hard to find $1/\mathrm{poly}(n)$-approximate equilibria in general games. Our results show that under the (widely believed) assumption that PPAD is not contained in quasi-polynomial time [17], such uniformly stable games are inherently easier for computation of approximate equilibria than general bimatrix games.[3] Moreover, variants of many games appearing commonly in experimental economics including the public goods game, matching pennies, and identical interest game [22] satisfy this condition. See Section 3, Section 5, and Appendix B for detailed examples.

Our second main result shows that there exist constants c_1, c_2 such that computing an ε-equilibrium in a game satisfying the $(\varepsilon, c_1\varepsilon^{1/4})$ perturbation-stability condition is as hard as computing a $c_2\varepsilon^{1/4}$-equilibrium in a general game. For our reduction, we show that any general game can be embedded into one having the $(\varepsilon, c_1\varepsilon^{1/4})$ perturbation-stability property such that an ε equilibrium in the new game yields a $c_2\varepsilon^{1/4}$-equilibrium in the original game. This result implies that the interesting range for the

[2] One should think of Δ as a function of ε, with both possibly depending on n. E. g., $\varepsilon = 1/\sqrt{n}$ and $\Delta = 6\varepsilon$. Additionally, note that any game that is stable for $\Delta < \varepsilon$ is also stable for $\Delta = \varepsilon$; therefore, the condition "$\varepsilon \le \Delta$" is not needed if we replace Δ with $\max\{\varepsilon,\Delta\}$ in the support size guaranteed and the number of equilibria allowed.

[3] The generic result of [28] achieves quasi-polynomial time only for $\varepsilon = \Omega(1/\mathrm{poly}(\log n))$.

(ε,Δ)-perturbation-stability condition i. e., where one could hope to do significantly better than in the general case, is $\Delta = o(\varepsilon^{1/4})$.

We also connect our perturbation-stability condition to a seemingly very different approximation-stability condition introduced by Awasthi et al. [5, 6]. Formally, a game satisfies the strong (ε,Δ)-approximation-stability condition if all ε-approximate equilibria are contained inside a small ball of radius Δ around a single equilibrium.[4] We prove that our perturbation-stability condition is equivalent to a much weaker version of this condition that we call *well-supported approximation stability*. This condition requires only that for any well-supported ε-approximate equilibrium[5] (p,q) there *exists* a Nash equilibrium (p^*,q^*) that is Δ-close to (p,q). Clearly, the well-supported approximation-stability condition is more general than strong (ε,Δ)-approximation-stability since rather than assuming that there exists a fixed Nash equilibrium (p^*,q^*) such that all ε-approximate equilibria are contained in a ball of radius Δ around (p^*,q^*), it requires only that for any well-supported ε-approximate equilibrium (p,q) there exists a Nash equilibrium (p^*,q^*) that is Δ-close to (p,q). Thus, perturbation-stable games are significantly more expressive than strongly approximation-stable games and Section 3 presents several examples of games satisfying the former but not the latter. However, our lower bound (showing that $\Delta = \omega(\varepsilon^{1/4})$ is as hard as the general case) also holds for the strong approximation-stability condition.

We also provide an interesting structural result showing that for $\Delta \geq \varepsilon$, each Nash equilibrium of an (ε,Δ) perturbation-stable game with $n^{O((\Delta/\varepsilon)^2)}$ Nash equilibria must be 8Δ-close to a well-supported ε-approximate equilibrium of support only $O((\Delta^2/\varepsilon^2)\log(1+\Delta^{-1})\log n)$. Similarly, a t-uniformly stable game with $n^{O(t^2)}$ equilibria has the property that for any Δ, each equilibrium is 8Δ-close to a well-supported Δ/t-approximate equilibrium with support of size $O(t^2\log(1+\Delta^{-1})\log n)$. This property implies that in quasi-polynomial time we can in fact find a *set* of approximate Nash equilibria that cover (within distance 8Δ) the set of *all* Nash equilibria in such games.

It is interesting to note that for our algorithmic results for finding approximate equilibria we do *not* require knowing the stability parameters. If the game happens to be reasonably stable, then we get improved running times over the Lipton et al. [28] guarantees; if this is not the case, then we fall back to the Lipton et al. [28] bound.[6] However, given a game, it might be interesting to know how stable it is. In this direction, we provide a characterization of stable constant-sum games along with an algorithm for computing the strong stability parameters of a given constant-sum game.

1.1 Related work

In addition to results on computing (approximate) equilibria in worst-case instances of general bimatrix games, there has also been a series of results on polynomial-time algorithms for computing (approximate)

[4] Awasthi et al. [5] argue that this condition is interesting since in situations where one would want to use an approximate Nash equilibrium for *predicting* how players will play (which is a common motivation for computing a Nash or an approximate Nash equilibrium), without such a condition the approximate equilibrium found might be far from the equilibrium played.

[5] Recall that in an ε-Nash equilibrium, the expected payoff of each player is within ε from her best response payoff; however the mixed strategies may include poorly-performing pure strategies in their support. By contrast, the support of a well-supported ε-approximate equilibrium may only contain strategies whose payoffs fall within ε of the player's best-response payoff.

[6] This is because algorithmically, we can simply try different support sizes in increasing order and stop when we find strategies forming a (well-supported) ε-equilibrium. In other words, given the desired approximation level ε, we can find an ε-approximate equilibrium in time $n^{O((\Delta^2/\varepsilon^2)\log(1+\Delta^{-1})\log n)}$ where Δ is the smallest value greater than or equal to ε such that the game is (ε,Δ) perturbation stable.

equilibria in specific classes of bimatrix games. For example, Bárány et al. [12] considered two-player games with randomly chosen payoff matrices, and showed that with high probability, such games have Nash equilibria with small support. Their result implies that in random two-player games, Nash equilibria can be computed in expected polynomial time. Kannan and Theobald [25] provide an FPTAS for the case where the sum of the two payoff matrices has constant rank and Adsul et al. [1] provide a polynomial time algorithm for computing an exact Nash equilibrium of a rank-1 bimatrix game. Note that it is possible for a game to be stable and yet have high rank, and on the other hand to be unstable and have rank 0; see Appendix B.

Awasthi et al. [5, 6] analyzed the question of finding an approximate Nash equilibrium in games that satisfy stability with respect to approximation. However, their condition is quite restrictive in that it focuses only on games that have the property that all the Nash equilibria are close together, thus eliminating from consideration most common games. By contrast, our perturbation stability notion, which (as mentioned above) can be shown to be a generalization of their notion, captures many more realistic situations. Our upper bound on approximate equilibria can be viewed as generalizing the corresponding result of [5, 6] and it is significantly more challenging technically. Moreover, our lower bounds also apply to the stability notion of [5, 6] and provide the first (nontrivial) results about the interesting range of parameters for that stability notion as well.

In a very different context, for clustering problems, Bilu and Linial [13] analyze maxcut clustering instances with the property that if the distances are perturbed by a multiplicative factor of α, then the optimum does not change. They show that under this condition, one can find the optimum solution in polynomial time so long as $\alpha \geq \min\{n/d_{\min}, \sqrt{nd_{\max}}\}$ where $d_{\min}$ and $d_{\max}$ are the minimum and maximum degrees in the graph, respectively. Following that work, a series of results [7, 11, 10, 3] have analyzed center-based clustering under this stability assumption, including k-median, k-means, and k-center objective functions, with the current best bounds finding the optimum solution when $\alpha \geq 2$ [3]. A survey of center-based clustering, including under various stability conditions, appears in [4]. Our notion of perturbation stability is inspired by this work, but is substantially less restrictive in two respects. First, we require stability only to small perturbations in the input, and second, we do not require the solutions (Nash equilibria) to stay fixed under perturbation, but rather just ask that they have a bounded degree of movement.

A very different direction of work in studying robust equilibrium solutions is given by Aghassi and Bertsimas [2]. Both [2] and the present paper aim to make the notion of equilibria in games more realistic. However, in [2] the equilibrium notion itself is modified. In standard probabilistic game theory, the players choose their strategies to maximize their *expected* payoff given the distribution of the other players' strategies. In the setting of [2], the players aim to maximize the *worst case* possible outcome over possible actions by other players. While our robustness notion refers to the robustness of the (standard probabilistic) equilibrium solution to perturbations of the game, the robustness notion of [2] refers to each player's payoff being robust against uncertainty in other players' behavior—thus it models the players as pessimistic rather than Bayesian agents.

The notion of perturbation stability we consider in our paper is also related to the stability notions considered by Lipton et al. [29] for economic solution concepts. The main focus of their work was understanding whether for a given solution concept or optimization problem *all* instances are stable. In this paper, our main focus is on understanding how rich the class of stable instances is, and what properties one can determine about their structure that can be leveraged to get better algorithms for computing approximate Nash equilibria.[7] We provide the first results showing better algorithms for computing approximate equilibria in such games.

2 Preliminaries

We consider 2-player general-sum n-action bimatrix games. Let R denote the payoff matrix of the row player and C denote the payoff matrix of the column player. If the row player chooses strategy i and the column player chooses strategy j, the payoffs are $R_{i,j}$ and $C_{i,j}$, respectively. We assume all payoffs are scaled to the range $[0,1]$.

A mixed strategy for a player is a probability distribution over the set of his pure strategies. The i-th pure strategy will be represented by the unit vector e_i, that has 1 in the ith coordinate and 0 elsewhere. For a pair (p,q) of mixed strategies, the payoff to the row player is the expected value of a random variable which is equal to $R_{i,j}$ with probability $p_i q_j$. Therefore the payoff to the row player is $p^T Rq$. Similarly the payoff to the column player is $p^T Cq$. Given strategies p and q for the row and column player, we denote by $\operatorname{supp}(p)$ and $\operatorname{supp}(q)$ the support of p and q, respectively.

A Nash equilibrium [31] is a pair (p^*,q^*) of strategies such that no player has an incentive to deviate unilaterally. Since mixed strategies are convex combinations of pure strategies, it suffices to consider only deviations to pure strategies. In particular, a pair (p^*,q^*) of mixed strategies is a *Nash-equilibrium* if for every pure strategy i of the row player we have $e_i^T Rq^* \le {p^*}^T Rq^*$, and for every pure strategy j of the column player we have ${p^*}^T Ce_j \le {p^*}^T Cq^*$. Note that in a Nash equilibrium (p^*,q^*), all rows i in the support of p^* satisfy $e_i^T Rq^* = {p^*}^T Rq^*$ and similarly all columns j in the support of q^* satisfy ${p^*}^T Ce_j = {p^*}^T Cq^*$.

Definition 2.1. A pair of mixed strategies (p,q) is an ε-equilibrium if both players have no more than ε incentive to deviate. Formally, for all rows i, $e_i^T Rq \le p^T Rq + \varepsilon$, and for all columns j, $p^T Ce_j \le p^T Cq + \varepsilon$.

Definition 2.2. A pair (p,q) of mixed strategies is a well-supported ε-equilibrium if for any $i \in \operatorname{supp}(p)$ (i. e., i s.t. $p_i > 0$) we have $e_i^T Rq \ge e_j^T Rq - \varepsilon$, for all j; similarly, for any $i \in \operatorname{supp}(q)$ (i. e., i s.t. $q_i > 0$) we have $p^T Ce_i \ge p^T Ce_j - \varepsilon$, for all j.

Definition 2.3. We say that a bimatrix game G' specified by R', C' is an L_∞ α-perturbation of G specified by R, C if we have $|R_{i,j} - R'_{i,j}| \le \alpha$ and $|C_{i,j} - C'_{i,j}| \le \alpha$ for all $i, j \in \{1, \ldots, n\}$.

Definition 2.4. For two probability distributions q and q', we define the distance between q and q' as the variation distance

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

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

$$d((p,q),(p',q')) = \max[d(p,p'), d(q,q')]\,.$$

[7] Just as in [29], one can show that for the stability conditions we consider in our paper, there exist unstable instances.

It is easy to see that d is a metric. If $d((p,q),(p',q')) \leq \Delta$, then we say that (p',q') is Δ-close to (p,q). Finally, throughout this paper we use "log" to mean the natural logarithm.

3 Stable games

The main notion of stability we introduce and study in this paper requires that any Nash equilibrium in a perturbed game be close to a Nash equilibrium in the original game. This is an especially motivated condition since in many real world situations the entries of the game we analyze are merely based on measurements and thus only approximately reflect the players' payoffs. In order to be useful for prediction, we would like that equilibria in the game we operate with be close to equilibria in the real game played by the players. Otherwise, in games where certain equilibria of slightly perturbed games are far from all equilibria in the original game, the analysis of behavior (or prediction) will be meaningless. We give the formal definition.

