

Nash Equilibria in Perturbation Resilient Games

Maria-Florina Balcan*

Mark Braverman †

March 14, 2012

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 to 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. [18], by an amount that depends on the degree of stability of the game and that reduces to their bound in the worst case. Furthermore, 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 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 [20, 22]. 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 [10] and Chen, Deng, and Teng [7, 8] 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 $n^{O((\log n)/\epsilon^2)}$, based on a structural result of Lipton et al. [18] showing that there always exist ϵ -approximate equilibria with support over at most $O((\log n)/\epsilon^2)$ strategies. This structural result has been shown to be existentially tight [15]. Even for large values of ϵ , despite considerable effort [11, 12, 21, 15, 6, 17], polynomial-time algorithms for computing ϵ -approximate equilibria are known

*School of Computer Science, College of Computing, Georgia Institute of Technology, Atlanta, Georgia.

†Princeton University.

only for $\epsilon \geq 0.3393$ [21]. 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 to 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 what is 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 to 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. 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 the (ϵ, Δ) *perturbation stability* condition, 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 $(\epsilon, 1)$ 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 all n -action (ϵ, Δ) perturbation-stable games with at most $n^{O((\Delta/\epsilon)^2)}$ Nash equilibria must have a well-supported ϵ -equilibrium of support $O(\frac{\Delta^2 \log(1+\Delta^{-1})}{\epsilon^2} \log n)$. This yields an $n^{O(\frac{\Delta^2 \log(1+\Delta^{-1})}{\epsilon^2} \log n)}$ -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 [18] for games satisfying this condition (and reducing to the bound of [18] 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 $\epsilon_0 = 1/\text{poly}(n)$ and some t satisfy the $(\epsilon, t\epsilon)$ stability to perturbations condition for all $\epsilon < \epsilon_0$. For games satisfying *t-uniform stability to perturbations* with $t = \text{poly}(\log(n))$, our results imply that we can find $1/\text{poly}(n)$ -approximate equilibria in $n^{\text{poly}(\log(n))}$ time, i.e., achieve a fully quasi-polynomial-time approximation scheme (FQP-TAS). This is especially interesting because the results of [8] prove that it is PPAD-hard to find $1/\text{poly}(n)$ -approximate equilibria in general games. Our results shows that under the (widely believed) assumption that PPAD is not contained in quasi-polynomial time [9], such uniformly stable game are inherently easier for

¹For 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.

²One should think of Δ as a function of ϵ , with both possibly depending on n . E.g., $\epsilon = 1/\sqrt{n}$ and $\Delta = 6\epsilon$.

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 [14] satisfy this condition. See Sections 3, 5, and Appendix C for detailed examples.

Our second main result shows that computing an ϵ -equilibrium in a game satisfying the $(\epsilon, \Theta(\epsilon^{1/4}))$ perturbation stability condition is as hard as computing a $\Theta(\epsilon^{1/4})$ -equilibrium in a general game. For our reduction, we show that any general game can be embedded into one having the $(\epsilon, \Theta(\epsilon^{1/4}))$ perturbation stability property such that an ϵ equilibrium in the new game yields a $\Theta(\epsilon^{1/4})$ -equilibrium in the original game. This result implies that the interesting range for the (ϵ, Δ) -perturbation stability condition, where one could hope to do significantly better than in the general case, is $\Delta = o(\epsilon^{1/4})$.

We also connect our stability to perturbations condition to a seemingly very different stability to approximations condition introduced by Awasthi et al. [2]. 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 stability to perturbations condition is equivalent to a much weaker version of this condition that we call the well-supported approximation stability. This condition requires only that for any well-supported ϵ -approximate equilibrium⁵ (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 considered by [2] 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 the interesting range of parameters is $\Delta = O(\epsilon^{1/4})$) also holds for the strong stability to approximations condition.

We also provide an interesting structural result showing that each Nash equilibrium of an (ϵ, Δ) perturbation stable game with $n^{O((\Delta/\epsilon)^2)}$ Nash equilibria must be 8Δ -close to well-supported ϵ -approximate equilibrium of support only $O(\frac{\Delta^2 \log(1+\Delta^{-1})}{\epsilon^2} \log n)$. 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. [18] guarantees; if this is not the case, then we fall back to the Lipton et al. [18] guarantees.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 in Section 6 and an algorithm for computing the

³The generic result of [18] achieves quasi-polynomial time only for $\epsilon = \Omega(1/\text{poly}(\log n))$.

⁴[2] 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.

⁵Recall 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.

⁶This 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 ϵ , the desired approximation level, we can find an ϵ -approximate equilibrium in time $n^{O(\frac{\Delta^2 \log(1+\Delta^{-1})}{\epsilon^2} \log n)}$ where Δ is the smallest value such that the game is (ϵ, Δ) perturbation stable.

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) equilibria in specific classes of bimatrix games. For example, Bárány et al. [4] 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 [16] 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.

Awasthi et al. [2] 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 bounds on approximate equilibria can be viewed as generalizing the corresponding result of [2] and it is significantly more challenging technically. Moreover, our lower bounds also apply to the stability notion of [2] 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 [5] 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, for $\alpha = \sqrt{n}$, one can find the optimum solution in polynomial time. Recently, Awasthi et al. [3], have shown a similar result for the k -median clustering problem and showed a similar result for $\alpha = \sqrt{3}$. Our stability to perturbations notion 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.

