On Stability Properties of Economic Solution Concepts

Richard J. Lipton[*] *Vangelis Markakis*[†] *Aranyak Mehta*[‡]

Abstract

In this note we investigate the stability of game theoretic and economic solution concepts under input perturbations. The existence of many of these solutions is proved using fixed point theorems. We note that fixed point theorems, in general, do not provide stable solutions, and we show how some commonly studied solution concepts are indeed not stable under perturbation of the inputs.

1 Introduction

Recently, there is great interest in the computational complexity of games and more generally of economic problems. We have seen many interesting results on various game and economic problems. We will call both types of problems economic problems. They include: games, auctions, markets, pricing schemes, and many others.

As often is the case, many of these problems are "intractable": sometimes this can be proved based on standard complexity assumptions. Other times, it remains open whether or not they have efficient algorithms. In either case, we are led in a natural manner, to consider approximation algorithms.

The goal of this paper is point out a fundamental problem with approximation techniques for many economic problems. Our main discovery is: *many natural economic problems have a strong "instability"*. That is: small perturbations to the problems yield huge changes to the solutions. We call such problems *unstable*. This is especially disturbing, we believe, because much of the motivation of the area is based on the fact that "real-world" economic problems need to be solved.

The reason that this is upsetting is two-fold. First, what does it even mean to solve such a problem? Even an "exact" algorithm for such a problem is potentially useless. If the algorithm finds the exact answer, but the problem is unstable, then what does the answer mean? Imagine, a game that is being played by real companies. Suppose that Delta Airline's management team sets up a 2-player non-zero sum game to model a business decision that must be made. Even if they can exactly solve the game, what use is the solution strategy if small changes to the payoff matrix result in huge variations in the value of the game? Perhaps some of the entries are

[*]Georgia Institute of Technology, Atlanta, GA
[†]University of Toronto.
[‡]IBM Almaden Research Center, San Jose, CA

estimates of the price of oil, for example. What if the value of their game varies tremendously when the value of oil changes by pennies? In this case how useful is the exact answer? How useful is the model?

Second, while instability is important to exact algorithms, it is even more relevant to approximation algorithms. Here it is more problematic. What does it mean to approximate a value that is unstable?

These are the reasons we find our results interesting. The main result, again, is that we can show that many economic problems are unstable. However, not all are. We can also show that many other problems are stable. Right now we are unable to give a clear rationale for which problems fall into which class. We think that this leads to many interesting open problems.

There is one very interesting connection that we can make. Many economic problems use a "fixed-point" argument to prove the existence of the "solution". For example, the most famous of these is probably the proof that all games have a Nash Equilibrium [7]. Also the first proof by von Neumann that every 2-player zero-sum game has a min-max solution used a fixed-point argument. Fixed-point arguments abound throughout economic theory. They are also used in the famous proof that markets will always "clear".

When we say fixed-point theorems we usually mean the Brouwer Fixed Point Theorem (BFPT). Sometimes "stronger" versions of this theorem are used. But often BFPT suffices. Recall, BFPT states that any continuous map from a finite dimension cube to itself must have a fixed-point. One way to think about the BFPT is that it is a mapping B from such functions f to points $B(f) = x$, so that $f(x) = x$.

The key to our discovery of lack of stability is this: the mapping $B(f) = x$ is itself *not continuous*. That is small changes to the function f can lead to huge changes to the point x. For example (see figure 1) there is a continuous function $f : I \to I$ (I is the 1-dimesion cube, i.e. the interval [0,1]) so that one of its fixed point $B(f) = x$ vanishes after a tiny perturbation of the function f.

There are two points about this property of BFPT. First, the fact that $B(f)$ is not continuous does not seem to be well known. Second, the fact that so much of economic theory uses the BFPT to prove existence of "solutions" means that much of economic theory could be unstable. Note, it does not imply that it is. The difficulty is that while there may be a proof that shows that BFPT implies an economic result that does not imply instability. There could of course be a "better" proof that avoids the use of BFPT. We will see this exact situation shortly.

Finally, a word about the notion of stability and our results proved here. The proofs of the results themselves are not very complicated. The difficulty was in finding the right new questions. We do not claim that any proofs are very hard. However, we do believe that these results are quite interesting. They raise new fundamental questions about much of economic theory.