Definition 3.1. A game G satisfies the (ε,Δ)-perturbation-stability condition if for any G' that is an L_∞ ε-perturbation of G and for any Nash equilibrium (p,q) in G', there exists a Nash equilibrium (p^*,q^*) in G such that (p,q) is Δ-close to (p^*,q^*).[8]

Observe that fixing ε, a smaller Δ means a stronger condition and a larger Δ means a weaker condition. Every game is $(\varepsilon,1)$-perturbation stable, and as Δ gets smaller, we might expect for the game to exhibit more useful structure.

Another stability condition we consider in this work is approximation stability.[9]

Definition 3.2. A game satisfies the (ε,Δ)-approximation-stability condition if for any ε-equilibrium (p,q) there exists a Nash equilibrium (p^*,q^*) such that (p,q) is Δ-close to (p^*,q^*).

A game satisfies the well-supported (ε,Δ)-approximation-stability condition if for any well-supported ε-equilibrium (p,q) there exists a Nash equilibrium (p^*,q^*) such that (p,q) is Δ-close to (p^*,q^*).

Clearly, if a game satisfies the (ε,Δ)-approximation-stability condition, then it also satisfies the well-supported (ε,Δ)-approximation-stability condition. Interestingly, we show that the perturbation-stability condition is equivalent to the well-supported approximation-stability condition. Specifically, we prove the following.

Theorem 3.3. *A game satisfies the well-supported $(2\varepsilon,\Delta)$-approximation-stability condition if and only if it satisfies the (ε,Δ)-perturbation-stability condition.*

Proof. Consider an $n \times n$ bimatrix game specified by R and C and assume it satisfies the well-supported $(2\varepsilon,\Delta)$-approximation-stability condition; we show it also satisfies the (ε,Δ)-perturbation-stability condition. Consider $R' = R + \Gamma$ and $C' = C + \Lambda$, where $|\Gamma_{i,j}| \leq \varepsilon$ and $|\Lambda_{i,j}| \leq \varepsilon$, for all i,j. Let (p,q) be an arbitrary Nash equilibrium in the new game specified by R' and C'. We will show that (p,q) is a

[8] Note that the entries of the perturbed game are not restricted to the $[0,1]$ interval, and are allowed to belong to $[-\varepsilon,1+\varepsilon]$. This is a proper way to formulate the notion because it implies, for instance, that if G is (ε,Δ) stable to perturbations, then for any $\alpha > 0$, αG is $(\alpha\varepsilon,\Delta)$ stable to perturbations. Theorem 3.3 provides further evidence that this definition is proper.

[9] This definition is inspired by a related notion in clustering initiated by [8, 9] and further studied by, e. g., [34] and [24].

well supported 2ε-approximate Nash equilibrium in the original game specified by R and C. To see this, note that by definition, (since (p,q) is a Nash equilibrium in the game specified by R' and C') we have $e_j^T R'q \leq p^T R'q \equiv v_R$ for all j; therefore $e_j^T Rq + e_j^T \Gamma q \leq v_R$, so $e_j^T Rq \leq v_R + \varepsilon$, for all j. On the other hand we also have $e_i^T Rq = e_i^T R'q - e_i^T \Gamma q \geq v_R - \varepsilon$ for all $i \in \operatorname{supp}(p)$. Therefore, $e_i^T Rq \geq e_j^T Rq - 2\varepsilon$, for all $i \in \operatorname{supp}(p)$ and for all j. Similarly we can show $p^T Ce_i \geq p^T Ce_j - 2\varepsilon$, for all $i \in \operatorname{supp}(q)$ and for all j. This implies that (p,q) is well supported 2ε-approximate Nash in the original game, and so by assumption is Δ-close to a Nash equilibrium of the game specified by R and C. So, this game satisfies the (ε,Δ)-perturbation-stability condition.

In the reverse direction, consider an $n \times n$ bimatrix game specified by R and C and assume it satisfies the (ε,Δ)-stability to perturbations condition. Let (p,q) be an arbitrary well supported 2ε Nash equilibrium in this game. Let us define matrices R' and C' such that $e_i^T R'q = \max_{i'} {e_{i'}}^T Rq - \varepsilon$ for all $i \in \operatorname{supp}(p)$ and $e_i^T R'q \leq \max_{i'} e_{i'}^T Rq - \varepsilon$ for all $i \notin \operatorname{supp}(p)$, $p^T C'e_j = \max_{j'} p^T Ce_{j'} - \varepsilon$ for all $j \in \operatorname{supp}(q)$ and $p^T C'e_{j'} \leq \max_j p^T Ce_{j'} - \varepsilon$ for all $j \notin \operatorname{supp}(q)$. Since (p,q) is a well supported 2ε Nash equilibrium we know this can be done such that $|(R'-R)_{i,j}| \leq \varepsilon$ and $|(C'-C)_{i,j}| \leq \varepsilon$, for all i, j (in particular, we have to add quantities in $[-\varepsilon,\varepsilon]$ to all the elements in rows i of R in the support of p and subtract quantities in $[0,\varepsilon]$ from all the elements in rows i of R not in the support of p; similarly for q). By design, (p,q) is a Nash equilibrium in the game defined by R', C', and from the (ε,Δ)-perturbation-stability condition, we obtain that it is Δ-close to a true Nash equilibrium of the game specified by R and C. Thus, any well supported 2ε Nash equilibrium in the game specified by R and C is Δ-close to a true Nash equilibrium of this game, as desired. $\square$

One can show that well-supported approximation stability is a strict relaxation of the approximation-stability condition. For example, consider the 2×2 bimatrix game

$$R = \begin{bmatrix} 1 & 1 \\ 1-\varepsilon_0 & 1-\varepsilon_0 \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 1-\varepsilon_0 \\ 1 & 1-\varepsilon_0 \end{bmatrix}.$$

For $\varepsilon < \varepsilon_0$ this game satisfies the well-supported $(\varepsilon,0)$-approximation-stability condition, but does not satisfy (ε,Δ)-approximation-stability for any $\Delta < \varepsilon/\varepsilon_0$. To see this note that $e_1^T Rq = 1$, $e_2^T Rq = 1-\varepsilon_0$, $p^T Ce_1 = 1$, and $p^T Ce_2 = 1-\varepsilon_0$ for any p and q. This implies that the only well supported ε-Nash equilibrium is identical to the Nash equilibrium $(1,0)^T,(1,0)^T$, thus the game is well-supported $(\varepsilon,0)$-approximation stable. On the other hand, the pair of mixed strategies (p,q) with $p = (1-\varepsilon/\varepsilon_0, \varepsilon/\varepsilon_0)^T$ and $q = (1-\varepsilon/\varepsilon_0, \varepsilon/\varepsilon_0)^T$ is an ε-Nash equilibrium. The distance between (p,q) and the unique Nash is $\varepsilon/\varepsilon_0$, thus this game is not (ε,Δ)-approximation stable for any $\Delta < \varepsilon/\varepsilon_0$.

Interestingly, approximation stability (which is a restriction of well-supported approximation stability and perturbation stability) is itself a relaxation of the stability condition considered by Awasthi et al. [5] which requires that all approximate equilibria be contained in a ball of radius Δ around a single Nash equilibrium. In this direction, we define the *strong version* of the stability conditions given in Definitions 3.1 and 3.2 to be a reversal of quantifiers that asks that there be a single (p^*,q^*) such that each relevant (p,q) (equilibrium in an ε-perturbed game, ε-approximate equilibrium, or well-supported ε-approximate equilibrium) is Δ-close to (p^*,q^*). Here we give the formal definition.

Definition 3.4. A game G satisfies the strong (ε,Δ)-perturbation-stability condition if there exists (p^*,q^*) a Nash equilibrium of G such that for any G' that is an L_∞ ε-perturbation of G we have that any Nash equilibrium in G' is Δ-close to (p^*,q^*).

A game G satisfies the strong (well-supported) (ε,Δ)-approximation-stability condition if there exists (p^*,q^*) a Nash equilibrium of G such that any (well supported) ε-equilibrium (p,q) is Δ-close to (p^*,q^*).

It is immediate from its proof that Theorem 3.3 applies to the strong versions of the definitions as well. We also note that our generic upper bounds in Section 4 apply to the most relaxed version (perturbation-stability) and our generic lower bound in Section 5 apply to the most stringent version (strong approximation stability).

Range of parameters. As shown in [5], if a game $\mathcal{G}$ satisfies the strong (ε,Δ)-approximation-stability condition and has a non-trivial Nash equilibrium (an equilibrium in which the players do not both have full support), then we must have $\Delta \geq \varepsilon$. We can show that if a game $\mathcal{G}$ satisfies (ε,Δ)-approximation-stability and if the union of all Δ-balls around all Nash equilibria does not cover the whole space,[10] then we must have $3\Delta \geq \varepsilon$; see Lemma C.1 and Lemma C.2 in Appendix C. In Section 5 we further discuss the meaningful range of parameters from the point of view of the equilibrium computation problem.

Examples. Variants of many classic games including the public goods game, matching pennies, and identical interest games are stable. As an example, consider the following modified identical interest game. Both players have n available actions. The first action is to stay home, and the other actions correspond to $n-1$ different possible meeting locations. If a player chooses action 1 (stay home), his payoff is $1/2$ no matter what the other player is doing. If the player chooses to go out to a meeting location, his payoff is 1 if the other player is there as well and it is 0 otherwise. This game has n pure equilibria (all (e_i,e_i)) and $\binom{n}{2}$ mixed equilibria (all $(1/2e_i+1/2e_j, 1/2e_i+1/2e_j)$) and it is well-supported $(\varepsilon,2\varepsilon)$-approximation stable for all $\varepsilon < 1/6$. Note that it does not satisfy *strong* (well-supported) stability because it has multiple very distinct equilibria. For further examples see Lemma 5.1 in Section 5, as well as Appendix B.

4 Equilibria in stable games

In this section we show we can leverage the structure implied by perturbation stability to improve over the best known generic bound of [28]. We start by considering ε and Δ as given.

Theorem 4.1. *Let $0 \leq \varepsilon \leq \Delta \leq 1$. Consider a game with at most $n^{O((\Delta/\varepsilon)^2)}$ Nash equilibria which satisfies the well-supported (ε,Δ)-approximation-stability condition (or the $(\varepsilon/2,\Delta)$-perturbation-stability condition). Then there exists a well supported ε-equilibrium where each player's strategy has support $O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1}) \log n)$.*

This improves by a factor $O(1/(\Delta^2 \log(1+\Delta^{-1})))$ in the exponent over the bound of [28] for games satisfying these conditions (and reduces to the bound of [28] in the worst-case when $\Delta = 1$).

[10] If the union of all Δ-balls around all Nash equilibria does cover the whole space, this is an easy case from our perspective. Any (p,q) would be a ε-equilibria.

Remark. Note that any game satisfying well-supported (ε,Δ)-approximation-stability for $\Delta<\varepsilon$ also satisfies well-supported $(\varepsilon,\varepsilon)$-approximation-stability. Therefore, the condition "$\varepsilon\le\Delta$" can be removed from Theorem 2 by replacing the number of equilibria allowed with $n^{O(\max\{1,(\Delta/\varepsilon)^2\})}$ and the support-size guarantee with $O(\max\{1,(\Delta/\varepsilon)^2\}\log(1+\max\{\varepsilon,\Delta\}^{-1})\log n)$.

Proof idea. We start by showing that any Nash equilibrium (p^*,q^*) of G must be highly concentrated. In particular, we show that for each of p^*,q^*, any portion of the distribution with substantial L_1 norm (having total probability at least 8Δ) must also have high L_2 norm. Specifically, the ratio of L_2 norm to L_1 norm must be $\Omega((\varepsilon/\Delta)(\log n)^{-1/2})$. This in turn can be used to show that each of p^*,q^* has all but at most 8Δ of its total probability mass concentrated in a set (which we call the high part) of size $O((\Delta/\varepsilon)^2\log(1+1/\Delta)\log n)$. Once the desired concentration is proven, we can then perform a version of the [28] sampling procedure on the low parts of p^* and q^* (with an accuracy of only ε/Δ) to produce overall an ε-approximate equilibrium of support only a constant factor larger. The primary challenge in this argument is to prove that p^* and q^* are concentrated.[11] This is done through our key lemma, Lemma 4.2 below. In particular, Lemma 4.2 can be used to show that if p^* (or q^*) had a portion with substantial L_1 norm and low L_2 norm, then there must exist a deviation from p^* (or q^*) that is far from *all* equilibria and yet is a well-supported approximate-Nash equilibrium, violating the stability condition. Proving the existence of such a deviation is challenging because of the large number of equilibria that may exist. Lemma 4.2 synthesizes the key points of this argument.