The notion of stability to perturbations we consider in our paper is also related to the stability notions considered by Lipton et al. [19] 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 to the row player and C denote the payoff matrix of the column player. If the row player chooses strategy i and the

⁷Just as in [19], one can show that for the stability conditions we consider in our paper, there exist unstable instances.

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 i th coordinate and 0 elsewhere. For a mixed strategy pair (p, q) , 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 $p^T R q$. Similarly the payoff to the column player is $p^T C q$. Given strategies p and q for the row and column player, we denote by $\text{supp}(p)$ and $\text{supp}(q)$ the support of p and q , respectively.

A Nash equilibrium [20] is a pair of strategies (p^*, q^*) 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 of mixed strategies (p^*, q^*) is a Nash-equilibrium if for every pure strategy i of the row player we have $e_i^T R q^* \leq p^{*T} R q^*$, and for every pure strategy j of the column player we have $p^{*T} C e_j \leq p^{*T} C q^*$. Note that in a Nash equilibrium (p^*, q^*) , all rows i in the support of p^* satisfy $e_i^T R q^* = p^{*T} R q^*$ and similarly all columns j in the support of q^* satisfy $p^{*T} C e_j = p^{*T} C q^*$.

Definition 1 A pair of mixed strategies (p, q) is an ϵ -equilibrium if both players have no more than ϵ incentive to deviate. Formally, (p, q) is an ϵ -equilibrium if for all rows i , we have $e_i^T R q \leq p^T R q + \epsilon$, and for all columns j , we have $p^T C e_j \leq p^T C q + \epsilon$.

Definition 2 A pair of mixed strategies (p, q) is a well supported ϵ -equilibrium if for any $i \in \text{supp}(p)$ (i.e., i s. t. $p_i > 0$) we have $e_i^T R q \geq p^T R q - \epsilon$, for all j ; similarly, for any $i \in \text{supp}(q)$ (i.e., i s. t. $q_i > 0$) we have $p^T C e_i \geq p^T C e_j - \epsilon$, for all j .

Definition 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 $|R_{i,j} - R'_{i,j}| \leq \alpha$ and $|C_{i,j} - C'_{i,j}| \leq \alpha$ for all $i, j \in \{1, \dots, n\}$.

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

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) .

Throughout this paper we use “log” to mean log-base-e.

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. Formally:

Definition 5 *A game G satisfies the (ϵ, Δ) stability to perturbations 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 $(\epsilon, 1)$ -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 stability to approximations:

Definition 6 *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 stability to perturbations condition is equivalent to the well supported approximation stability condition. Specifically:

Theorem 1 *A game satisfies the well supported $(2\epsilon, \Delta)$ -approximation stability condition if and only if it satisfies the (ϵ, Δ) -stability to perturbations condition.*

Proof: Consider an $n \times n$ bimatrix game specified by R and C and assume it satisfies the well supported $(2\epsilon, \Delta)$ -approximation stability condition; we show it also satisfies the (ϵ, Δ) -stability to perturbations condition. Consider $R' = R + \Gamma$ and $C' = C + \Lambda$, where $|\Gamma_{i,j}| \leq \epsilon$ and $|\Lambda_{i,j}| \leq \epsilon$, 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 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 $e_j^T R q + e_j^T \Gamma q \leq v_R$, so $e_j^T R q \leq v_R + \epsilon$, for all j . On the other hand we also have $e_i^T R q = e_i^T R' q - e_i^T \Gamma q \geq v_R - \epsilon$ for all $i \in \text{supp}(p)$. Therefore, $e_i^T R q \geq e_j^T R q - 2\epsilon$, for all $i \in \text{supp}(p)$ and for all j . Similarly we can show $p^T C e_i \geq p^T C e_j - 2\epsilon$, for all $i \in \text{supp}(q)$ 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) -stability to perturbations 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

⁸Note that the entries of the perturbed game are not restricted to the $[0, 1]$ interval, and are allowed to belong to $[-\epsilon, 1 + \epsilon]$. 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 $(\alpha\epsilon, \Delta)$ stable to perturbations. Theorem 1 provides further evidence that this definition is proper.

in this game. Let us define matrices R' and C' such that $e_i^T R' q = \max_{i'} e_{i'}^T R q - \epsilon$ for all $i \in \text{supp}(p)$ and $e_i^T R' q \leq \max_{i'} e_{i'}^T R q - \epsilon$ for all $i \notin \text{supp}(p)$, $p^T C' e_j = \max_{j'} p^T C e_{j'} - \epsilon$ for all $j \in \text{supp}(q)$ and $p^T C' e_{j'} \leq \max_{j'} p^T C e_{j'} - \epsilon$ for all $j \notin \text{supp}(q)$. Since (p, q) is a well supported 2ϵ Nash equilibrium we know this can be done such that $|(R' - R)_{i,j}| \leq \epsilon$ and $|(C' - C)_{i,j}| \leq \epsilon$, for all i, j (in particular, we have to add quantities in $[-\epsilon, \epsilon]$ to all the elements in rows i of R in the support of p and subtract quantities in $[0, \epsilon]$ 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) -stability to perturbations 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. ■

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

$$R = \begin{bmatrix} 1 & 1 \\ 1 - \epsilon_0 & 1 - \epsilon_0 \end{bmatrix} \quad C = \begin{bmatrix} 1 & 1 - \epsilon_0 \\ 1 & 1 - \epsilon_0 \end{bmatrix}$$

For $\epsilon < \epsilon_0$ this game satisfies the well supported $(\epsilon, 0)$ -approximation stability condition, but does not satisfy (ϵ, Δ) -approximation stability for any $\Delta < \epsilon/\epsilon_0$. To see this note that $e_1^T R q = 1$, $e_2^T R q = 1 - \epsilon_0$, $p^T C e_1 = 1$, and $p^T C e_2 = 1 - \epsilon_0$ 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 $(\epsilon, 0)$ -approximation stable. On the other hand, the pair of mixed strategies (p, q) with $p = (1 - \epsilon/\epsilon_0, \epsilon/\epsilon_0)^T$ and $q = (1 - \epsilon/\epsilon_0, \epsilon/\epsilon_0)^T$ is an ϵ -Nash equilibrium. The distance between (p, q) and the unique Nash is ϵ/ϵ_0 , thus this game is not (ϵ, Δ) -approximation stable for any $\Delta < \epsilon/\epsilon_0$.