They also help answer partially some hard questions: For example, our methods show that it is quite unlikely for there to be a reduction from finding the equilibrium values of a non-zero sum game to finding the value of a zero-sum game. This holds even if one allows the zero-sum game to be exponentially larger that the non-zero sum game. Note, the recent work [5, 2] on Nash Equilibrium does not rule out such a reduction. They do, of course, supply strong

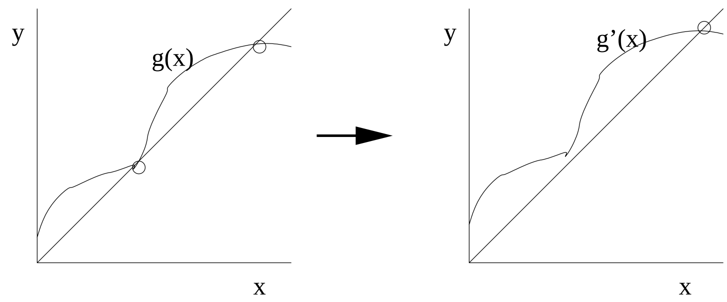

Figure 1: Fixed points are not stable.

evidence that Nash is hard. But, they do not rule out such exponential reductions between non-zero and zero sum games.

In some recent related work, a stability analysis has been initiated by Even-Dar *et al.*[6] for the dynamics of limit-order mechanisms in financial markets. The authors propose different models for limit-order mechanisms and study the stability properties of various quantities of interest in such markets. For discrete optimization problems, Bilu and Linial [1] have addressed the question of whether stable instances are easier to solve. Their definition of stability is different than ours.

2 Preliminaries

Let $\mathcal{P}$ be a problem which takes as input a vector of rational numbers $I = (x_1, ..., x_n)$ (e.g., if $\mathcal{P}$ is the problem of computing a Nash equilibrium, then I specifies the entries of the payoff matrices). We will denote by $\psi(I)$ the set of solutions to the problem. An element of $\psi(I)$ can be seen as another vector of numbers, corresponding to the quantitites that the problem asks us to compute. For example, if the problem at hand is an optimization problem, then $\psi(I)$ can be taken to be either the optimal value of the objective function or the value followed by a witness achieving this value. As a second example, if $\mathcal{P}$ is a 0-sum game, then an element of $\psi(I)$ is a vector that consists of the value of the game followed by the equilibrium probabilities of the players. In general, $\psi(I)$ can be empty. However, we will mainly focus on problems for which a solution always exists, i.e., $\psi(I) \neq \emptyset$ for every instance I of $\mathcal{P}$.

The following definition is an attempt to capture the notion of stability:

Definition 1 *A problem $\mathcal{P}$ is said to be stable in l_p-norm if there exists a constant c such that for any $\delta > 0$, for any two instances I, $\tilde{I}$ of $\mathcal{P}$ for which $||I - \tilde{I}||_p \leq \delta$, and for any solution $S \in \psi(I)$, there exists a solution $\tilde{S} \in \psi(\tilde{I})$ for which $||S - \tilde{S}||_p \leq c\delta$.*

In other words, if we perturb the instance slightly, the solution changes only slightly. If we are interested only in one (or in a subset) of the quantities that are included in a solution of $\mathcal{P}$, we can define $\mathcal{P}$ to be stable with respect to the desirable quantities, in a similar manner. We can also talk about an algorithm being stable. Given an algorithm $\mathcal{A}$ for $\mathcal{P}$ and an instance I, let $\mathcal{A}(I)$ denote the solution returned by the algorithm.

$$\begin{bmatrix} 1+e & e \\ 1 & e/2 \end{bmatrix} \begin{bmatrix} e & 0 \\ 0 & 1 \end{bmatrix} \longrightarrow \begin{bmatrix} 1-e & 0 \\ 1 & e/2 \end{bmatrix} \begin{bmatrix} e & 0 \\ 0 & 1 \end{bmatrix}$$

Figure 2: An Unstable Game. Here $e > 0$ is an arbitrarily small number.