Lemma 4.2. *Let $0\le\varepsilon\le\Delta\le 1$. Let $\tilde p$ be an arbitrary distribution over $\{1,2,\dots,n\}$. Let $S=c(\Delta/\varepsilon)^2\log n$ for $c\ge 56^2$, and fix $\beta\le 1$ such that $1-\beta\ge 8\Delta$. Assume that the entries of $\tilde p$ can be partitioned into two sets H and L such that*

$$\|\tilde p_L\|_1=1-\beta\,,\qquad \|\tilde p_H\|_1=\beta\,,\quad \textit{and}\quad \|\tilde p_L\|_2^2\le(1-\beta)^2/S\,.$$

Let us fix k_1 n-dimensional vectors $v^{(1)},\dots,v^{(k_1)}$ with entries in $[-1,1]$ and k_2 distributions $p^{(1)},\dots,p^{(k_2)}$, where $k_1=n^2$ and $k_2\le n^{c'(\Delta/\varepsilon)^2}$ for $c'=(2/225)c-\log(6)/\log(n)$. Then there exists $\tilde p'$ with $\operatorname{supp}(\tilde p')\subseteq\operatorname{supp}(\tilde p)$ *such that*

(1) $d(\tilde p,\tilde p')=3\Delta$,

(2) $\tilde p'\cdot v^{(i)}\le\tilde p\cdot v^{(i)}+\varepsilon$ for all $i\in\{1,\dots,k_1\}$, and

(3) $d(\tilde p',p^{(i)})>d(\tilde p,p^{(i)})-\Delta$ for all $i\in\{1,\dots,k_2\}$.

Proof. We show the desired result by using the probabilistic method. We randomly partition the entries in $\tilde p_L$ into two sets A and B, moving 3Δ probability mass from A to B to create $\tilde p'$ (moving an amount of probability to/from $\tilde p_i$ proportional to the value of $\tilde p_i$); for $i\in H$, we let $\tilde p'_i=\tilde p_i$. We next carefully use concentration bounds together with both our L_1 and L_2 bounds on $\tilde p_L$ to show that with high probability, conditions (2) and (3) are both satisfied.

[11] We note that [5] prove an upper bound for the strong approximation-stability condition using the same concentration idea. However, proving the desired concentration is significantly more challenging in our case since we deal with many equilibria.

Formally, let us define the random variable $X_i = 1$ with probability $1/2$ and $X_i = 0$ with probability $1/2$. Define

$$\tilde{p}'_i = \begin{cases} \tilde{p}_i & \text{for } i \in H, \\ \tilde{p}_i + \frac{3\Delta \tilde{p}_i X_i}{\sum_{i\in L} \tilde{p}_i X_i} - \frac{3\Delta \tilde{p}_i (1-X_i)}{\sum_{i\in L} \tilde{p}_i (1-X_i)} & \text{for } i \in L. \end{cases}$$

We have

$$\mathbf{E}\left[\sum_{i\in L} \tilde{p}_i X_i\right] = \frac{1-\beta}{2}.$$

By applying McDiarmid's inequality (see Theorem A.1 in Appendix A) and using the fact that $\|\tilde{p}_L\|_2^2 \le (1-\beta)^2/S$, we obtain that with probability at least $3/4$ we have:

$$\left|\sum_{i\in L} \tilde{p}_i X_i - \frac{1-\beta}{2}\right| \le \frac{1-\beta}{12} \qquad \text{which also implies} \qquad \left|\sum_{i\in L} \tilde{p}_i (1-X_i) - \frac{1-\beta}{2}\right| \le \frac{1-\beta}{12}. \tag{4.1}$$

Assume that this happens. In this case, $\tilde{p}'$ is a legal mixed strategy for the row player and by construction we have $d(\tilde{p}, \tilde{p}') = 3\Delta$.

Let v be an arbitrary vector in $\{v^{(1)}, \ldots, v^{(k_1)}\}$. We have:

$$\tilde{p}' \cdot v = \tilde{p} \cdot v + 3\Delta \left(\frac{\sum_{i\in L} \tilde{p}_i X_i v_i}{\sum_{i\in L} \tilde{p}_i X_i} - \frac{\sum_{i\in L} \tilde{p}_i (1-X_i) v_i}{\sum_{i\in L} \tilde{p}_i (1-X_i)}\right).$$

Define

$$Z_1 = \sum_{i\in L} \tilde{p}_i X_i v_i, \quad Z_2 = \sum_{i\in L} \tilde{p}_i X_i, \quad Z_3 = \sum_{i\in L} \tilde{p}_i (1-X_i) v_i, \quad Z_4 = \sum_{i\in L} \tilde{p}_i (1-X_i),$$

so we have:

$$\tilde{p}' \cdot v = \tilde{p} \cdot v + 3\Delta \left(\frac{Z_1}{Z_2} - \frac{Z_3}{Z_4}\right).$$

Using McDiarmid's inequality we get that with probability at least $1 - 1/n^3$, each of the quantities Z_1, Z_2, Z_3, Z_4 is within $((1-\beta)/28)(\varepsilon/\Delta)$ of its expectation; we are using here the fact that the value of X_i can change any one of the quantities by at most $\tilde{p}_i$, so the exponent in the McDiarmid bound is

$$\frac{-\left(\frac{\varepsilon}{\Delta}\right)^2 \left(\frac{1-\beta}{28}\right)^2}{\sum_{i\in L} \tilde{p}_i^2} \le -\left(\frac{c}{28^2}\right) \log n.$$

Also, we have

$$\mathbf{E}[Z_2] = \mathbf{E}[Z_4] = \frac{1-\beta}{2} \quad \text{and} \quad \mathbf{E}[Z_1] = \mathbf{E}[Z_3].$$

So, we get that with probability at least $1-1/n^3$ we have

$$\begin{aligned}
\tilde{p}'\cdot v-\tilde{p}\cdot v &\leq 3\Delta\left(\frac{\mathbf{E}[Z_1]+\frac{(1-\beta)\varepsilon}{28\Delta}}{\frac{1-\beta}{2}-\frac{(1-\beta)\varepsilon}{28\Delta}}-\frac{\mathbf{E}[Z_3]-\frac{(1-\beta)\varepsilon}{28\Delta}}{\frac{1-\beta}{2}+\frac{(1-\beta)\varepsilon}{28\Delta}}\right)\\
&= 3\Delta\left(\frac{(\frac{2}{1-\beta})\mathbf{E}[Z_1]+\frac{\varepsilon}{14\Delta}}{1-\frac{\varepsilon}{14\Delta}}-\frac{(\frac{2}{1-\beta})\mathbf{E}[Z_3]-\frac{\varepsilon}{14\Delta}}{1+\frac{\varepsilon}{14\Delta}}\right)\\
&= 3\Delta\left[\left(\frac{2}{1-\beta}\mathbf{E}[Z_1]+\frac{\varepsilon}{14\Delta}\right)\left(1+\frac{\varepsilon}{14\Delta}\right)\right.\\
&\qquad\left.-\left(\frac{2}{1-\beta}\mathbf{E}[Z_3]-\frac{\varepsilon}{14\Delta}\right)\left(1-\frac{\varepsilon}{14\Delta}\right)\right]\left(\frac{1}{1-(\frac{\varepsilon}{14\Delta})^2}\right)\\
&\leq 3.1\Delta\left[\left(\frac{2}{1-\beta}\mathbf{E}[Z_1]+\frac{\varepsilon}{14\Delta}\right)\left(1+\frac{\varepsilon}{14\Delta}\right)\right.\\
&\qquad\left.-\left(\frac{2}{1-\beta}\mathbf{E}[Z_3]-\frac{\varepsilon}{14\Delta}\right)\left(1-\frac{\varepsilon}{14\Delta}\right)\right]\\
&= 3.1\Delta\left(\frac{2}{1-\beta}\left(\mathbf{E}[Z_1]+\mathbf{E}[Z_3]\right)\left(\frac{\varepsilon}{14\Delta}\right)+\frac{\varepsilon}{7\Delta}\right).
\end{aligned}$$

Finally, using the fact that $Z_1+Z_3\leq\sum_{i\in L}\tilde{p}_i=1-\beta$, we get

$$\tilde{p}'\cdot v-\tilde{p}\cdot v \quad\leq\quad 3.1\Delta\left(\frac{\varepsilon}{7\Delta}+\frac{\varepsilon}{7\Delta}\right)$$

yielding the desired bound $\tilde{p}'\cdot v\leq\tilde{p}\cdot v+\varepsilon$. Applying the union bound over all $i\in\{1,\ldots,k_1\}$ we obtain that the probability that there exists v in $v^{(1)},\ldots,v^{(k_1)}$ such that $\tilde{p}'\cdot v\geq\tilde{p}\cdot v+\varepsilon$ is at most $1/3$.

Consider an arbitrary distribution p in $\{p^{(1)},\ldots,p^{(k_2)}\}$. Assume that $p=\tilde{p}+g$. By definition, we have:

$$\begin{aligned}
d(p,\tilde{p}') &= \frac{1}{2}\sum_{i\in L}\left|\tilde{p}_i+g_i-\tilde{p}_i-\frac{3\Delta\tilde{p}_iX_i}{\sum_{i\in L}\tilde{p}_iX_i}+\frac{3\Delta\tilde{p}_i(1-X_i)}{\sum_{i\in L}^{n}\tilde{p}_i(1-X_i)}\right|+\frac{1}{2}\sum_{i\in H}|g_i|\\
&= \frac{1}{2}\sum_{i\in L}\left|g_i-\frac{3\Delta\tilde{p}_iX_i}{\sum_{i\in L}\tilde{p}_iX_i}+\frac{3\Delta\tilde{p}_i(1-X_i)}{\sum_{i\in L}\tilde{p}_i(1-X_i)}\right|+\frac{1}{2}\sum_{i\in H}|g_i|\\
&\geq \frac{1}{2}\sum_{i\in L}\left|g_i-\frac{6\Delta\tilde{p}_iX_i}{1-\beta}+\frac{6\Delta\tilde{p}_i(1-X_i)}{1-\beta}\right|-\frac{1}{2}\sum_{i\in L}\left[\frac{6\Delta\tilde{p}_iX_i}{5(1-\beta)}+\frac{6\Delta\tilde{p}_i(1-X_i)}{7(1-\beta)}\right]+\frac{1}{2}\sum_{i\in H}|g_i|\\
&\geq \frac{1}{2}\sum_{i\in L}\left|g_i-\frac{6\Delta\tilde{p}_iX_i}{1-\beta}+\frac{6\Delta\tilde{p}_i(1-X_i)}{1-\beta}\right|-\frac{1}{2}\sum_{i\in L}\frac{6\Delta\tilde{p}_i}{5(1-\beta)}+\frac{1}{2}\sum_{i\in H}|g_i|\\
&= \frac{1}{2}\sum_{i\in L}\left|g_i-\frac{6\Delta\tilde{p}_iX_i}{1-\beta}+\frac{6\Delta\tilde{p}_i(1-X_i)}{1-\beta}\right|-\frac{3\Delta}{5}+\frac{1}{2}\sum_{i\in H}|g_i|,
\end{aligned}$$

where the first inequality follows from applying relation (4.1) to the denominators and the fact that the sum of the two denominators equals $1-\beta$. The last equality also follows from the fact that $\|\tilde{p}_L\|_1=1-\beta$.

Let us denote by

$$Z = \frac{1}{2}\sum_{i\in L}\left|g_i - \frac{6\Delta\tilde{p}_i X_i}{1-\beta} + \frac{6\Delta\tilde{p}_i(1-X_i)}{1-\beta}\right|.$$

We have:

$$E[Z] = \frac{1}{2}\sum_{i\in L}\left[\frac{1}{2}\left|g_i - \frac{6\Delta\tilde{p}_i}{1-\beta}\right| + \frac{1}{2}\left|g_i + \frac{6\Delta\tilde{p}_i}{1-\beta}\right|\right] \geq \frac{1}{2}\sum_{i\in L}|g_i|,$$

therefore

$$E[d(p,\tilde{p}')] \geq E[Z] + \frac{1}{2}\sum_{i\in H}|g_i| - \frac{3\Delta}{5} \geq \frac{1}{2}\sum_{i\in L}|g_i| + \frac{1}{2}\sum_{i\in H}|g_i| - \frac{3\Delta}{5} = d(p,\tilde{p}) - \frac{3\Delta}{5}.$$

We can now apply McDiarmid's inequality (see Theorem A.1) to argue that with high probability Z is within $2\Delta/5$ of its expectation. Note that $c_i = 6\Delta\tilde{p}_i/(1-\beta)$. Therefore:

$$\Pr\{|Z-\mathbf{E}[Z]| \geq 2\Delta/5\} \leq 2e^{-2\Delta^2(1-\beta)^2/(225\sum_{i\in L}\tilde{p}_i^2\Delta^2)} \leq 2e^{-(2/225)S} \leq \frac{1}{3k_2},$$

where the last inequality follows from the definition of c'. This then implies that

$$\Pr\left\{d(p,\tilde{p}') \leq d(p,\tilde{p}') - \Delta\right\} \leq \frac{1}{3k_2}.$$

By the union bound we get that the probability that there exists a p in $\{p^{(1)},\ldots,p^{(k_2)}\}$ such that $d(p,\tilde{p}') \leq d(p,\tilde{p}') - \Delta$ is at most $1/3$. Summing up overall all possible events we get that there is a non-zero probability of (1), (2), (3) happening, as desired.