Interestingly, the notion of approximation stability which is a restriction of the well-supported approximation stability and of stability to perturbations conditions is a relaxation of the stability condition considered by Awasthi et al. [2] 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 stability conditions given in Definitions 3 and 4 to be a reversal of quantifiers that asks 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^*) . Formally:

Definition 7 *A game G satisfies the strong (ϵ, Δ) stability to perturbations 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 1 applies to the strong versions of the definitions as well. We also note that our generic upper bounds in Section 4 will apply to the most relaxed version (perturbation-stability) and our generic lower bound in Section 5 will be apply to the most stringent version (strong approximation stability).

Range of parameters As shown in [2], 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 $\Delta \geq \epsilon$. We can show that if a game \mathcal{G} satisfies the (ϵ, Δ) -approximation stability and if the union of all Δ -balls around all Nash equilibria does not cover the whole space,⁹ then we must have $3\Delta \geq \epsilon$ – see Lemma 5 in Appendix B. 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 $\binom{n}{2}$ equilibria (all $(1/2e_i + 1/2e_j, 1/2e_i + 1/2e_j)$) and it is well-supported $(\epsilon, 2\epsilon)$ -approximation stable for all $\epsilon < 1/6$. Note that it does not satisfy strong (well-supported) stability because it has multiple very distinct equilibria. For further examples see Lemma 2 in Section 5, as well as Appendix C.

4 Equilibria in Stable Games

In this section we show we can leverage the structure implied by stability to perturbations to improve over the best known generic bound of [18]. We start by considering ϵ and Δ as given. We can show:

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

This improves by a factor $O(1/(\Delta^2 \log(1 + \Delta^{-1})))$ in the exponent over the bound of [18] for games satisfying these conditions (and reduces to the bound of [18] in the worst-case when $\Delta = 1$).

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 $\Omega((\epsilon/\Delta)(\log n)^{-1/2})$. This in turn can be used to show that each of p^*, q^* has all but at most 8Δ of its total probability mass is concentrated in a set (which we call the high part) of size $O((\Delta/\epsilon)^2 \log(1 + 1/\Delta) \log n)$. Once the desired concentration is proven, we can then perform a version of the [18] 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.¹⁰ This is done through our key lemma, Lemma 1 below. In particular, Lemma 1 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

⁹If 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.

¹⁰We note that [2] 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.

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 1 synthesizes the key points of this argument and it is proven through a careful probabilistic argument.

In the following we consider $c = (56)^2$ and let $c' = 27$.

Lemma 1 *Let us fix ϵ and Δ , $0 \leq \epsilon \leq \Delta \leq 1$. Let \tilde{p} be an arbitrary distribution over $\{1, 2, \dots, n\}$. Let $S = c(\Delta/\epsilon)^2 \log n$ and fix $\beta \leq 1$ such that $1 - \beta \geq 8\Delta$. Assume that the entries of \tilde{p} can be partitioned into two sets H and L such that $\|\tilde{p}_L\|_1 = 1 - \beta$, $\|\tilde{p}_H\|_1 = \beta$, $\|\tilde{p}_L\|_2^2 \leq \frac{(1-\beta)^2}{S}$. 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 $k_2 \leq n^{c'(\Delta/\epsilon)^2}$. Then there exists \tilde{p}' with $\text{supp}(\tilde{p}') \subseteq \text{supp}(\tilde{p})$ such that:*

1. $d(\tilde{p}, \tilde{p}') = 3\Delta$ and
2. $\tilde{p}' \cdot v_i \leq \tilde{p} \cdot v_i + \epsilon$ for all $i \in \{1, \dots, k_1\}$.
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. Let us define the random variable $X_i = 1$ with probability $1/2$ and $X_i = 0$ with probability $1/2$. Define

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

We have $E[\sum_{i \in L} \tilde{p}_i X_i] = (1 - \beta)/2$. By applying McDiarmid's inequality (see Theorem 6) and using the fact that $\|\tilde{p}_L\|_2^2 \leq \frac{(1-\beta)^2}{S}$, we obtain that with probability at least $3/4$ we have both:

$$\left| \sum_{i \in L} \tilde{p}_i X_i - \frac{1 - \beta}{2} \right| \leq \frac{1 - \beta}{12} \quad \text{and} \quad \left| \sum_{i \in L} \tilde{p}_i(1 - X_i) - \frac{1 - \beta}{2} \right| \leq \frac{1 - \beta}{12}. \quad (1)$$