Definition 2 *An algorithm $\mathcal{A}$ for a problem $\mathcal{P}$ is said to be stable in l_p-norm if there exists a constant c such that for any $\delta > 0$ and for any two instances I, $\tilde{I}$ of $\mathcal{P}$ such that $||I - \tilde{I}||_p \leq \delta$, we have $||\mathcal{A}(I) - \mathcal{A}(\tilde{I})||_p \leq c\delta$.*

We immediately have:

Proposition 1 *Let $\mathcal{A}$ be an algorithm for a problem $\mathcal{P}$. If $\mathcal{A}$ is a stable algorithm, then $\mathcal{P}$ is a stable problem.*

Similarly, we can define a stable reduction as follows.

Definition 3 *Consider a reduction f from problem $\mathcal{P}$ to problem $\mathcal{P}'$ taking every input I of $\mathcal{P}$ to an input $f(I)$ of $\mathcal{P}'$. The reduction f is said to be stable in l_p-norm if there exists a constant c such that for any $\delta > 0$, for any two instances I, $\tilde{I}$ of $\mathcal{P}$ for which $||I - \tilde{I}||_p \leq \delta$, we have $||f(I) - f(\tilde{I})||_p \leq c\delta$.*

We now investigate the stability properties of solution concepts from game theory and economics. From now on we focus on stability in l_∞-norm, i.e., we consider perturbing every input of the problem by a small amount.

3 Nash equilibria in bimatrix games

In a bimatrix noncooperative game, we are given a payoff matrix R for the row player and a matrix C for the column player. The most popular solution concept is that of a Nash equilibrium – a pair of strategies such that no player has an incentive to deviate unilaterally. We show here that games are not stable with respect to Nash equilibrium strategies, and thay are also not stable with respect to the payoff that players receive in a Nash equilibrium.

Theorem 2 *Nash equilibria in general sum bimatrix games are not stable with respect to equilibrium strategies and are also not stable with respect to equilibrium payoffs. This is true even when the games in question have unique Nash equilibria.*

Proof :

We exhibit a 2×2 game and an ϵ-perturbation of its payoff matrices, $(0 < \epsilon < 1)$, such that each game has a unique Nash equilibrium, but the equilibrium strategies and the payoffs to the two players in the Nash equilibria in the two games game are arbitrarily far away from each other.

The game is as follows (see left side of the Figure 2): the payoff matrix for the row player is $R_{11} = 1 + \epsilon$, $R_{12} = \epsilon$, $R_{21} = 1$, $R_{22} = \epsilon/2$. For the column player, $C_{11} = \epsilon$, $C_{12} = C_{21} = 0$, $C_{22} = 1$.

Since the first row of R is dominating, it follows that the only Nash equilibrium of the game is achieved when the row player chooses the first row and the column player chooses the first column. This results in a payoff tuple of $(1+\epsilon, \epsilon)$. Suppose we modify R_{11} to $1-\epsilon$, and R_{12} to 0 (right side of Figure 2). The second row is now dominating, so the only Nash equilibrium is when the row player chooses the second row and the column player chooses the second column and the payoffs in the Nash equilibrium of the game are $\epsilon/2$ for the row player and 1 for the column player. In this example while the perturbations are small, the entries of the game are also small. This is not necessary for such a construction, and we can easily modify the game to have no small entries. □

Proposition 1 immediately yields:

Corollary 3 *No stable algorithm can compute a player's Nash equilibrium payoff or the Nash equilibrium strategies in a general bimatrix game.*

Suppose we do not insist on playing exact nash equilibria but instead we also allow approximate equilibria. We will call a pair of strategies an ϵ-equilibrium if a player cannot gain more than ϵ by deviating from her strategy. We then have the following:

Theorem 4 *Let A, B and R, C be two non-zero sum games such that* $\max_{i,j}\{a_{ij} - r_{ij}\} \leq \delta$ *and* $\max_{i,j}\{b_{ij} - c_{ij}\} \leq \delta$. *Let (x, y) be any δ-Nash equilibrium in (A, B). Then (x, y) is a 3δ-Nash equilibrium in (R, C).*

Proof : We have $x^T Ay \geq x'^T Ay - \delta$, for all x' and $x^T By \geq x^T By' - \delta$, for all y'.