□

Proof of Theorem 4.1. Let (p^*,q^*) be an arbitrary Nash equilibrium. We show that each of p^* and q^* are highly concentrated, meaning that all but at most 8Δ of their total probability mass is concentrated in a set of size $O((\Delta/\varepsilon)^2\log(1+1/\Delta)\log n)$. Let's consider p^* (the argument for q^* is similar). We begin by partitioning it into its *heavy* and *light* parts. Specifically, we greedily remove the largest entries of p^* and place them into a set H (the heavy elements) until either

(a) the remaining entries L (the light elements) satisfy the condition that $\forall i \in L$, $\Pr[i] \leq (1/S)\Pr[L]$ for S as in Lemma 4.2, or

(b) $\Pr[H] \geq 1-8\Delta$,

whichever comes first. Using the fact that the game satisfies the well-supported (ε,Δ)-approximation-stability condition, we will show that case (a) cannot occur first, which will imply that p^* is highly concentrated.

In the following, we denote β as the total probability mass over H. Assume by contradiction that case (a) occurs first. Note that we have $\|p_L\|_1 = 1-\beta$, $\|p_H\|_1 = \beta$, and

$$\sum_{i\in L}(p_i)^2 \leq \frac{1}{S}\sum_{i\in L}p_i\sum_{i\in L}p_i = \frac{1}{S}(1-\beta)^2,$$

so $\|p_L\|_2^2 \le (1-\beta)^2/S$. Let $v_{i,j} = C(e_i - e_j)$. Since (p^*, q^*) is a Nash equilibrium we know that $p^* \cdot v_{i,j} \le 0$ for all i and for all $j \in \operatorname{supp}(q^*)$.

By Lemma 4.2 there exists $\tilde{p}'$ such that (1) $d(p^*, \tilde{p}') = 3\Delta$, (2) $\tilde{p}' \cdot v_{i,j} \le p^* \cdot v_{i,j} + \varepsilon$ for all i and for all $j \in \operatorname{supp}(q^*)$ and (3) $d(\tilde{p}', p^{(i)}) > d(p^*, p^{(i)}) - \Delta$ for all Nash equilibria $p^{(i)}$. By (2) we have that $\tilde{p}' \cdot v_{i,j} \le \varepsilon$ for all i and for all $j \in \operatorname{supp}(q^*)$, which implies that $(\tilde{p}', q^*)$ is a well supported ε-approximate equilibrium (since by (2) the column player has at most an ε incentive to deviate and since $\operatorname{supp}(\tilde{p}') \subseteq \operatorname{supp}(p^*)$ we know that the row player has no incentive to deviate). By (1) we also have that $(\tilde{p}', q^*)$ is 3Δ-far from (p^*, q^*). We now use (3) to show that $(\tilde{p}', q^*)$ is Δ-far from all the other equilibria as well. Let p be such an equilibrium. Note that if $d(p, p^*) > 4\Delta$, then clearly, by the triangle inequality $d(p, \tilde{p}') > \Delta$. If $d(p, p^*) < 2\Delta$, clearly, by the triangle inequality, $d(p, \tilde{p}') > \Delta$. Finally if $d(p, p^*) \in [2\Delta, 4\Delta]$, then by (3), we that $d(p, \tilde{p}') > \Delta$, as desired.

Overall we get that $(\tilde{p}', q^*)$ is a well-supported ε-approximate equilibrium that is Δ-far from all the other equilibria of the game. This contradicts the well supported (ε, Δ)-approximation-stability condition, as desired.

Thus, case (b) occurs first, which implies that p^* is highly concentrated. In particular, the fact that case (a) has not yet occurred implies that the greedy construction of H has been able at each step to choose an entry i such that $\Pr[i] > (1/S)\Pr[L]$; so, when $\beta = \Pr[H] \ge 1 - 8\Delta$, H has at most $S \log(1 + (8\Delta)^{-1})$ elements. The key idea now is that since $1 - \beta \le 8\Delta$, we can apply the sampling argument of [28] to L with accuracy parameter $O(\varepsilon/\Delta)$ and then union the result with H.

Specifically, let us decompose p^* as

$$p^* = \beta p_H + (1-\beta) p_L .$$

Applying the sampling argument of [28] to p_L, we have that by sampling a multiset $\mathcal{X}$ of S elements from $L = \operatorname{supp}(p_L)$, we are guaranteed that for any column e_j, we have

$$\left|(U_{\mathcal{X}})^T C e_j - p_L^T C e_j\right| \le (\varepsilon/8\Delta) ,$$

where $U_{\mathcal{X}}$ is the uniform distribution over $\mathcal{X}$. This means that for $\tilde{p} = \beta p_H + (1-\beta) U_S$, all columns e_j satisfy

$$|p^{*T} C e_j - \tilde{p}^T C e_j| \le \varepsilon/2 .$$

We have thus found the row portion of an ε-equilibrium with support of size $S \log(1 + 1/(8\Delta))$ as desired. □

Corollary 4.3. *Let $0 \le \varepsilon \le \Delta \le 1$. Let G be a game with at most $n^{O((\Delta/\varepsilon)^2)}$ Nash equilibria satisfying the well-supported (ε, Δ)-approximation-stability condition (or the $(\varepsilon/2, \Delta)$-perturbation-stability condition). Then*

(1) *given G we can find a well-supported ε-equilibrium (p, q) of G in time $n^{O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1}) \log n)}$;*

(2) *given G', an L_∞ $\varepsilon/6$-perturbation of G, we can find a well supported ε-equilibrium (p, q) of G in time $n^{O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1}) \log n)}$.*

In both cases, (p, q) is Δ-close to a Nash equilibrium (p^*, q^*) of G.

Proof. (1) By Theorem 4.1, we can simply try all supports of size $O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1})\log n)$ and for each of them write an LP to search for a well-supported ε-Nash equilibrium.

(2) By Theorem 4.1, G has a well-supported $\varepsilon/3$-Nash equilibrium with support of size $O((\Delta/\varepsilon)^2 \cdot \log(1+\Delta^{-1})\log n)$. Since G' is an L_∞ $\varepsilon/6$-perturbation of G, then this is also a well-supported $2\varepsilon/3$-Nash equilibrium of G'. Thus by trying all supports of size $O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1})\log n)$ in G' we can find a well-supported $2\varepsilon/3$-Nash equilibrium of G'. Since G is an L_∞ $\varepsilon/6$-perturbation of G', this will be a well-supported ε-Nash equilibrium of G. □

Corollary 4.3 improves by a factor $O(1/(\Delta^2 \log(1+\Delta^{-1})))$ in the exponent over the bound of [28] for games satisfying this condition. The most interesting range of improvements happens when ε is a function on n and Δ is a function of ε; e. g., $\varepsilon = 1/\sqrt{n}$, $\Delta = 10\varepsilon$—in this case we obtain an improvement of $O(n/\log(n))$ in the exponent over the bound of [28].

The proof of Theorem 4.1 also implies an interesting structural result, namely, that *each* Nash equilibrium of such a game is close to a pair of strategies of small support and by the triangle inequality, the same will happen for any perturbation of G. We state the result.

Theorem 4.4. *Let $0 \le \varepsilon \le \Delta \le 1$. Consider a game G with at most $n^{O((\Delta/\varepsilon)^2)}$ Nash equilibria which satisfies well-supported (ε,Δ)-approximation-stability condition (or the $(\varepsilon/2,\Delta)$-perturbation-stability condition). Then it must be the case that*

(1) *any Nash equilibrium in G is 8Δ-close to a pair of mixed strategies, each with support of size at most $O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1})\log n)$;*

(2) *for any game G' with L_∞ distance $\varepsilon/2$ from G, any Nash equilibrium in G' is 9Δ-close to a pair of mixed strategies, each with support of size $O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1})\log n)$.*

Proof. Consider an arbitrary Nash equilibrium (p^*,q^*) of G. The proof of Theorem 4.1 shows that the supports of p^* and q^* can be partitioned into heavy and light sets so that $p^* = p_H + p_L$ and $q^* = q_H + q_L$, where p_H and q_H each have probability mass at least $1-8\Delta$ and each has support size $O((\Delta/\varepsilon)^2 \log(1+\Delta^{-1})\log n)$. Thus, (p^*,q^*) is 8Δ-close to $(p_H/|p_H|, q_H/|q_H|)$. Part (2) then uses that fact that any Nash equilibrium in G' must be Δ-close to some Nash equilibrium in G due to perturbation-stability. □

So far in Theorem 4.1 and Corollary 4.3 we have considered ε and Δ fixed. It is also interesting to consider games where the stability conditions hold uniformly for all sufficiently small ε. We call such games *uniformly stable.*

Definition 4.5. Consider $t \ge 1$. We say that a game is *t-uniformly perturbation-stable* (or t-uniformly well-supported approximation-stable) if there exists $\varepsilon_0 = 1/\mathrm{poly}(n)$ such that for all $\varepsilon \le \varepsilon_0$, G satisfies $(\varepsilon, t\varepsilon)$ stability to perturbations (or the well-supported $(\varepsilon, t\varepsilon)$-approximation-stability).

For games satisfying the t-uniform perturbation-stability condition with $t = O(\mathrm{poly}(\log(n)))$ we can find $1/\mathrm{poly}(n)$-approximate equilibria in $n^{\mathrm{poly}(\log(n))}$ time, and more generally ε-approximate equilibria in $n^{\log(1/\varepsilon)\mathrm{poly}(\log(n))}$ time, thus achieving a FQPTAS. This provides a dramatic improvement over the best worst-case bounds known.

Corollary 4.6. *Let $t = O(\mathrm{poly}(\log(n)))$. Then the following hold.*

(1) There is a FQPTAS to find approximate equilibria in games satisfying the t-uniform well-supported approximation-stability condition (or the t-uniform perturbation-stability condition) with at most $n^{O(t^2)}$ Nash equilibria.

(2) Games satisfying the t-uniform well-supported approximation-stability condition (or the t-uniform perturbation-stability condition) with at most $n^{O(t^2)}$ Nash equilibria have the property that for any Δ each equilibrium is 8Δ-close to a pair of mixed strategies each with support of size $O(t^2 \log(1+\Delta^{-1}) \log n)$; moreover, for any L_∞-perturbation of magnitude Δ/t of such games, it must be the case that any Nash equilibrium in G' is 9Δ-close to a pair of mixed strategies each with support of size $O(t^2 \log(1+\Delta^{-1}) \log n)$.

Corollary 4.6 is especially interesting because the results of [16] prove that it is PPAD-hard to find $1/\text{poly}(n)$-approximate equilibria in general bimatrix games. Our results shows that under the (widely believed) assumption that PPAD is not contained in quasi-polynomial time [17], such uniformly stable games are inherently easier for computation of approximate equilibria than general bimatrix games. Moreover, variants of many games appearing commonly in experimental economics including the public goods game and identical interest game [22] satisfy this condition (see Appendix B).

5 Converting the general case to the stable case

In this section we show that computing an ε-equilibrium in a game satisfying the strong $(\varepsilon, \Theta(\varepsilon^{1/4}))$-approximation-stability condition is as hard as computing a $\Theta(\varepsilon^{1/4})$-equilibrium in a general game. For our reduction, we show that any general game can be embedded into one having the strong $(\varepsilon, \Theta(\varepsilon^{1/4}))$-approximation-stability property such that an ε equilibrium in the new game yields a $\Theta(\varepsilon^{1/4})$ equilibrium in the original game. Since all the notions of stability considered in this paper generalize the strong approximation-stability condition, the main lower bound in this section (Theorem 5.2) applies to these notions as well.

We start by stating a key lemma that shows the existence of a family of modified matching pennies games that satisfy strong approximation stability and have certain properties that will be helpful in proving our main lower bound.

Lemma 5.1. *Assume that $\Delta \le 1/10$. Consider the games defined by the following matrices.*

$$R = \begin{bmatrix} 1+\alpha_{1,1} & 1+\alpha_{1,2} & \dots & 1+\alpha_{1,n} & 0 \\ 1+\alpha_{2,1} & 1+\alpha_{2,2} & \dots & 1+\alpha_{2,n} & 0 \\ & \dots & \dots & & \\ 1+\alpha_{n,1} & 1+\alpha_{n,2} & \dots & 1+\alpha_{n,n} & 0 \\ 0 & 0 & \dots & 0 & 2\Delta \end{bmatrix}, \quad C = \begin{bmatrix} \gamma_{1,1} & \gamma_{1,2} & \dots & \gamma_{1,n} & 1 \\ \gamma_{2,1} & \gamma_{2,2} & \dots & \gamma_{2,n} & 1 \\ & \dots & \dots & & \\ \gamma_{n,1} & \gamma_{n,2} & \dots & \gamma_{n,n} & 1 \\ 2\Delta & 2\Delta & \dots & 2\Delta & 0 \end{bmatrix},$$

where $\alpha_{i,j} \in [-\Delta, 0]$ and $\gamma_{i,j} \in [0, \Delta]$ for all i, j. Each such game satisfies strong $(\Delta^2, 4\Delta)$-approximation-stability. Moreover if (p,q) is a Δ^2-Nash equilibrium, then we must have

$$\Delta/2 \le p_1 + \dots + p_n \le 4\Delta \quad \textit{and} \quad \Delta/2 \le q_1 + \dots + q_n \le 4\Delta.$$

Since the proof of Lemma 5.1 is somewhat tedious we defer its proof at the end of this section. We now present the main result of this section.