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 $\{v_1, \dots, v_{k_1}\}$. 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 $(\frac{1-\beta}{28})(\frac{\epsilon}{\Delta})$ 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 $-(\frac{\epsilon}{\Delta})^2(\frac{1-\beta}{28})^2/\sum_{i \in L} \tilde{p}_i^2 \leq -(\frac{\epsilon}{28^2}) \log n$. Also, we have $\mathbf{E}[Z_2] = \mathbf{E}[Z_4] = \frac{1-\beta}{2}$, and $\mathbf{E}[Z_1] = \mathbf{E}[Z_3]$. So, we get that with probability at least $1 - 1/n^3$ we have

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

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

$$\tilde{p}' \cdot v \leq \tilde{p} \cdot v + 3.1\Delta \left(\frac{\epsilon}{7\Delta} + \frac{\epsilon}{7\Delta} \right),$$

yielding the desired bound $\tilde{p}' \cdot v \leq \tilde{p} \cdot v + \epsilon$. Applying the union bound over all $i \in \{1, \dots, k_1\}$ we obtain that the probability that there exists v in v_1, \dots, v_{k_1} such that $\tilde{p}' \cdot v \geq \tilde{p} \cdot v + \epsilon$ is at most $1/3$.

Consider an arbitrary distribution p in $\{p_1, \dots, p_{k_2}\}$. Assume that $p = \tilde{p} + g$. By definition, we have:

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

where the first inequality follows from applying relation (1) to the denominators, and the last equality 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 6) to argue that with high probability Z is within $2\Delta/5$ of its expectation. Note that $c_i = \frac{6\Delta \tilde{p}_i}{1-\beta}$. Therefore:

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

This then implies that

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

By the union bound we get that the probability that there exists a p in $\{p_1, \dots, p_{k_2}\}$ 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 (Theorem 2): 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 $O((\Delta/\epsilon)^2 \log(1 + 1/\Delta) \log n)$. 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 $\forall i \in L, \Pr[i] \leq \frac{1}{S} \Pr[L]$ for S as in Lemma 1, 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 $\|p_L\|_2^2 \leq \frac{(1-\beta)^2}{S}$. Let $v_{i,j} = C(e_i - e_j)$. Since (p^*, q^*) is a Nash equilibrium we know that $p^* \cdot v_{i,j} \leq 0$ for all i and for all $j \in \text{supp}(q^*)$.

By Lemma 1 there exists \tilde{p}' such that (1) $d(p^*, \tilde{p}') = 3\Delta$, (2) $\tilde{p}' \cdot v_{i,j} \leq p^* \cdot v_{i,j} + \epsilon$ for all i and for all $j \in \text{supp}(q^*)$ and (3) $d(\tilde{p}', p_i) > d(p^*, p_i) - \Delta$ for all $i \in \{1, \dots, k\}$ (here k is the number of equilibria of the game). By (2) we have that $\tilde{p}' \cdot v_{i,j} \leq \epsilon$ for all i and for all $j \in \text{supp}(q^*)$, 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 $\text{supp}(\tilde{p}') \subseteq \text{supp}(p^*)$ 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 equilibria. 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 equilibria 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. We clearly have $1 - \beta \leq 8\Delta$; moreover, it is easy to show that the set H has at most $S \log(1 + (8\Delta)^{-1})$ elements. The key idea is that since $1 - \beta \leq 8\Delta$, we can now apply the sampling argument of [18] to L with accuracy parameter $O(\epsilon/\Delta)$ 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 [18] to p_L , we have that by sampling a multiset \mathcal{X} of S elements from $L = \text{supp}(p_L)$, we are guaranteed that for any column e_j , we have:

$$|(U_{\mathcal{X}})^T C e_j - p_L^T C e_j| \leq (\epsilon/8\Delta),$$

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:

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

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

Corollary 1 *Let us fix ϵ and Δ , $0 \leq \epsilon \leq \Delta \leq 1$. Let G be a game with at most $n^{O((\Delta/\epsilon)^2)}$ Nash equilibria satisfying the well supported (ϵ, Δ) -approximation stability condition (or the $(\epsilon/2, \Delta)$ -stability to perturbations condition).*

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

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

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

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

$2\epsilon/3$ -Nash equilibrium of G' . Since G is an L_∞ $\epsilon/6$ -perturbation of G' , this will be a well-supported ϵ -Nash equilibrium of G . ■

Corollary 1 improves by a factor $O(1/(\Delta^2 \log(1 + \Delta^{-1})))$ in the exponent over the bound of [18] 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., $\epsilon = 1/\sqrt{n}$, $\Delta = 10\epsilon$ – in this case we obtain an improvement of $O(n/\log(n))$ in the exponent over the bound of [18].

The proof of Theorem 2 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 . Formally:

Theorem 3 *Let us fix ϵ and Δ , $0 \leq \epsilon \leq \Delta \leq 1$. Consider a game G with at most $n^{O((\Delta/\epsilon)^2)}$ Nash equilibria which satisfies the well supported (ϵ, Δ) -approximation stability condition (or the $(\epsilon/2, \Delta)$ -stability to perturbations 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 $O((\Delta/\epsilon)^2 \log(1 + \Delta^{-1}) \log n)$.*
- (2) *For any game G' with L_∞ distance $\epsilon/2$ of G , any Nash equilibrium in G' is 9Δ -close to a pair of mixed strategies each with support of size $O((\Delta/\epsilon)^2 \log(1 + \Delta^{-1}) \log n)$.*

So far in Theorem 2 and Corollary 1 we have considered ϵ and Δ fixed. It is also interesting to consider games where the stability conditions hold uniformly for all ϵ small enough. We call such games uniformly stable games. Formally:

Definition 8 *Consider $t \geq 1$. We say that a game is t -uniformly stable to perturbations (or t -uniformly well supported approximation stable) if there exists $\epsilon_0 = 1/\text{poly}(n)$ such that for all $\epsilon \leq \epsilon_0$, G satisfies $(\epsilon, t\epsilon)$ stability to perturbations (or the well supported $(\epsilon, t\epsilon)$ approximation stability).*

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

Corollary 2 (1) *Let $t = O(\text{poly}(\log(n)))$. There is a FQPTAS to find approximate-equilibria in games satisfying the t -uniform well supported approximation stability condition (or the t -uniform stability to perturbations 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 stability to perturbations 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 2 is especially interesting because the results of [8] 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 [9], such uniformly stable game 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 [14] satisfy this condition.