Let $D = R - A$. Then $x^T Ry - x^T Ay = x^T Dy$. But since x and y are probability distributions and $|d_{ij}| \leq \delta$, we have $|x^T Ry - x^T Ay| \leq \delta$.

Similarly, $|x'^T Ry - x'^T Ay| \leq \delta$ for all x'. This immediately gives $x^T Ry \geq x'^T Ry - 3\delta$.

Similarly we get $x^T Cy \geq x^T Cy' - 3\delta$ for all y'. □

The instability result of Theorem 2 is alleviated to some extent by Theorem 4. The former says that if players are constrained to play exact Nash equilibria, then small perturbations to the payoff matrices can cause a dramatic change in equilibrium strategies and equilibrium payoffs. However, the latter says that if players are allowed to play approximate equilibria then small perturbations do not cause much loss in the approximation. We believe that this is a good motivation to restrict ourselves to consider only approximate equilibria.

4 Zero-Sum Games

Finding a Nash equilibrium in a two-player game is a non-linear programming question. One may conjecture that this is the reason that Nash equilibria are unstable. However, this reasoning is incorrect:

Fact: *Linear programs in general are not stable with respect to their optimal vectors and optimal values. For example, it is easy to construct instances in which slight perturbations of the hyperplanes result in a feasible LP becoming infeasible.*

This leads us to the question of the stability of zero-sum games, which are known to be equivalent to linear programming. Surprisingly, we have the following stability result for zero-sum games.

Consider a 2-player zero-sum game determined by payoff matrix $\mathbf{A}$, which determines for every choice of strategies, the payment that the column player makes to the row player. The payoff to the row player in a minmax solution is known as the *value* of the game.

Theorem 5 *The value of a zero-sum game is stable.*

Proof :

Let $\mathbf{A}$ and $\mathbf{B}$ be two zero-sum games with payoffs bounded between -1 and 1, i.e. $a_{ij}, b_{ij} \in [-1, 1]$. Let $\delta = \max_{i,j}\{|a_{ij} - b_{ij}|\}$. Let u be the value of game $\mathbf{A}$ and w be the value of game $\mathbf{B}$. We will show that $|u - w| < \delta$, hence proving the theorem.

By writing the minmax solution as an LP in the standard manner we get that u is the optimum value of the following LP (1):

$$\begin{aligned} &\max v \\ &\mathbf{A}\mathbf{x} \geq v\mathbb{1} \\ &\sum_{i=1}^{n} x_i = 1 \\ &x_i \geq 0 \end{aligned}$$

Similarly w is the optimum value of the LP (2):

$$\begin{aligned} &\max v \\ &\mathbf{B}\mathbf{x} \geq v\mathbb{1} \\ &\sum_{i=1}^{n} x_i = 1 \\ &x_i \geq 0 \end{aligned}$$

Let $\mathbf{x}^*$ be a value of $\mathbf{x}$ which yields a value of u for v in LP (1). Let $\mathbf{D} = \mathbf{B} - \mathbf{A}$. Then

$$\mathbf{B}\mathbf{x}^* = (\mathbf{A} + \mathbf{D})\mathbf{x}^* = \mathbf{A}\mathbf{x}^* + \mathbf{D}\mathbf{x}^*$$

Since $\sum_{i=1}^{n} x_i^* = 1$ and $x_i^* > 0$, we have that $|(\mathbf{D}\mathbf{x}^*)_i| < \delta$. Hence

$$\mathbf{B}\mathbf{x}^* \geq \mathbf{A}\mathbf{x}^* - \delta\mathbb{1} \geq (u - \delta)\mathbb{1}$$

Hence $(\mathbf{x}^*, u - \delta)$ is a feasible solution for LP (2). Hence w, the value of game $\mathbf{B}$, is at least $u - \delta$. But symmetrically, we get that $u \geq w - \delta$. Hence we get $|u - w| < \delta$.

