

The Price of Uncertainty

Maria-Florina Balcan
Microsoft Research
New England

Avrim Blum*
Dept of Computer Science
Carnegie Mellon University

Yishay Mansour†
Tel Aviv University
and Google Research

ABSTRACT

We study the degree to which small fluctuations in costs in well-studied potential games can impact the result of natural best-response and improved-response dynamics. We call this the *Price of Uncertainty* and study it in a wide variety of potential games (including fair cost-sharing games, set-cover games, routing games, and job-scheduling games), finding a number of surprising results. In particular, we show that in certain cases, even extremely small fluctuations can cause these dynamics to spin out of control and move to states of much higher social cost, whereas in other cases these dynamics are much more stable even to large degrees of fluctuation.

We also consider the resilience of these dynamics to a small number of Byzantine players about which no assumptions are made. We show again a contrast between different games. In certain cases (e.g., fair cost-sharing, set-covering, job-scheduling) even a single Byzantine player can cause best-response dynamics to transition to states of substantially higher cost, whereas in others (e.g., the class of β -nice games which includes routing, market-sharing and many others) these dynamics are much more resilient.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; F.2 [Analysis of Algorithms and Problem Complexity]

General Terms: Algorithms, Theory, Economics.

Keywords: Algorithmic Game Theory, Best Response Dynamics, Social Cost

1. INTRODUCTION

It is widely accepted that rational agents in competitive environments can be viewed as *utility maximizers*. Economic theory has gone to great lengths to justify this assumption, and deriving it from basic plausible axioms. Major milestones in this line of research include von-Neumann and Morgenstein [23], de Finetti [9] and Savage [19]. In essence, these results connect between agents'

*Supported in part by NSF grant CCF-0830540.

†Supported in part by the IST Programme of the EC, under the PASCAL Network of Excellence, by the Israel Science Foundation and by the United States-Israel Binational Science Foundation (BSF). This publication reflects the authors' views only.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'09, July 6–10, 2009, Stanford, California, USA.

Copyright 2009 ACM 978-1-60558-458-4/09/07 ...\$5.00.

preferences, likelihoods of events, and utility functions. (We should remark that there is a line of work in behavioral economics which challenges this approach, for example the well known works of Kahneman and Tversky [21].) In this work we explore how small fluctuations or uncertainties about agents' own utilities can substantially affect social welfare when players follow natural dynamics.

In many cases we can view the agents' utility functions as being based on measurements of some physical quantities. For example in job scheduling, the speed of each machine is a physical quantity which determines the load on each machine. An agent has to approximate this speed from its observation, where the simplest way is to consider the ratio of output quantity and time. Even if the output quantity is given, different agents might have (slightly) different measurements of time, which will cause them to compute slightly different machine speeds. Even the same agent might hypothesize different speeds for the same machine at different times, due to imperfection in its clock. We can model this phenomenon as follows. We assume that each machine has an absolute speed s , and at each time t each agent i observes a speed $s_i^t \in [s/(1 + \epsilon), s(1 + \epsilon)]$, for some uncertainty parameter $\epsilon > 0$. This modeling of uncertainty is reminiscent of the statistical query learning model of Kearns [12].

In other situations, even without uncertainty in measurement, the underlying game itself may exhibit small fluctuations in cost. For example, consider a transportation problem where each agent selects a route. We might model edges as having delay functions that are, say, linear in the amount of traffic on them, but in reality delays may also depend on external environmental factors which have been abstracted out of the model. Therefore delays would not be exactly identical at two different times even if the amount of traffic at those times is the same. We can again view this similarly: at each time t , resource j has cost $c_j^t \in [c_j/(1 + \epsilon), c_j(1 + \epsilon)]$ where c_j is the "base" cost of that resource, and ϵ is a degree-of-fluctuation parameter. In fact, this same issue of fluctuations in actual (in addition to perceived) costs may occur in job scheduling as well.

The question we are interested in is: can these fluctuations (in either perceived or actual costs) cause best-response and improved-response dynamics to spin out of control and move from high-quality states to states of much worse social welfare? What values of ϵ can different well-studied games tolerate? We focus on potential games, where such dynamics are especially natural, and define a new measure we call the *Price of Uncertainty (PoU)*, that models the effect these perturbations can cause. To define the PoU we assume at each step t an agent does a best (or improved) response to the cost function at time t . The PoU of a given game bounds the ratio of the initial social welfare to the social welfare of any state that can be reached in such a dynamic.¹ A small PoU

¹As mentioned above, it could be the *actual* cost of resource j at time t is slightly different from its cost at time 0, or else these could just be differences in measurement, but either way the effect on an

implies that the system can not deteriorate due to such perturbations in costs, and provides a certain degree of robustness to the system; a large PoU means that deterioration could be severe. Note that all the games we consider have only a small gap between potential and cost, so without any fluctuations, these dynamics would never cause social welfare to deteriorate substantially. In addition, all the games we consider have near-optimal equilibria (their price of stability is low), and they continue to do so even after perturbation: thus, the effect we are studying is not that the system moves to a poor state because good equilibria no longer exist, but rather whether small perturbations can lead natural dynamics astray.