5 Converting the general case to the stable case

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

We start by stating a useful lemma that shows the existence of a family of modified matching pennies games that are strong approximation stable games with certain properties that will be helpful in proving our main lower bound.

Lemma 2 *Assume that $\Delta \leq 1/10$. Consider the games defined by the 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 the strong $(\Delta^2, 4\Delta)$ approximation stability condition. Moreover if (p, q) is a Δ^2 -Nash equilibrium, then we must have

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

See Appendix B for a proof. We now present the main result of this section.

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

Proof: The main idea is to construct a linear embedding of any given game into a larger game with one more strategy per player played with large probability, thereby compressing the incentives of the original game into a smaller scale. In particular, consider $\Delta = (8\epsilon)^{1/4}$ and consider a general game with payoff matrices R and C . Let us construct a new game with the payoff matrix for the row player R' defined as:

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

and for the column player C' defined as:

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

(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 2, the new game defined by R' and C' satisfies the strong $(\Delta^2, 4\Delta)$ approximation stability condition, which in turn implies satisfying $(\epsilon, 8\epsilon^{1/4})$ approximation stability (since $\epsilon \leq \Delta^2$ 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 $\Delta = (8\epsilon)^{1/4}$, this implies the desired result.

Let (p, q) be an $\Delta^4/8$ -equilibrium in the new game. By Lemma 2 (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:

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

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

$$\beta\Delta e_i^T (\Delta/2) R q_f > p_f (\Delta/2) R q_f + \Delta^4/8. \quad (2)$$

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 2 we get:

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

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 $1, 2, \dots, n$ 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:

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

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:

$$p_f^T ((\Delta/2)C + (\Delta/2)\mathbf{1}_{n,n})\alpha\Delta e_j > p_f^T ((\Delta/2)C + (\Delta/2)\mathbf{1}_{n,n})q_f + \Delta^4/8. \quad (3)$$

We also have $p_{n+1}(2\Delta, \dots, 2\Delta)\alpha\Delta e_j = p_{n+1}(2\Delta, \dots, 2\Delta)q_f$. By adding this quantity as well as the term $p_f(1, \dots, 1)q_{n+1}$ to the both sides of the inequality 3 inequality 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 $1, 2, \dots, n$ 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. ■

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

6 Stability in constant-sum games

Consider a game defined by R and C . Let

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

and

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

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 [22]. Lemmas 3 and 4 below characterize constant sum games satisfying approximation stability terms of properties of the space of mixed strategies for the row player and column player separately. Theorem 5 gives a polynomial time algorithm for determining the approximately best parameters for the strong approximation stability property for a given game.

Lemma 3 *If for any p that is Δ -far from \mathcal{P}^* there exists e_j such that $p^T R e_j < v_R - \alpha$ and for any q that is Δ -far from \mathcal{Q}^* there exists e_j such that $e_j^T C q < v_C - \alpha$, then the game satisfies the $(\alpha/2, \Delta)$ approximation stability.*

Proof: We show that 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}^* .¹¹ Assume WLOG that p is Δ -far from \mathcal{P}^* . We know that there exists e_j such that $p^T R e_j < v_R - \alpha$. We show that (p, q) cannot be an $\alpha/2$ Nash equilibrium. If $p^T R q < v_R - \alpha/2$, 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 $p^T R q \geq v_R - \alpha/2$, then $p^T C q \leq v_C + \alpha/2$, but we know that $p^T C e_j > v_C + \alpha$, so the column player would have an $\alpha/2$ incentive to deviate, as desired. ■

Lemma 4 *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. Moreover, if, in the former case, $\text{supp}(p) \subseteq \text{supp}(p^*)$ for some $p^* \in \mathcal{P}^*$, or if, in the latter case, $\text{supp}(q) \subseteq \text{supp}(q^*)$ for some $q^* \in \mathcal{Q}^*$, then the game cannot be well-supported (α, Δ) approximation stable.*

¹¹This 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 that note that if both p and q are close to \mathcal{P}^* and \mathcal{Q}^* respectively, then there exists a Nash equilibria $(p_1, q_1), (p_2, q_2)$ such that $d(p, p_1) \leq \Delta$ and $d(q, q_1) \leq \Delta$. By interchangeability, we get that (p_1, q_2) is a Nash equilibrium and we also have $d((p_1, q_2), (p, q)) \leq \Delta$.