This shows that for any $\epsilon > 0$, perturbing the entries of a payoff matrix by at most ϵ, results in a game in which the value is ϵ-close to the original value. That is, zero-sum games are stable in l_∞-norm with respect to the value of the game. □

Recalling Theorem 2 and the definition of stable reductions, we get:

Corollary 6 *There is no stable reduction from a value of a general-sum game to the value of a zero-sum game. This holds even if the reduction is not polynomial time and even if the reduction is to an exponential size zero-sum game.*

Note that the statement of the corollary does not depend on any complexity assumptions. Although we proved that the value of a zero-sum game is stable, we know that zero-sum games are equivalent to linear programming, and since we have already noted that linear programming is unstable, we expect the following instability result for zero-sum games:

Theorem 7 *The minmax strategies of zero-sum games are not stable.*

Proof : To see this, consider the following 2×2 game given by $A_{11} = 1$, $A_{12} = 1 - \epsilon$, $A_{21} = 1 - \epsilon$, $A_{22} = 1 - 2\epsilon$. We perturb A to the following matrix B: $B_{11} = 1 - 2\epsilon$, $B_{12} = 1 - \epsilon$, $B_{21} = 1 - \epsilon$, $B_{22} = 1$. In the first game the first row dominates the second, hence the only minmax solution is that the row player plays row 1 and the column player plays column 2. In the second game, a small perturbation makes the second row dominate the first. Even though the value of the game is exactly the same as before, the minmax strategies now are row 2 for the row player and column 1 for the column player. □

5 Correlated equilibrium

Apart from Nash equilibria, several other "nicer" solution concepts have been proposed in the literature. One of the most important of these is the notion of Correlated Equilibrium due to Aumann, which we investigate next (see [8] for the definition of correlated equilibrium). The set of correlated equilibria is a superset of the set of Nash equilibria of a game. A nice property of the set of correlated equilibria is that it is a polytope. Hence a correlated equilibrium can be easily computed, and furthermore, one can optimize linear functions over the set of correlated equilibria, e.g., maximizing the sum of the player's payoffs. However, we have:

Theorem 8 *Correlated Equilibria and payoffs in correlated equilibria are not stable.*

Proof : We show that there exists a 2×2 game and an ϵ-perturbation of its payoff matrices, such that the correlated equilibria strategies and the equilibrium payoffs in the two games are arbitrarily far away from each other. This is the same example as in Theorem 2 – it is easy to see that the only correlated equilibrium in both the original and the perturbed game is the (unique) Nash equilibrium in each. □

6 Market equilibria

In a market equilibrium problem, a set of n agents come to an exchange market with an initial endowment of goods. Let m denote the number of goods and let $w_i = (w_{i1}, ..., w_{im})$ be the endowment of an agent. Each agent has a utility function u_i that specifies her preferences over different bundles of goods. Given a price vector π for the goods, each agent would like to obtain the bundle of goods x_i that maximizes her utility subject to the budget constraint $\pi x_i \leq \pi w_i$. An equilibrium in such a market is a price vector such that every agent buys a bundle of goods that maximizes her utility and all the goods are sold (market clears).

Under certain assumptions on the utilities of the agents (e.g., linearity or concavity) the celebrated work of Arrow and Debreu guarantees that an equilibirum always exists. We have recently seen a surge of activity towards finding efficient algorithms for computing equilibria. For a survey, see [3].

In this Section we restrict ourselves to utility functions for which an equilibrium always exists and we show that the problem is unstable for a certain type of nonlinear utilities. In particular, perturbing the utilities slightly may result in prices that are arbitrarily far from the original ones. To illustrate this we consider a special case of the Leontief utility model in which each agent comes with an initial endowment of 1 unit of a distinct good and the utility of an agent for a bundle $x_i = (x_{i1}, ..., x_{im})$ is the minimum of the linear functions

$$u_i(x_i) = min_k \frac{x_{ik}}{f_{ik}}$$

where f_{ik} is the coefficient of agent i with respect to good k. The input to the problem is the matrix $F = (f_{ik})$.

Theorem 9 *Market equilibrium prices in a Leontief economy are not stable.*

The proof is based on the reduction of [4] which shows that for an instance of the Leontief setting described above, market equilibria are in one-to-one correspondence with the Nash equilibria of a certain related game. The goal of finding such a reduction there was to show that computing market equilibrium prices in this model is hard (PPAD-complete). By looking at the construction of their reduction, we point out that not only is the problem hard, but the situation is worse – the problem is itself unstable.