Theorem 5.2. *Computing an ε-equilibrium in a game satisfying the strong $(\varepsilon, 8\varepsilon^{1/4})$-approximation-stability condition is as hard as computing an $(8\varepsilon)^{1/4}$-equilibrium in a general game.*

Proof. The main idea is to construct a linear embedding of any given game into a larger, stable game in such a way that incentives are not magnified too greatly when approximate equilibria of the new game are translated back into strategies for the original game. In particular, consider $\Delta = (8\varepsilon)^{1/4}$ and consider a general game with $n \times n$ payoff matrices R and C. Let us construct a new game with $(n+1) \times (n+1)$ payoff matrices R' and C' defined as follows.

$$R' = \begin{bmatrix} \mathbf{1}_{n,n} - (\Delta/2)\mathbf{1}_{n,n} + (\Delta/2)R & \mathbf{0}_{n,1} \\ \mathbf{0}_{1,n} & 2\Delta \end{bmatrix}$$

and

$$C' = \begin{bmatrix} (\Delta/2)\mathbf{1}_{n,n} + (\Delta/2)C & \mathbf{1}_{n,1} \\ 2\Delta\mathbf{1}_{1,n} & 0 \end{bmatrix}.$$

(For $s, r > 0$, the matrix $\mathbf{1}_{s,r}$ is the $s \times r$ matrix with all entries set to 1 and the matrix $\mathbf{0}_{s,r}$ is the $s \times r$ matrix with all entries set to 0.) By Lemma 5.1, the new game defined by R' and C' satisfies the strong $(\Delta^2, 4\Delta)$-approximation-stability condition, which in turn implies satisfying $(\varepsilon, 8\varepsilon^{1/4})$-approximation-stability (since $\varepsilon \leq \Delta^2$ and $4(8)^{1/4} \leq 8$). We show next that any $\Delta^4/8$-equilibrium in this new game (defined by R' and C') induces a Δ-equilibrium in the original game (defined by R and C). Since $\Delta = (8\varepsilon)^{1/4}$, this implies the desired result.

Let (p, q) be a $\Delta^4/8$-equilibrium in the new game. By Lemma 5.1 (since $\Delta^4/8 \leq \Delta^2$), p must have $\beta\Delta$ probability mass in the first n rows and q must have $\alpha\Delta$ probability mass in the first n columns, where $\alpha, \beta \in [1/2, 4]$. Let p_f, q_f denote p restricted to the first n rows and q restricted to the first n columns. Let $\tilde{p}_f = p_f/|p_f|$ and $\tilde{q}_f = q_f/|q_f|$, where $|p_f| = \beta\Delta$ and $|q_f| = \alpha\Delta$. We show that that $(\tilde{p}_f, \tilde{q}_f)$ is a Δ-equilibrium in the original game defined by R and C. We prove this by contradiction. Assume this is not the case. Assume first that the row player has an Δ incentive to deviate. There must exist e_i such that

$$e_i^T R\tilde{q}_f > \tilde{p}_f R\tilde{q}_f + \Delta.$$

Multiplying both sides by $\alpha\beta\Delta^3/2$ and using the fact that $\alpha\beta \geq 1/4$ we obtain the following inequality.

$$\beta\Delta e_i^T (\Delta/2)Rq_f > p_f(\Delta/2)Rq_f + \Delta^4/8. \tag{5.1}$$

We clearly have $\beta\Delta e_i^T(\mathbf{1}_{n,n} - (\Delta/2)\mathbf{1}_{n,n})q_f = p_f^T(\mathbf{1}_{n,n} - (\Delta/2)\mathbf{1}_{n,n})q_f$, and by adding this quantity as well as $p_{n+1}(2\Delta)q_{n+1}$ to both sides of inequality (5.1) we get

$$\begin{aligned} &\beta\Delta e_i^T(\mathbf{1}_{n,n} - (\Delta/2)\mathbf{1}_{n,n} + (\Delta/2)R)q_f + p_{n+1}(2\Delta)q_{n+1} \\ &\qquad > p_f(\mathbf{1}_{n,n} - (\Delta/2)\mathbf{1}_{n,n} + (\Delta/2)R)q_f + p_{n+1}(2\Delta)q_{n+1} + \Delta^4/8, \end{aligned}$$

which implies

$$(\beta\Delta e_i + p_{n+1}e_{n+1})^T R' q > p^T R' q + \Delta^4/8\,.$$

Therefore there exists a deviation for the row player (namely moving all $\beta\Delta$ probability mass from rows $1,2,\ldots,n$ onto row i), yielding a benefit of $\Delta^4/8$ to the row player. This contradicts the assumption that (p,q) is an $\Delta^4/8$-equilibrium in the new game, as desired.

Assume now that the column player has an Δ incentive to deviate. There must exist e_j such that

$$\tilde{p}_f^T C e_j > \tilde{p}_f^T C \tilde{q}_f + \Delta\,.$$

Multiplying both sides by $\alpha\beta\Delta^3/2$ and using the fact that $\alpha\beta \geq 1/4$ we get:

$$p_f^T(\Delta/2)C(\alpha\Delta e_j) > p_f^T(\Delta/2)Cq_f + \Delta^4/8\,.$$

We have $p_f^T(\Delta/2)\mathbf{1}_{n,n}\alpha\Delta e_j = p_f^T(\Delta/2)\mathbf{1}_{n,n}q_f$, so:

$$p_f^T((\Delta/2)C + (\Delta/2)\mathbf{1}_{n,n})\alpha\Delta e_j > p_f^T((\Delta/2)C + (\Delta/2)\mathbf{1}_{n,n})q_f + \Delta^4/8\,. \tag{5.2}$$

We also have $p_{n+1}(2\Delta,\ldots,2\Delta)\alpha\Delta e_j = p_{n+1}(2\Delta,\ldots,2\Delta)q_f$. By adding this quantity as well as the term $p_f^T(1,\ldots,1)^T q_{n+1}$ to both sides of inequality (5.2) we get:

$$p^T C'(\alpha\Delta e_j + q_{n+1}e_{n+1}) > p^T C' q + \Delta^4/8\,.$$

Therefore there exists a deviation for the column player (namely moving all $\alpha\Delta$ probability mass from columns $1,2,\ldots,n$ onto column i), yielding a benefit of $\Delta^4/8$ to the column player. This contradicts the assumption that (p,q) is an $\Delta^4/8$-equilibrium in the new game, as desired. □

Note. Theorem 5.2 implies that for any $\varepsilon \leq (1/8)(0.3393)^4$, an algorithm for finding an ε-equilibrium in a game satisfying the strong $(\varepsilon, 8\varepsilon^{1/4})$-approximation-stability condition would imply a better than best currently known algorithm for finding approximate equilibria in general games (the best bound is currently 0.3393).

We end this section by proving Lemma 5.1.

Proof of Lemma 5.1. First note that $e_{n+1}^T Rq = 2\Delta q_{n+1}$ and

$$e_i^T Rq = (1+\alpha_{i,1})q_1 + (1+\alpha_{i,2})q_2 + \cdots + (1+\alpha_{i,n})q_n \quad \text{for} \quad 1 \leq i \leq n\,.$$

Also $p^T Ce_{n+1} = p_1 + \cdots + p_n$ and

$$p^T Ce_j = p_1\gamma_{1,j} + \cdots + p_n\gamma_{n,j} + 2\Delta p_{n+1} \quad \text{for} \quad 1 \leq j \leq n\,.$$

By a simple case analysis, one can show that any Nash equilibrium (p,q) must have $0 < p_{n+1} < 1$ and $0 < q_{n+1} < 1$. (This is also implicit in our analysis on Δ^2-Nash equilibria below.) This then implies that in any Nash equilibrium (p,q) such that $p_i \neq 0$ we must have:

$$2\Delta q_{n+1} = (1+\alpha_{i,1})q_1 + (1+\alpha_{i,2})q_2 + \cdots + (1+\alpha_{i,n})q_n\,. \tag{5.3}$$

Similarly, in any Nash equilibrium (p,q) such that $q_j \neq 0$ we must have:

$$p_1 + \cdots + p_n = p_1\gamma_{1,j} + \cdots + p_n\gamma_{n,j} + 2\Delta p_{n+1}\,. \tag{5.4}$$

Identities (5.3) and (5.4) together with the fact that $\alpha_{i,j} \in [-\Delta, 0]$ and $\gamma_{i,j} \in [0,\Delta]$ for all i,j, imply that there must exist a Nash equilibrium (p,q) satisfying:

$$\frac{2\Delta}{1+2\Delta} \leq p_1 + \cdots + p_n \leq \frac{2\Delta}{1+\Delta} \quad \text{and} \quad \frac{2\Delta}{1+2\Delta} \leq q_1 + \cdots + q_n \leq \frac{2\Delta}{1+\Delta}\,. \tag{5.5}$$

To get the desired stability guarantee we now show that any Δ^2-equilibrium must have $\Delta/2 \leq p_1 + \cdots + p_n \leq 4\Delta$ and $\Delta/2 \leq q_1 + \cdots + q_n \leq 4\Delta$. (This in turn implies that any Δ^2-equilibrium must be at distance at most 4Δ from a Nash equilibrium satisfying relation (5.5).) We prove this by contradiction. Consider an arbitrary Δ^2-equilibrium (p,q). We analyze a few cases.

Case 1. Suppose $p_{n+1} > 1-\Delta/2$. Then the column player's payoff for column $n+1$ is $p^T Ce_{n+1} = \sum_{i=1}^n p_i \leq \Delta/2$. But the column player's payoff for a column $j \in \{1,\ldots,n\}$ is:

$$p^T Ce_j = \sum_{i=1}^n \gamma_{i,j} p_i + 2\Delta p_{n+1} \geq 2\Delta(1-\Delta/2)\,.$$

If $q_{n+1} > 1/2$ then the column player has incentive to deviate at least:

$$p^T Ce_j - p^T Cq \geq (1/2)[2\Delta(1-\Delta/2) - \Delta/2] > \Delta/2 > \Delta^2\,,$$

which cannot happen since (p,q) is a Δ^2-equilibrium. On the other hand if $q_{n+1} \leq 1/2$, then the row player has huge incentive to deviate. Specifically, the row's player payoff for row 1 is $e_1^T Rq \geq (1/2)(1-\Delta)$, but row's player payoff for row $n+1$ is $e_{n+1}^T Rq \leq (1/2)2\Delta = \Delta$. Thus in this case the row player has incentive to deviate at least:

$$e_1^T Rq - p^T Rq \geq (1-\Delta/2)[1/2(1-\Delta) - \Delta] > \Delta\,,$$

which cannot happen since (p,q) is a Δ^2-equilibrium.

Case 2. Suppose $p_{n+1} < 1-4\Delta$. Then the column player's payoff for column $n+1$ is $p^T Ce_{n+1} = \sum_{i=1}^n p_i \geq 4\Delta$, whereas the column player's payoff for a column $j \in \{1,\ldots,n\}$ is

$$p^T Ce_j = p_1\gamma_{1,j} + \cdots + p_n\gamma_{n,j} + 2\Delta p_{n+1} \leq \Delta\sum_{i=1}^n p_i + 2\Delta p_{n+1} = \Delta(1-p_{n+1}) + 2\Delta p_{n+1} \leq 2\Delta\,.$$

So, if $q_{n+1} < 1-\Delta/2$, then the column player has incentive to deviate at least

$$p^T Ce_{n+1} - p^T Cq > (\Delta/2)[4\Delta - 2\Delta] \geq \Delta^2\,,$$

a contradiction. On the other hand, if $q_{n+1} > 1-\Delta/2$, then the row player's payoff for row $n+1$ is $e_{n+1}^T Rq \geq 2\Delta(1-\Delta/2)$, but the row player's payoff for a row $i \in \{1,\ldots,n\}$ is:

$$e_i^T Rq = \sum_{j=1}^n (1+\alpha_{i,j})q_j \leq 1 - q_{n+1} \leq \Delta/2\,.$$

So, in this case, the row player has incentive to deviate at least

$$e_{n+1}^T Rq - p^T Rq \geq 4\Delta[2\Delta(1-\Delta/2) - \Delta/2] > \Delta^2\,,$$

which cannot happen since (p,q) is a Δ^2-equilibrium.