Proof: Assume that there exists p that is Δ -far from \mathcal{P}^* such that $\min_j p^T R e_j \geq v_R - \alpha$. Let $p^* \in \mathcal{P}^*$ be such that $\text{supp}(p) \subseteq \text{supp}(p^*)$ if such p^* exists, else let $p^* \in \mathcal{P}^*$ be arbitrary; let (p^*, q^*) be an equilibrium that certifies that p^* that is in \mathcal{P}^* . We know that $p^{*T} R q^* = v_R$ and $p^{*T} C q^* = v_C$. Clearly, (p, q^*) that is Δ -far from (p^*, q^*) . We show that (p, q^*) is an α -Nash equilibrium, i.e., neither player has more than an α -incentive to deviate. We have $p^T R q^* \geq v_R - \alpha$ and $p^{tT} R q^* \leq v_R$ for any p' (since q^* is minimax optimal) so the row player has at most α -incentive to deviate. We also have $p^T C q^* \geq v_C$ (since q^* is minimax optimal) and the most the column player could get is $v_C + \alpha$ since $\min_j p^T R e_j \geq v_R - \alpha$, so $\max_j p^T C e_j \leq v_C + \alpha$. ■

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

Theorem 5 *Given any $0 < \alpha < 1$, we can use Algorithm 1 to whp determine Δ such that the game satisfies the $(\alpha/2, 2\Delta)$ strong approximation stability property, but not $(\alpha, \Delta/2)$ strong 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 (p', q') . In Step 3 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') . From the perspective of the row player, as shown in Lemma 3, if there exists e_j such that $p^T R e_j < v_R - \alpha$, then (p, q) cannot be an $\alpha/2$ Nash equilibrium for any q , so all $\alpha/2$ Nash equilibria must satisfy $p^T R e_j \geq v_R - \alpha$ for all j . As shown in Lemma 4, for any p such that $p^T R e_j \geq v_R - \alpha$, for all j we have that (p, q^*) is an α Nash equilibrium. Similarly for the column player. So all $\alpha/2$ -equilibria must be at distance at most Δ from (p', q') , where Δ is the output of Algorithm 1, and there exists an α -Nash equilibrium that is at distance at 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. Note that the running time is polynomial since we we perform steps (A), (B) at most $n^{O(1/\alpha^2)}$ times, so overall the running time is polynomial $n^{O(1/\alpha^2)}$. ■

In order to determine the approximate parameters for the strong well supported approximation stability property for a given game, we can adapt Algorithm 1 as follows. Given any $0 < \alpha < 1$, we can use Algorithm 1 to whp determine Δ_h . We then we re-run Algorithm 1, but in the LP in (A) we add the constraint that $p_i = 0$ for all $i \notin \text{supp}(p^*)$ (and similarly for q in the LP (B)) to get a value Δ_l . Then by lemmas 3 and 4, we are guaranteed that the game satisfies the $(\alpha/2, 2\Delta_h)$, but not the $(\alpha, \Delta_l/2)$ strong well supported approximation stability property.

Acknowledgments

We thank Avrim Blum, Dick Lipton, Yishay Mansour, Shanghua Teng, and Santosh Vempala for useful discussions. We also thank Vangelis Markakis for pointing [19] to us.

This research was supported in part by NSF grant CCF-0953192, ONR grant N00014-09-1-0751, and AFOSR grant FA9550-09-1-0538. This work was done in part while the first author was visiting Microsoft Research NE and while the second author was a member of Microsoft Research NE.

References

- [1] ADSUL, B., GARG, J., MEHTA, R., AND SOHONI, M. 2011. Rank-1 bi-matrix games: A homeomorphism and a polynomial time algorithm. In *STOC*.
- [2] AWASTHI, P., BALCAN, M. F., BLUM, A., SHEFFET, O., AND VEMPALA, S. 2010. On nash-equilibria of approximation-stable games. In *SAGT*.
- [3] AWASTHI, P., BLUM, A., AND SHEFFET, O. 2011. Center-based clustering under perturbation stability. *Information Processing Letters*.
- [4] BÁRÁNY, I., VEMPALA, S., AND VETTA, A. 2007. Nash equilibria in random games. *Random Structures & Algorithms* 31, 4, 391–405.
- [5] BILU, Y. AND LINIAL, N. 2010. Approximate clustering without the approximation. In *Innovations in Computer Science*. 332 – 341.
- [6] BOSSE, H., BYRKA, J., AND MARKAKIS, E. 2007. New algorithms for approximate nash equilibria in bimatrix games. In *WINE*.
- [7] CHEN, X. AND DENG, X. 2006. Settling the complexity of two-player nash equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*.
- [8] CHEN, X., DENG, X., AND TENG, S. 2009. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM* 56, 3.
- [9] DASKALAKIS, C. 2010. On the complexity of approximating a nash equilibrium. In *Annual ACM-SIAM Symposium on Discrete Algorithms*.
- [10] DASKALAKIS, C., GOLDBERG, P. W., AND PAPADIMITRIOU, C. H. 2009. The complexity of computing a nash equilibrium. *SIAM J. Comput.* 39, 1, 195–259.
- [11] DASKALAKIS, C., MEHTA, A., AND PAPADIMITRIOU, C. H. 2006. A note on approximate nash equilibria. In *WINE*.
- [12] DASKALAKIS, C., MEHTA, A., AND PAPADIMITRIOU, C. H. 2007. Progress in approximate nash equilibria. In *Proceedings of the 8th ACM Conference on Electronic Commerce*.
- [13] DEVROYE, L., GYORFI, L., AND LUGOSI, G. 1996. *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag.
- [14] DURLAUF, S. N. AND BLUME, L. E. 2008. *Game Theory*. Palgrave Macmillan.
- [15] FEDER, T., NAZERZADEH, H., AND SABERI, A. 2007. Approximating nash equilibria using small-support strategies. In *Proceedings of the 8th ACM conference on Electronic Commerce*.
- [16] KANNAN, R. AND THEOBALD, T. 2010. Games of fixed rank: A hierarchy of bimatrix games. *Econom. Theory* 42, 157–173.
- [17] KONTOGIANNIS, S. C., PANAGOPOULOU, P. N., AND SPIRAKIS, P. G. 2006. Polynomial algorithms for approximating nash equilibria of bimatrix games. In *WINE*.
- [18] LIPTON, R. J., MARKAKIS, E., AND MEHTA, A. 2003. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic Commerce*.
- [19] LIPTON, R. J., MARKAKIS, E., AND MEHTA, A. 2006. On stability properties of economic solution concepts. Manuscript.