Proof : Assume for contradiction that the problem is stable. We will show that this implies that Nash equilibria are stable, which contradicts the results of Section 3. Consider a bimatrix game with payoff matrices (R, C). Let $(\tilde{R}, \tilde{C})$ be a δ-perturbation of the matrices. In the reduction of [4], the instance of the Leontief economy that is produced from (R, C) is a block matrix F, where the upper right block is R, the bottom left block is C^T and the rest is 0. Clearly then, the δ-perturbed game $(\tilde{R}, \tilde{C})$ corresponds to a δ-perturbation $\tilde{F}$ of the instance F. Let π be an equilibrium price vector for F. By the definition of stability there exists an equilibrium $\tilde{\pi}$ for $\tilde{F}$ such that $|\pi_j - \tilde{\pi}_j| \leq c\,\delta$. As shown in [4] the Nash equilibria of the bimatrix game are in one-to-one correspondence with the market equilibria and moreover, given π, the Nash equilibrium strategies are proportional to the quantities:

$$\beta_j = \frac{\pi_j}{\sum_k f_{kj} \pi_k}$$

However, since $\tilde{\pi}$ is close to π, simple calculations can show that for any equilibrium of (R, C), we can find a nearby equilibrium of the perturbed game, contradicting the fact that Nash equilibria are not stable. □

We conjecture that when the utilities are linear, the problem is stable.

7 Cooperative Games and Solution Concepts that do not Always Exist

In this section we briefly note that we cannot hope to prove any nice stability properties for solution concepts that are not always guaranteed to exist. To illustrate this, we continue our study with *cooperative* games. One important solution concepts in such games is the notion of the *core*. For the relevant definitions see [8]. As with correlated equilibria, the set of the core outcomes of a cooperative game forms a polytope. However, the core of a game may be empty.

Theorem 10 *There exists an instance of a cooperative game with transferable payoff, which has a non-empty core such that an ϵ-perturbation of it results in a game with an empty core.*

Proof : We can carry out the proof by giving an explicit construction. However we can employ a more generic topological argument, which can be used for proving instability of other concepts. A cooperative game of n agents with transferable payoff is specified by $N = 2^n$ numbers, where each such number corresponds to the payoff that a certain coalition of agents receives. Hence we can view a game as a point in $\mathbb{R}^N$. Now consider the class of all cooperative games with n agents that have a non-empty core. This is essentially a set of connected components of $\mathbb{R}^N$. Pick such a connected component, say $\mathcal{U}$. At the boundary of $\mathcal{U}$ lie games that have a non-empty core. However, we can perturb these games appropriately so as to end up outside of $\mathcal{U}$ and not inside any of the remaining connected components. □

8 Conclusions and Open Questions

We have analyzed the stability of several important economic solution concepts. It is important to carry out such a stability analysis for other economic solution concepts. One example is: are market equilibrium prices for the case of linear utilities stable under perturbation of buyer utility functions? One approach here would be to use the explicit convex programs that have been suggested for this problem. Another approach would be to investigate if any of the recent algorithms for this problem are stable algorithms.

Another important question would be to investigate if there is any connection between stability and efficient computation. Consider for example the subclass of general sum noncooperative games that are stable under perturbations. Is there a poly-time algorithm for computing a Nash equilibrium in such games? Such a study has already been initiated in [1] for some combinatorial problems.

References

[1] Y. Bilu and N. Linial. Are stable instances easy? manuscript, 2006.

[2] X. Chen and X. Deng. Settling the complexity of two-player nash equilibrium. ECCC, 2005.

[3] B. Codenotti, S. Pemmaraju, and K. Varadarajan. Algorithms column: The computation of market equilibria. *ACM SIGACT News*, 35(4), 2004.

[4] B. Codenotti, A. Saberi, K. Varadarajan, and Y. Ye. Leontief economies encode nonzero sum two-player games. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[5] C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a nash equilibrium. In *STOC*, 2006.

[6] E. Even-Dar, S. Kakade, M. Kearns, and Y. Mansour. (in)stability properties of limit order dynamics. In *Annual ACM Conference on Electronic Commerce*, 2006.

[7] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[8] M. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, Cambridge, Massachusetts, 1994.