Case 3. Suppose $q_{n+1} > 1-\Delta/2$. As in the bottom-half of the case 2 analysis we have that the row player's payoff for row $n+1$ is $e_{n+1}^T Rq \geq 2\Delta(1-\Delta/2)$, but the row player's payoffs for rows $1,\ldots,n$ are $\leq \Delta/2$. So, if $p_{n+1} < 1-\Delta$ then the row player has incentive to deviate at least

$$e_{n+1}^T Rq - p^T Rq \geq \Delta[2\Delta(1-\Delta/2) - \Delta/2] > \Delta^2\,,$$

which cannot happen since (p,q) is a Δ^2-equilibrium. On the other hand, if $p_{n+1} \geq 1-\Delta$, then the column player's payoff for column $n+1$ is $p^T Ce_{n+1} = \sum_{i=1}^n p_i \leq \Delta$, but the column player's payoff for a columns $j \in \{1,\ldots,n\}$ is:

$$p^T Ce_j = p_1\gamma_{1,j} + \cdots + p_n\gamma_{n,j} + 2\Delta p_{n+1} \geq (1-\Delta)(2\Delta)\,.$$

So, the column player has incentive to deviate at least

$$p^T Ce_j - p^T Cq \geq q_{n+1}[2\Delta(1-\Delta) - \Delta] \geq \Delta/2 > \Delta^2\,,$$

which cannot happen since (p,q) is a Δ^2-equilibrium.

Case 4. Finally assume that $q_{n+1} < 1-4\Delta$. Then the row player's payoff for row n+1 is $e_{n+1}^T Rq \leq 2\Delta(1-4\Delta)$, but row player's payoffs for rows $1,\ldots,n$ are $\geq (4\Delta)(1-\Delta)$. So, if $p_{n+1} > 1/2$ then Row has incentive to deviate at least

$$e_{n+1}^T Rq - p^T Rq \geq (1/2)[4\Delta(1-\Delta) - 2\Delta(1-4\Delta)] \geq \Delta/2 > \Delta^2\,,$$

which cannot happen since (p,q) is a Δ^2-equilibrium. Finally if $p_{n+1} < 1/2 < 1-4\Delta$ we apply the analysis in case 2.

Thus any Δ^2-equilibrium must have $\Delta/2 \leq p_1 + \cdots + p_n \leq 4\Delta$ and $\Delta/2 \leq q_1 + \cdots + q_n \leq 4\Delta$, as desired. This concludes the proof. □

6 Stability in constant-sum games

It is interesting to note that for our algorithmic results for finding approximate equilibria, we do *not* require knowing the stability parameters. If the game happens to be reasonably stable, then we get improved running times over the Lipton et al. [28] guarantees; if this is not the case, then we fall back to the Lipton et al. [28] guarantees. This is because algorithmically, we can simply try different support sizes in increasing order and stop when we find strategies forming a (well-supported) ε-equilibrium. In other words, given the desired approximation level ε, we can find an ε-approximate equilibrium in time $n^{O((\Delta^2/\varepsilon^2)\log(1+\Delta^{-1})\log n)}$ where Δ is the smallest value greater than or equal to ε such that the game is (ε,Δ)-perturbation stable. However, given a game, it might be interesting to know how stable it is. In this direction, in this section we provide a characterization of stable constant-sum games and an algorithm for computing the strong stability parameters of a given constant-sum game.

Consider a constant-sum game defined by R and C. Let

$$\mathcal{P}^* = \{p, \exists\, q \text{ such that } (p,q) \text{ is a Nash equilibrium}\}$$

and

$$\mathcal{Q}^* = \{q, \exists\, p \text{ such that } (p,q) \text{ is a Nash equilibrium}\}.$$

We say that p is Δ-far from $\mathcal{P}^*$ if the minimum distance between p and $p' \in \mathcal{P}^*$ is $> \Delta$. Let v_R and v_C be the unique values of the row and column player, respectively, in a Nash equilibrium [35]. Lemma 6.1 and Lemma 6.2 below characterize constant sum games satisfying approximation stability in terms of properties of the space of mixed strategies for the row player and column player separately. Theorem 6.3 gives a polynomial time algorithm for determining the approximately best parameters for the strong approximation stability property for a given game.

Lemma 6.1. *We have:*

(a) If $p^T Re_j < v_R - \alpha$ for some j then (p,q) is not an $\alpha/2$-Nash equilibrium for any q. Similarly, if $e_j^T Cq < v_C - \alpha$ for some j then (p,q) is not an $\alpha/2$-Nash equilibrium for any p.

(b) If for any p that is Δ-far from $\mathcal{P}^*$ there exists e_j such that $p^T Re_j < v_R - \alpha$ and for any q that is Δ-far from $\mathcal{Q}^*$ there exists e_j such that $e_j^T Cq < v_C - \alpha$, then the game satisfies $(\alpha/2, \Delta)$-approximation-stability.

Proof. Part (a): Consider p such that $p^T Re_j < v_R - \alpha$ for some j, and let q be arbitrary. If $p^T Rq < v_R - \alpha/2$, then this is not an $\alpha/2$-equilibrium since the row player could play its minimax optimal strategy and get v_R. On the other hand if $p^T Rq \geq v_R - \alpha/2$, then $p^T Cq \leq v_C + \alpha/2$, so this is also not an $\alpha/2$-equilibrium because $p^T Ce_j > v_C + \alpha$, so the column player has more than an $\alpha/2$ incentive to deviate. The argument for q such that $e_j^T Cq < v_C - \alpha$ for some j is completely symmetric.

Part (b): We show that under this condition, any (p,q) that is Δ-far from all Nash equilibria cannot be an $\alpha/2$-equilibrium. Consider (p,q) that is Δ-far from all Nash equilibria. Then either p is Δ-far from $\mathcal{P}^*$ or q is Δ-far from $\mathcal{Q}^*$.[12] This implies that by assumption either there exists e_j such that $p^T Re_j < v_R - \alpha$, or there exists e_j such that $e_j^T Cq < v_C - \alpha$. Therefore, by part (a), (p,q) cannot be an $\alpha/2$-Nash equilibrium. □

Lemma 6.2. *Let (p^*, q^*) be a Nash equilibrium.*

(a) If p satisfies $\min_j p^T Re_j \geq v_R - \alpha$ then (p, q^*) is an α Nash equilibrium. Similarly, if q satisfies $\min_j e_j^T Cq \geq v_C - \alpha$ then (p^*, q) is an α Nash equilibrium.

[12] This follows from the well known interchangeability property of constant-sum games, meaning that given two Nash equilibria points (p_1,q_1) and (p_2,q_2), the strategy pairs (p_1,q_2) and (p_2,q_1) are also Nash equilibria. To see why this follows, note that if both p and q are close to $\mathcal{P}^*$ and $\mathcal{Q}^*$, respectively, then there exist Nash equilibria (p_1,q_1), (p_2,q_2) such that $d(p,p_1) \leq \Delta$ and $d(q,q_2) \leq \Delta$. By interchangeability, we get that (p_1,q_2) is a Nash equilibrium and we also have $d((p_1,q_2),(p,q)) \leq \Delta$.

(b) *If there exists p that is Δ-far from $\mathcal{P}^*$ such that $\min_j p^T R e_j \geq v_R - \alpha$ or if there exists q that is Δ-far from $\mathcal{Q}^*$ such that $\min_j e_j^T C q \geq v_C - \alpha$, then the game cannot be (α,Δ)-approximation-stable.*

Proof. Part (a): Consider p such that $\min_j p^T R e_j \geq v_R - \alpha$. This implies $p^T R q^* \geq v_R - \alpha$. Moreover, $p'^T R q^* \leq v_R$ for any p' since q^* is minimax optimal for the column player. Therefore, the row player has at most α-incentive to deviate from (p,q^*). Similarly, $p^T C q^* \geq v_C$ (since q^* is minimax optimal) and since $\min_j p^T R e_j \geq v_R - \alpha$, we have $\max_j p^T C e_j \leq v_C + \alpha$. So the column player has at most α-incentive to deviate from (p,q^*) as well. The argument for q such that $\min_j e_j^T C q \geq v_C - \alpha$ is completely symmetric.

Part (b): Assume that there exists p that is Δ-far from $\mathcal{P}^*$ such that $\min_j p^T R e_j \geq v_R - \alpha$. By part (a), we know (p,q^*) is an α-Nash equilibrium. Since this is Δ-far from the set of Nash equilibria, the game is not (α,Δ)-approximation-stable. The case for q that is Δ-far from $\mathcal{Q}^*$ such that $\min_j e_j^T C q \geq v_C - \alpha$ is completely symmetric. □

If the game satisfies the strong approximation-stability condition, then we can efficiently compute good approximations for the stability parameters. Specifically, we have the following result.

Theorem 6.3. *Given any $0 < \alpha < 1$, we can use Algorithm 1 so that we determine whp (at least $1 - 1/\mathrm{poly}(n)$) a value Δ such that the game satisfies the strong $(\alpha/2, 2\Delta)$-approximation-stability condition but not strong $(\alpha, \Delta/2)$-approximation-stability. The running time is polynomial $n^{O(1/\alpha^2)}$.*

Proof. We first find a minimax optimal solution (p^*,q^*) and then in Step 2, a small support α-Nash equilibrium (p',q'). Note that we can do Step 2 efficiently since we have (p^*,q^*) and it succeeds with probability $1 - 1/\mathrm{poly}(n)$.

In Step 3 we find the maximum Δ_p such that for some p of distance Δ_p from p' we have $p^T R e_j \geq v_R - \alpha$ for all j, and in Step 4 we find the maximum Δ_q such that for some q of distance Δ_q from q' we have $e_j^T C q \geq v_C - \alpha$ for all j. So, for all (p,q) of distance greater than $\Delta = \max(\Delta_p, \Delta_q)$ from (p',q'), either $p^T R e_j < v_R - \alpha$ for some j or $e_j^T C q < v_C - \alpha$ for some j. By Lemma 6.1(a), this implies that all $\alpha/2$-Nash equilibria are within distance Δ of (p',q'). By Lemma 6.2(a), this also implies that there exists an α-Nash equilibrium that is at distance at least Δ from (p',q'); in particular, either (p,q^*) or (p^*,q) depending on whether the maximum Δ occurred in Step 3 or Step 4, respectively.

In sum, in Steps 3 and 4 we find Δ such that all $\alpha/2$-Nash equilibria are within distance Δ of (p',q') and there exists an α-Nash equilibrium at distance Δ from (p',q'). By triangle inequality, we obtain that the game is $(\alpha/2, 2\Delta)$ stable and it is *not* $(\alpha, \Delta/2)$ stable. In particular, since all $\alpha/2$-Nash equilibria are within distance Δ of (p',q'), and this includes (p^*,q^*), we get that all $\alpha/2$-Nash equilibria are within distance 2Δ of (p^*,q^*), so the game is $(\alpha/2, 2\Delta)$ stable. We also know that there exists $(\tilde{p},\tilde{q})$ an α-Nash equilibrium at distance Δ from (p',q') (which is also an α-Nash equilibrium), so by triangle inequality, at least one of the two α-Nash equilibria $(\tilde{p},\tilde{q})$ and (p',q') is at distance at least $\Delta/2$ from (p^*,q^*), so the game is not $(\alpha, \Delta/2)$ stable.

Note that the running time is polynomial since we we perform steps 3 and 4 at most $2^{O(\log n/\alpha^2)} = n^{O(1/\alpha^2)}$ times, so overall the running time is $n^{O(1/\alpha^2)}$. □

Algorithm 1 Determining the strong stability parameters of a constant sum game.

Input: R, C, parameter α.

1. Solve for minimax optimal (p^*, q^*) and compute minimax values v_R and v_C.
2. Apply the sampling procedure in [28] from (p^*, q^*) to get (p', q') with support of size $O((\log n)/\alpha^2)$ that is an α-Nash.
3. For each partition Π_p of the support of p' into supp_+ and supp_-, solve the following LP:
$$\begin{aligned} \text{maximize } v(\Pi_p) = \sum_{i \in \text{supp}_+} (p_i - p'_i) + \sum_{i \in \text{supp}_-} (p'_i - p_i) + \sum_{i \in \{1,2,\dots,n\} \setminus \text{supp}(p')} p_i \\ \text{s.t.} \qquad p_i \geq p'_i \text{ for all } i \in \text{supp}_+, \\ p_i \leq p'_i \text{ for all } i \in \text{supp}_-, \\ p_i \geq 0 \text{ for all } i \text{ and } \textstyle\sum_i p_i = 1, \\ p^T R e_j \geq v_R - \alpha \text{ for all } j. \end{aligned}$$
Let $\Delta_p = \max_{\Pi_p} v(\Pi_p)$.
4. For each partition Π_q of the support of q' into supp_+ and supp_-, solve the following LP:
$$\begin{aligned} \text{maximize } v(\Pi_q) = \sum_{i \in \text{supp}_+} (q_i - q'_i) + \sum_{i \in \text{supp}_-} (q'_i - q_i) + \sum_{i \in \{1,2,\dots,n\} \setminus \text{supp}(q')} q_i \\ \text{s.t.} \qquad q_i \geq q'_i \text{ for all } i \in \text{supp}_+, \\ q_i \leq q'_i \text{ for all } i \in \text{supp}_-, \\ q_i \geq 0 \text{ for all } i \text{ and } \textstyle\sum_i q_i = 1, \\ e_j^T C q \geq v_C - \alpha \text{ for all } j. \end{aligned}$$
Let $\Delta_q = \max_{\Pi_q} v(\Pi_q)$.