- [20] NASH, J. F. 1951. Non-cooperative games. *Annals of Mathematics* 54, 286–295.
- [21] TSAKNAKIS, H. AND SPIRAKIS, P. 2007. An optimization approach for approximate nash equilibria. In *Proceedings of the 3rd International Workshop On Internet And Network Economics*.
- [22] VON NEUMANN, J. AND MORGENSTERN, O. 1994. *Theory of Games and Economic Behavior*. Princeton University Press.

A Standard Facts

We start by stating the McDiarmid inequality (see [13]) we use in our proofs:

Theorem 6 *Let Y_1, \dots, Y_n be independent random variables taking values in some set A , and assume that $t : A^n \rightarrow R$ satisfies:*

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

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

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

We now state a well known fact showing any pair of strategies that is sufficiently close to a Nash equilibrium is a sufficiently good approximate Nash equilibrium.

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

B Additional Proofs

Lemma 2 Assume that $\Delta \leq 1/10$. Consider the games defined by the matrices:

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

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

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

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

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

Also $p^T C e_{n+1} = p_1 + \dots + p_n$ and

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

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:

$$2\Delta q_{n+1} = (1 + \alpha_{i,1})q_1 + (1 + \alpha_{i,2})q_2 + \dots + (1 + \alpha_{i,n})q_n. \quad (4)$$

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

$$p_1 + \dots + p_n = p_1 \gamma_{1,j} + \dots + p_n \gamma_{n,j} + 2\Delta p_{n+1}. \quad (5)$$

Identities 4 and 5 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:

$$\frac{2\Delta}{1 + 2\Delta} \leq p_1 + \dots + p_n \leq \frac{2\Delta}{1 + \Delta} \quad \text{and} \quad \frac{2\Delta}{1 + 2\Delta} \leq q_1 + \dots + q_n \leq \frac{2\Delta}{1 + \Delta}. \quad (6)$$

To get the desired stability guarantee we now show that any Δ^2 -equilibrium must have $\Delta/2 \leq p_1 + \dots + p_n \leq 4\Delta$ and $\Delta/2 \leq q_1 + \dots + q_n \leq 4\Delta$. (This in turn implies that any Δ^2 -equilibrium must be at distance at most 4Δ from a Nash equilibrium satisfying relation (6).) 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 $p^T C e_{n+1} = \sum_{i=1}^n p_i \leq \Delta/2$. But the column player's payoff for a column $j \in \{1, \dots, n\}$ is:

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

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

$$p^T C e_j - p^T C q \geq (1/2)[2\Delta(1 - \Delta/2) - \Delta/2] > \Delta/2 > \Delta^2,$$

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 $e_1^T R q \geq (1/2)(1 - \Delta)$, but row's player payoff for row $n + 1$ is $e_{n+1}^T R q \leq (1/2)2\Delta = \Delta$. Thus in this case the row player has incentive to deviate at least:

$$e_1^T R q - e_{n+1}^T R q \geq (1 - \Delta/2)[1/2(1 - \Delta) - \Delta] > \Delta,$$

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 $p^T C e_{n+1} = \sum_{i=1}^n p_i \geq 4\Delta$, whereas the column player's payoff for a column $j \in \{1, \dots, n\}$ is:

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

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

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

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

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

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

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

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 $e_{n+1}^T R q \geq 2\Delta(1 - \Delta/2)$, but the row player's payoffs for rows $1, \dots, n$ are $\leq \Delta/2$. So, if $p_{n+1} < 1 - \Delta$ then the row player has incentive to deviate at least:

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

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 $p^T C e_{n+1} = \sum_{i=1}^n p_i \leq \Delta$, but the column player's payoff for a columns $j \in \{1, \dots, n\}$ is:

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

So, the column player has incentive to deviate at least:

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

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 $e_{n+1}^T R q \leq 2\Delta(1 - 4\Delta)$, but row player's payoffs for rows $1, \dots, n$ are $\geq (4\Delta)(1 - \Delta)$. So, if $p_{n+1} > 1/2$ then Row has incentive to deviate at least:

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

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 $\Delta/2 \leq p_1 + \dots + p_n \leq 4\Delta$ and $\Delta/2 \leq q_1 + \dots + q_n \leq 4\Delta$, as desired. This concludes the proof. ■

Lemma 5 *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 $3\Delta \geq \epsilon$.*

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 1 we also have that this is a 3Δ Nash equilibrium. This then implies the desired result. ■

Lemma 6 *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 $\Delta \geq \epsilon/4$.*
- (b) *If \mathcal{G} satisfies the strong (ϵ, Δ) -stability to perturbations condition, then we must have $\Delta \geq \epsilon/8$.*

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 WLOG that p is a mixed strategy. Consider an α internal deviation of the row player, i.e., consider p' with $\text{supp}(p') \subseteq \text{supp}(p)$ such that $d(p, p') = \alpha$. Since $\text{supp}(p') \subseteq \text{supp}(p)$ we have $p'^T Rq = p^T Rq$. Since (p, q) is a Nash equilibrium we have $p^T C e_j = p^T C q \equiv v_C$ for all $j \in \text{supp}(q)$ and $p^T C e_j \leq v_C$ for all $j \notin \text{supp}(q)$. Since $d(p, p') = \alpha$ we have

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

for all j , so $p'^T C e_j \geq v_C - \alpha$, for all $j \in \text{supp}(q)$ and $p'^T C e_j \leq v_C + \alpha$, for all $j \notin \text{supp}(q)$. 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 $\Delta \geq \epsilon/4$, as desired. By Theorem 1, we immediately get (b) as well. ■

C Examples

To illustrate our notions of stability, we present two n by n games satisfying uniform stability to perturbations.

Example 1 A classic game from experimental economics is the public goods game which is defined 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 $i \in \{0, 1, \dots, n - 1\}$ and $j \in \{0, 1, \dots, n - 1\}$. This has payoffs ranging from 0 up to $0.75(n - 1)$, so to scale to the range $[0, 1]$ as we do in our paper, we multiply all the payoffs by $1/n$. I.e., if the Row player plays i and the Column player plays j then the payoff to the Row player is $[0.75j - 0.25i]/n$ and the payoff to the Column player is $[0.75i - 0.25j]/n$.

First note that this game is $(\epsilon, 0)$ stable to perturbations for all $\epsilon < 1/(8n)$. To see this note that without any perturbation, for any j and any $i \geq 1$ we have $e_0^T R e_j - e_i^T R e_j = 0.25i/n \geq 0.25/n$. 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 our game we get: $e_0^T R' e_j - e_i^T R' e_j \geq 0.25/n - 2\epsilon > 0$. 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 $(\epsilon, 0.99)$ stable for any $\epsilon > 1/(4n)$. To see this consider adding ϵ to $R[1, 0]$. I.e., if the Row player plays action 1 and the Column player plays action 0, then the payoff to Row player is $-0.25/n + \epsilon > 0$. 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, $R[1, j] = 1/2$ for all j , $R[i, i] = 1$ for all $i > 1$, $R[i, j] = 0$ for $i > 1, j \neq i$. and similarly $C[i, 1] = 1/2$, $C[j, j] = 1$ for $j > 1$, $C[i, j] = 0$ for $j > 1, i \neq j$. We claim that this game is well supported $(\epsilon, 2\epsilon)$ -stable for all $\epsilon < 1/6$, 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, $p^T C e_1 = 1/2$ and $p^T C e_i = p_i$ for $i > 1$. Note that if (p, q) is an well supported ϵ -Nash equilibrium and if $q_i < 1/2 - \epsilon$ 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 $< 1/2 - \epsilon$. 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 $q_i > 1/2 + \epsilon$ for $i > 0$ then both $p_i = 1$ and $q_i = 1$. If $q_i > 1/2 + \epsilon$ we have $e_i^T Rq = 1/2 + \epsilon$ 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 $i > 1, i \in \text{supp}(q)$, we have $1/2 - \epsilon \leq q_i \leq 1/2 + \epsilon$ and $1/2 - \epsilon \leq p_i \leq 1/2 + \epsilon$. Similarly, for any action $i > 1, i \in \text{supp}(p)$, we have $1/2 - \epsilon \leq q_i \leq 1/2 + \epsilon$ and $1/2 - \epsilon \leq p_i \leq 1/2 + \epsilon$. We have two cases. The first one is if there is exactly one action $i > 1$ in $\text{supp}(q)$. 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 $\text{supp}(q)$. 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.

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

Input: R, C , parameter α .

1. Solve for minimax optimal (p^*, q^*) .
2. Step 2: Apply the sampling procedure in [18] from (p^*, q^*) to get (p', q') with support of size $O((\log n)/\alpha^2)$ that is an α -Nash. Set $\Delta = 0$.
3. Find Δ as follows:

(A) For each partition of the support of p' into supp_+ and supp_- do:

i. Solve the following LP:

$$\begin{aligned} \max \Delta &= \sum_{i \in \text{supp}_+} (p_i - p'_i) + \sum_{i \in \text{supp}_-} (p'_i - p_i) + \sum_{i \in \{1, 2, \dots, n\} \setminus \text{supp}(p')} p_i \\ \text{s.t. } p_i &\geq p'_i \text{ for all } i \in \text{supp}_+ \\ p_i &\leq p'_i \text{ for all } i \in \text{supp}_- \\ p^T R e_j &\geq v_R - \alpha \text{ for all } j \end{aligned}$$

ii. If Δ is smaller than v , the value of the previous LP, reset Δ to be v .

(B) for each partition of the support of q' into supp_+ and supp_- do:

i. Solve the following LP:

$$\begin{aligned} \max \Delta &= \sum_{i \in \text{supp}_+} (q_i - q'_i) + \sum_{i \in \text{supp}_-} (q'_i - q_i) + \sum_{i \in \{1, 2, \dots, n\} \setminus \text{supp}(q')} q_i \\ \text{s.t. } q_i &\geq q'_i \text{ for all } i \in \text{supp}_+ \\ q_i &\leq q'_i \text{ for all } i \in \text{supp}_- \\ e_j^T C q &\geq v_C - \alpha \text{ for all } j \end{aligned}$$

ii. If Δ is smaller than v , the value of the previous LP, reset Δ to be v .

Output: Radius Δ .