Output: Radius $\Delta = \max(\Delta_p, \Delta_q)$.

7 Discussion and open questions

Our main results show that for $0 \leq \varepsilon \leq \Delta \leq 1$, all n-action (ε, Δ)-perturbation-stable games with at most $n^{O((\Delta/\varepsilon)^2)}$ Nash equilibria must have at least one well-supported ε-equilibrium of support
$$O\left(\frac{\Delta^2 \log(1+\Delta^{-1})}{\varepsilon^2} \log n\right).$$
Our current proof crucially uses the limited number of Nash-equilibria. In particular, the key technical component (Lemma 4.2) shows that if p^* (or q^*) had a portion with substantial L_1 norm and low L_2 norm, then there must exist a deviation from p^* (or q^*) that is far from *all* equilibria and yet is a well-supported

approximate Nash equilibrium, violating the stability condition. Proving the existence of such a deviation is based on probabilistic argument and requires doing a union bound over all such equilibria. It would be very interesting if the upper bound on the number of equilibria could be relaxed by using a more refined way to guard against all the equilibria or by using a different technique.

Another interesting question is whether the games that are the result of reductions proving PPAD-hardness are just $(\varepsilon,1)$-perturbation stable or whether Δ can be <1. Our results imply that the games used in the reduction of PPAD-hardness cannot be particularly stable, in particular they cannot satisfy the *t-uniform perturbation-stability* condition since for games satisfying t-uniform perturbation stability with $t=\mathrm{poly}(\log(n))$, our results imply that we can find $1/\mathrm{poly}(n)$-approximate equilibria in $n^{\mathrm{poly}(\log(n))}$ time, i. e., achieve a fully quasi-polynomial-time approximation scheme (FQPTAS). This is especially interesting because the results of [16] prove that it is PPAD-hard to find $1/\mathrm{poly}(n)$-approximate equilibria in general games. Our results show that under the (widely believed) assumption that PPAD is not contained in quasi-polynomial time [17], such uniformly stable games are inherently easier for computation of approximate equilibria than general bimatrix games. It is an interesting open question to understand the level of stability of the games used in PPAD hardness.

Additionally, we show that determining whether a given game satisfies strong approximation stability (and computing approximations to the stability parameters) can be done in polynomial time for constant-sum games. It is an open question whether this can be done efficiently for broader classes of games. For example, it would be interesting to consider the class of bimatrix games satisfying the condition of *mutual concavity* [27], or equivalently the class of *strategically zero-sum* games [30].

Acknowledgments

We thank the reviewers for their helpful comments. We thank Avrim Blum, Dick Lipton, Yishay Mansour, Shanghua Teng, and Santosh Vempala for useful discussions. We also thank Vangelis Markakis for pointing [29] to us.

A Standard facts

The following is the McDiarmid inequality (see [21]) we use in our proofs.

Theorem A.1. *Let $Y_1,\dots,Y_n$ be independent random variables taking values in some set A, and assume that $t: A^n \to R$ satisfies:*

$$\sup_{y_1,\dots,y_n\in A,\bar{y}_i\in A} |t(y_1,\dots,y_n)-t(y_1,\dots,y_{i-1},\bar{y}_i,y_{i+1},y_n)| \le c_i\,,$$

for all i, $1\le i\le n$. Then for all $\gamma>0$ we have:

$$\Pr\{|t(Y_1,\dots,Y_n)-\mathbf{E}[t(Y_1,\dots,Y_n)]|\ge\gamma\}\le 2e^{-2\gamma^2/\sum_{i=1}^{n}c_i^2}\,.$$

Also, it is a well known fact that any pair of strategies that is sufficiently close to a Nash equilibrium is a sufficiently good approximate Nash equilibrium.

Claim A.2. *If (p,q) is α-close to a Nash equilibrium (p^*,q^*) (i. e., if $d((p,q),(p^*,q^*))\le\alpha$), then (p,q) is a 3α-Nash equilibrium.*

B Examples

We present here several examples of games satisfying our stability condition.

Example 1. A classic game from experimental economics is the public goods game which is defined is as follows. We have two players and each can choose to play a number between 0 and $n-1$ corresponding to an amount of money to contribute. If the Row player contributes i dollars and the Column player contributes j dollars, then each gets back $0.75(i+j)$. So the payoff to the Row player is $0.75(i+j)-i$ and the payoff to the Column player is $0.75(i+j)-j$, where $i \in \{0,1,\ldots,n-1\}$ and $j \in \{0,1,\ldots,n-1\}$. This has payoffs ranging from 0 up to $0.75(n-1)$, so to scale to the range $[0,1]$ (as we do in this paper), we multiply all the payoffs by $1/n$. That is, if the Row player plays i and the Column player plays j then the payoff to the Row player is $R[i,j]=[0.75j-0.25i]/n$ and the payoff to the Column player is $C[i,j]=[0.75i-0.25j]/n$.

We first claim that this game is $(\varepsilon,0)$-stable under perturbations for all $\varepsilon<1/(8n)$. To see this, note that without any perturbation, for any j and any $i \geq 1$ we have $e_0^T Re_j - e_i^T Re_j = 0.25i/n \geq 0.25/n$. That means that the Row player prefers playing action 0 compared to action i by $0.25i/n \geq 0.25/n$. So, in a game R',C' that is an L_∞ ε-perturbation of the original game we get: $e_0^T R'e_j - e_i^T R'e_j \geq 0.25/n - 2\varepsilon > 0$. That means that in the perturbed game, the Row player still prefers playing action 0. This implies that the only equilibrium in the perturbed game has the Row player playing action 0, and similarly the Column player playing 0, so the only equilibrium is $(0,0)$.

We now claim that this game is not $(\varepsilon,0.99)$ stable for any $\varepsilon>1/(4n)$. To see this consider adding ε to $R[1,0]$. That is, if the Row player plays action 1 and the Column player plays action 0, then the payoff to Row player is $-0.25/n+\varepsilon>0$. Now, $(1,0)$ is a Nash equilibrium since this payoff is strictly greater than $R[i,0]$ for any $i \neq 1$. In particular, $R[0,0]=0$ and $R[i,0]<0$ for all $i \geq 2$. So, there is now a Nash equilibrium (actually the unique Nash equilibrium) at variation distance 1 from the original Nash equilibrium.

Example 2. We present here a variant of the identical interest game. Both players have n available actions. The first action is to stay home, and the other actions correspond to $n-1$ different possible meeting locations. If a player chooses action 1 (stay home), his payoff is $1/2$ no matter what the other player is doing. If the player chooses to go out to a meeting location, his payoff is 1 if the other player is there as well and it is 0 otherwise. Formally, if the Row player plays i and the Column player plays j then the payoff to the Row player is

$$R[1,j]=1/2 \text{ for all } j,\ R[i,i]=1 \text{ for all } i>1,\ R[i,j]=0 \text{ for } i>1,\ j\neq i.$$

and the payoff to the Column player is

$$C[i,1]=1/2,\ C[j,j]=1 \text{ for } j>1,\ C[i,j]=0 \text{ for } j>1, i\neq j.$$

We claim that this game is well supported $(\varepsilon,2\varepsilon)$-stable for all $\varepsilon<1/6$, so it is 2-uniformly stable.

Note that $e_1^T Rq = 1/2$ and $e_i^T Rq = q_i$ for $i>1$. Similarly, $p^T Ce_1 = 1/2$ and $p^T Ce_i = p_i$ for $i>1$. Note that if (p,q) is an well supported ε-Nash equilibrium and if $q_i < 1/2-\varepsilon$ for $i>1$ then both $p_i=0$ and $q_i=0$. This follows immediately since $e_1^T Rq = 1/2$ and $e_i^T Rq = q_i$ for $i>1$, so p_i must equal 0 on any action whose expected payoff is $< 1/2-\varepsilon$. Since $p_i = 0$, q_i must equal 0 as well in order to be well-supported. Also note that if (p,q) is an well supported ε-Nash equilibrium and if $q_i > 1/2+\varepsilon$ for $i > 0$ then both $p_i = 1$ and $q_i = 1$. If $q_i > 1/2+\varepsilon$ we have $e_i^T Rq = 1/2+\varepsilon$ and since $e_j^T Rq \leq 1/2$ for $j \neq i$ we must have $p_i = 1$. This in turn implies $q_i = 1$. Similarly, we can show the same for the row player as well.

These imply that the well supported ε-Nash equilibria that are not already Nash equilibria must satisfy: for any action $i > 1$, $i \in \operatorname{supp}(q)$, we have $1/2-\varepsilon \leq q_i \leq 1/2+\varepsilon$ and $1/2-\varepsilon \leq p_i \leq 1/2+\varepsilon$. Similarly, for any action $i > 1$, $i \in \operatorname{supp}(p)$, we have $1/2-\varepsilon \leq q_i \leq 1/2+\varepsilon$ and $1/2-\varepsilon \leq p_i \leq 1/2+\varepsilon$. We have two cases. The first one is if there is exactly one action $i > 1$ in $\operatorname{supp}(q)$. In that case, (p,q) has distance at most ε from the Nash equilibrium $(1/2e_0+1/2e_i, 1/2e_0+1/2e_i)$. The second one is if there are two such actions $i, j > 0$ in $\operatorname{supp}(q)$. In that case, (p,q) has distance at most 2ε from the Nash equilibrium $(1/2e_i+1/2e_j, 1/2e_i+1/2e_j)$, as desired.

Example 3. To relate to other conditions considered in the literature, we point out here that perturbation stability is incomparable to the class of constant-rank games studied, e. g., in [25, 1]. One can generate n-by-n matrices R and C that are highly stable and yet where $R+C$ has large rank (rank n), and one can also generate zero-sum games that are very unstable under perturbations.

In the first direction, consider any pair of n-by-n matrices R and C such that:

- The first row of R is all 1s.
- The first columns of C is all 1s.
- All other entries are in the range $[0,1/2)$.

Such matrices are $(1/4,0)$-stable under perturbations, since under such perturbations, the only Nash equilibrium is still $p = q = e_1$. Moreover, if the entries in $[0,1/2)$ are filled in randomly (or even filled in arbitrarily but with infinitesimal perturbations added), we will have $\operatorname{rank}(R+C) = n$ with probability 1.

In the other direction, there are zero-sum games that are very unstable under perturbations. For example, let $\delta > 0$ be a small value and consider

$$R = -C = \begin{bmatrix} 1 & 1-\delta \\ 1-\delta & 1-2\delta \end{bmatrix}; \quad R' = -C' = \begin{bmatrix} 1-2\delta & 1-\delta \\ 1-\delta & 1 \end{bmatrix}.$$

Here, (R',C') is a 2δ-perturbation of (R,C), and yet the equilibria are far apart. In particular, the unique Nash equilibrium of (R,C) is (e_1,e_2), yet the unique Nash equilibrium of (R',C') is (e_2,e_1).

C Range of parameters

Lemma C.1. *Assume that the game $\mathcal{G}$ satisfies the (ε,Δ)-approximation-stability and that the union of all Δ-balls around all Nash equilibria do not cover the whole space. Then we must have $3\Delta \geq \varepsilon$.*

Proof. Since the union of all Δ-balls around all Nash equilibria do not cover the whole space, we must have a (p,q) that is at distance exactly Δ from some fixed Nash equilibrium and that is Δ-far from all the other Nash equilibria. By Claim A.2 we also have that this is a 3Δ-Nash equilibrium. This then implies the desired result. ☐

Lemma C.2. *Assume that the bimatrix game $\mathcal{G}$ specified by R and C has a non-pure Nash equilibrium.*

(a) If $\mathcal{G}$ satisfies the strong well-supported (ε,Δ)-approximation-stability condition, then we must have $\Delta \geq \varepsilon/4$.

(b) If $\mathcal{G}$ satisfies the strong (ε,Δ)-perturbation-stability condition, then we must have $\Delta \geq \varepsilon/8$.

Proof. Assume $\mathcal{G}$ satisfies the strong well-supported (ε,Δ)-approximation-stability condition. By definition, there exists a Nash equilibrium (p^*,q^*) such that any well supported ε-equilibrium is Δ-close to (p^*,q^*). Let (p,q) be an arbitrary non-pure Nash equilibrium of $\mathcal{G}$ and assume without loss of generality that p is a mixed strategy. Consider an α internal deviation of the row player, i. e., consider p' with $\operatorname{supp}(p') \subseteq \operatorname{supp}(p)$ such that $d(p,p') = \alpha$. Since $\operatorname{supp}(p') \subseteq \operatorname{supp}(p)$ we have $p'^T Rq = p^T Rq$. Since (p,q) is a Nash equilibrium we have $p^T Ce_j = p^T Cq \equiv v_C$ for all $j \in \operatorname{supp}(q)$ and $p^T Ce_j \leq v_C$ for all $j \notin \operatorname{supp}(q)$. Since $d(p,p') = \alpha$ we have

$$|p'^T Ce_j - p^T Ce_j| \leq |(p'-p)^T Ce_j| \leq \alpha\,,$$

for all j, so $p'^T Ce_j \geq v_C - \alpha$, for all $j \in \operatorname{supp}(q)$ and $p'^T Ce_j \leq v_C + \alpha$, for all $j \notin \operatorname{supp}(q)$. Thus (p',q) is a well supported 2α-Nash equilibrium. By construction, we have $d((p',q),(p,q)) = \alpha$. Since d is a metric, by the triangle inequality, we get that at least one of the pairs (p',q) and (p,q) is at least $\alpha/2$ far from (p^*,q^*); however they are both 2α well supported Nash equilibria. This implies that we must have $\Delta \geq \varepsilon/4$, as desired. By Theorem 3.3, we immediately get (b) as well. □

References

[1] BHARAT ADSUL, JUGAL GARG, RUTA MEHTA, AND MILIND SOHONI: Rank-1 bimatrix games: A homeomorphism and a polynomial time algorithm. In *Proc. 43rd STOC*, pp. 195–204. ACM Press, 2011. [doi:10.1145/1993636.1993664, arXiv:1010.3083] 5, 26

[2] MICHELE AGHASSI AND DIMITRIS BERTSIMAS: Robust game theory. *Math. Program.*, 107(1-2):231–273, 2006. [doi:10.1007/s10107-005-0686-0] 5

[3] HARIS ANGELIDAKIS, KONSTANTIN MAKARYCHEV, AND YURY MAKARYCHEV: Algorithms for stable and perturbation-resilient problems. In *Proc. 49th STOC*, pp. 438–451. ACM Press, 2017. [doi:10.1145/3055399.3055487] 5

[4] PRANJAL AWASTHI AND MARIA-FLORINA BALCAN: Center based clustering: A foundational perspective. In HENNIG, MEILA, MURTAGH, AND ROCCI, editors, *Handbook of Cluster Analysis*. CRC Press, 2015. Available from CMU. 5

[5] PRANJAL AWASTHI, MARIA-FLORINA BALCAN, AVRIM BLUM, OR SHEFFET, AND SANTOSH VEMPALA: On Nash-equilibria of approximation-stable games. In *Proc. 3rd Internat. Symp. Algorithmic Game Theory (SAGT'10)*, pp. 78–89, 2010. [doi:10.1007/978-3-642-16170-4_8] 4, 5, 8, 9, 10

[6] PRANJAL AWASTHI, MARIA-FLORINA BALCAN, AVRIM BLUM, OR SHEFFET, AND SANTOSH VEMPALA: On Nash-equilibria of approximation-stable games. *Current Science J.*, 103(9):1014–1020, 2012. Available at JSTOR. Preliminary version in SAGT'10. 4, 5

[7] PRANJAL AWASTHI, AVRIM BLUM, AND OR SHEFFET: Center-based clustering under perturbation stability. *Inform. Process. Lett.*, 112(1-2):49–54, 2012. [doi:10.1016/j.ipl.2011.10.006, arXiv:1009.3594] 5

[8] MARIA-FLORINA BALCAN, AVRIM BLUM, AND ANUPAM GUPTA: Approximate clustering without the approximation. In *Proc. 20th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'09)*, pp. 1068–1077. ACM Press, 2009. [doi:10.1137/1.9781611973068.116] 7

[9] MARIA-FLORINA BALCAN, AVRIM BLUM, AND ANUPAM GUPTA: Clustering under approximation stability. *J. ACM*, 60(2):8:1–8:34, 2013. [doi:10.1145/2450142.2450144] 7

[10] MARIA-FLORINA BALCAN, NIKA HAGHTALAB, AND COLIN WHITE: k-center clustering under perturbation resilience. In *Proc. 43rd Internat. Colloq. on Automata, Languages and Programming (ICALP'16)*, volume 55 of *LIPIcs*, pp. 68:1–68:14. Schloss Dagstuhl, 2016. [doi:10.4230/LIPIcs.ICALP.2016.68, arXiv:1505.03924] 5

[11] MARIA-FLORINA BALCAN AND YINGYU LIANG: Clustering under perturbation resilience. *SIAM J. Comput.*, 45(1):102–155, 2016. Preliminary version in ICALP'12. [doi:10.1137/140981575, arXiv:1112.0826] 5

[12] IMRE BÁRÁNY, SANTOSH VEMPALA, AND ADRIAN VETTA: Nash equilibria in random games. *Random Struct. Algorithms*, 31(4):391–405, 2007. Preliminary version in FOCS'05. [doi:10.1002/rsa.20199] 5

[13] YONATAN BILU AND NATHAN LINIAL: Are stable instances easy? *Combinatorics, Probab. Comput.*, 21(5):643–660, 2012. Preliminary version appeared in ICS 2010. [doi:10.1017/S0963548312000193, arXiv:0906.3162] 5

[14] HARTWIG BOSSE, JAROSLAW BYRKA, AND EVANGELOS MARKAKIS: New algorithms for approximate Nash equilibria in bimatrix games. *Theor. Comput. Sci.*, 411(1):164–173, 2010. Preliminary version in WINE'07. [doi:10.1016/j.tcs.2009.09.023] 2

[15] XI CHEN AND XIAOTIE DENG: Settling the complexity of two-player Nash equilibrium. In *Proc. 47th FOCS*, pp. 261–272. IEEE Comp. Soc. Press, 2006. [doi:10.1109/FOCS.2006.69, arXiv:0704.1678] 2

[16] XI CHEN, XIAOTIE DENG, AND SHANG-HUA TENG: Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3):14:1–14:57, 2009. [doi:10.1145/1516512.1516516, arXiv:0704.1678] 2, 3, 16, 24

[17] CONSTANTINOS DASKALAKIS: On the complexity of approximating a Nash equilibrium. *ACM Trans. Algorithms*, 9(3):23:1–23:35, 2013. Preliminary version in SODA'11. [doi:10.1145/2483699.2483703] 3, 16, 24

[18] CONSTANTINOS DASKALAKIS, PAUL W. GOLDBERG, AND CHRISTOS H. PAPADIMITRIOU: The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009. Preliminary version in STOC'06. [doi:10.1137/070699652] 2

[19] CONSTANTINOS DASKALAKIS, ARANYAK MEHTA, AND CHRISTOS H. PAPADIMITRIOU: Progress in approximate Nash equilibria. In *Proc. 8th ACM Conf. on Electronic Commerce (EC'07)*, pp. 355–358, 2007. [doi:10.1145/1250910.1250962] 2

[20] CONSTANTINOS DASKALAKIS, ARANYAK MEHTA, AND CHRISTOS H. PAPADIMITRIOU: A note on approximate Nash equilibria. *Theoret. Comput. Sci.*, 410(17):1581–1588, 2009. Preliminary version in WINE'06. [doi:10.1016/j.tcs.2008.12.031] 2

[21] LUC DEVROYE, LÁSZLÓ GYŐRFI, AND GÁBOR LUGOSI: *A Probabilistic Theory of Pattern Recognition*. Springer, 1996. [doi:10.1007/978-1-4612-0711-5] 24

[22] STEVEN N. DURLAUF AND LAWRENCE E. BLUME: *Game Theory*. Palgrave Macmillan, 2008. [doi:10.1057/9780230280847] 3, 16

[23] TOMÁS FEDER, HAMID NAZERZADEH, AND AMIN SABERI: Approximating Nash equilibria using small-support strategies. In *Proc. 8th ACM Conf. on Electronic Commerce (EC'07)*, pp. 352–354. ACM, 2007. [doi:10.1145/1250910.1250961] 2

[24] RISHI GUPTA, TIM ROUGHGARDEN, AND C. SESHADHRI: Decompositions of triangle-dense graphs. *SIAM J. Comput.*, 45(2):197–215, 2016. Preliminary version in ITCS'14. [doi:10.1137/140955331, arXiv:1309.7440] 7

[25] RAVI KANNAN AND THORSTEN THEOBALD: Games of fixed rank: A hierarchy of bimatrix games. *Economic Theory*, 42(1):157–173, 2010. Preliminary version in SODA'07. [doi:10.1007/s00199-009-0436-2, arXiv:cs/0511021] 5, 26

[26] SPYROS C. KONTOGIANNIS, PANAGIOTA N. PANAGOPOULOU, AND PAUL G. SPIRAKIS: Polynomial algorithms for approximating Nash equilibria of bimatrix games. *Theor. Comput. Sci.*, 410(17):1599–1606, 2009. Preliminary version in WINE'06. [doi:10.1016/j.tcs.2008.12.033] 2

[27] SPYROS C. KONTOGIANNIS AND PAUL G. SPIRAKIS: On mutual concavity and strategically-zero-sum bimatrix games. *Theor. Comput. Sci.*, 432:64–76, 2012. Preliminary version in APPROX'10. [doi:10.1016/j.tcs.2012.01.016] 24

[28] RICHARD J. LIPTON, EVANGELOS MARKAKIS, AND ARANYAK MEHTA: Playing large games using simple strategies. In *Proc. 4th ACM Conf. on Electronic Commerce (EC'03)*, pp. 36–41. ACM, 2003. [doi:10.1145/779928.779933] 2, 3, 4, 9, 10, 14, 15, 20, 23

[29] RICHARD J. LIPTON, EVANGELOS MARKAKIS, AND ARANYAK MEHTA: On stability properties of economic solution concepts. Manuscript, 2006. 5, 6, 24

[30] HERVÉ MOULIN AND JEAN-PHILIPPE VIAL: Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *Internat. J. Game Theory*, 7(3–4):201–221, 1978. [doi:10.1007/BF01769190] 24

[31] JOHN F. NASH: Non-cooperative games. *Ann. of Math.*, 54(2):286–295, 1951. [doi:10.2307/1969529] 2, 6

[32] AVIAD RUBINSTEIN: Settling the complexity of computing approximate two-player Nash equilibria. *ACM SIGecom Exchanges*, 15(2):45–49, 2017. Preliminary version in FOCS'16. [doi:10.1145/3055589.3055596, arXiv:1606.04550] 2

[33] HARALAMPOS TSAKNAKIS AND PAUL G. SPIRAKIS: An optimization approach for approximate Nash equilibria. *Internet Mathematics*, 5(4):365–382, 2007. Preliminary version in WINE'07. [doi:10.1080/15427951.2008.10129172] 2

[34] KONSTANTIN VOEVODSKI, MARIA-FLORINA BALCAN, HEIKO RÖGLIN, SHANG-HUA TENG, AND YU XIA: Efficient clustering with limited distance information. In *Proc. 26th Ann. Conf. on Uncertainty in Artificial Intelligence (UAI'10)*, pp. 632–640. AUAI Press, 2010. Available from AUAI Press. [arXiv:1408.2045] 7

[35] JOHN VON NEUMANN AND OSKAR MORGENSTERN: *Theory of Games and Economic Behavior (60th-Anniversary Edition)*. Princeton Univ. Press, 2007. [http://press.princeton.edu/titles/7802.html] 2, 21

AUTHORS

Maria-Florina Balcan
Associate professor
School of Computer Science
Carnegie Mellon University, Pittsburgh, PA
ninamf@cs.cmu.edu
http://www.cs.cmu.edu/~ninamf

Mark Braverman
Professor
Department of Computer Science
Princeton University, Princeton, NJ
mbraverm@cs.princeton.edu
http://www.cs.princeton.edu/~mbraverm

ABOUT THE AUTHORS

MARIA-FLORINA BALCAN is an Associate Professor in the School of Computer Science at Carnegie Mellon University. Her main research interests are machine learning, computational aspects in economics and game theory, and algorithms. Her honors include the CMU SCS Distinguished Dissertation Award, an NSF CAREER Award, a Microsoft Faculty Research Fellowship, a Sloan Research Fellowship, and several paper awards. She has held several leadership positions including being a Program Committee Co-chair for COLT 2014, a Program Committee Co-chair for ICML 2016, and a board member of the International Machine Learning Society.

MARK BRAVERMAN received his Ph. D. from the University of Toronto in 2008 under the supervision of Stephen Cook. He is currently a Professor of Computer Science at Princeton University. He is interested in most aspects of theoretical computer science—particularly in its connections to other disciplines such as information theory, game theory, and mechanism design. His work has been recognized by an NSF CAREER Award, a Packard Fellowship in Science and Engineering, and a Presburger Award in Theoretical Computer Science.