Our Results: We analyze a number of well-studied potential games from this perspective, including cost-sharing [2], matroid games [1], job-scheduling [13, 8], and the class of β -nice games [3], proving both upper and lower bounds on how resilient they are to such perturbations under both best-response and improved-response dynamics. Our analysis shows a number of surprising distinctions between these games, as well as between these two dynamics. In some cases, we show even exponentially small perturbations can result in drastic increases in cost; in others, polynomially-small perturbations are sufficient to ensure good behavior, and finally some are resilient even to constant-sized fluctuations. For example, in fair cost-sharing games with many players of each type, we show that with best-response dynamics, the Price of Uncertainty is constant even for constant $\epsilon > 0$. However, with improved-response dynamics, even exponentially small fluctuations can cause an exponentially large increase in the cost of the system. On the other hand, with few players of each type, the game becomes less resilient, and constant-size fluctuations can cause even best-response dynamics to incur a PoU of $\Omega(n)$, where n is the number of players, moving from an equilibrium of cost $O(\mathbf{OPT})$ to one of cost $\Omega(n \cdot \mathbf{OPT})$, matching the Price of Anarchy of the game. For set-covering games, a natural special case of fair cost-sharing studied by [6] where players must choose which set to belong to and split the cost with others making the same choice, we show both best-response and improved-response dynamics have a logarithmic Price of Uncertainty for $\epsilon = O(1/(mn))$. However, for matroid games (a broad class that generalizes set-covering and many other games), while best-response dynamics continues to exhibit good behavior, improved-response dynamics again has exponentially large PoU even for exponentially small ϵ (exponential in the rank of the matroid). We also give a technically intricate lower bound of $\Omega(\epsilon\sqrt{n}/\log(1/\epsilon))$ on the PoU of best-response dynamics for set-covering games, showing that $\epsilon < \text{polylog}(n)/\sqrt{n}$ is necessary for polylogarithmic PoU. Finally, for the class of β -nice games, which for constant β includes congestion games with linear (or constant-degree polynomial) latency functions, market-sharing games, and many others [3], we show that for random order best-response dynamics, the price of uncertainty is $\tilde{O}(\beta)$. However, again for improved-response dynamics the PoU can be exponential. We also present results for job scheduling on unrelated machines and consensus games.

We also explore a different kind of robustness, which is a robustness to adversarial (Byzantine) players. This can be viewed as a *fault-tolerance* of the system to a few misbehaved agents. For our lower bounds, we concentrate on the case of a single adversarial (Byzantine) player, and measure to what degree can such a player can cause the system’s social welfare to deteriorate. We show

agent’s behavior is the same. Also, in general we assume that the initial state is arbitrary. It could be an equilibrium, in which case the PoU can be viewed as studying the stability of the equilibrium. Alternatively, the initial state might result from a change in the system (adding or removing a link in routing, or adding or removing a machine in job scheduling). In such a case the agents dynamics start from a more arbitrary state.

that for set cover games, a single Byzantine player can cause best-response dynamics to move from an equilibrium of cost $O(\mathbf{OPT})$ to one of cost $\Omega(n \cdot \mathbf{OPT})$. (This is as bad a situation as possible since the Price of Anarchy in such games is $O(n)$.) For job scheduling on two unrelated machines we show that a single Byzantine player can induce a dynamic where the cost increases from 1 to $\Omega(n)$, and that the same is true for consensus games. For our positive results, we allow any number of Byzantine players. We show that for job scheduling on identical machines the effect of Byzantine players is very limited. For β -nice games, under additional assumptions on the impact the Byzantine players can have on the social cost of any *given* state, and assuming random-order best-response dynamics, we can show that at any time t , the expected social cost can not be above $6\beta \cdot \mathbf{OPT} \cdot \text{GAP}$, where GAP bounds the ratio between potential and cost for the game (e.g., for fair cost sharing games $\text{GAP} = O(\log n)$).

Our work can be thought of as asking what kinds of fault-tolerance properties we can guarantee in multi-agent environments. That is, if the system is currently in a low-cost equilibrium state, when can we expect it to remain so even in the presence of slight perturbations in costs or in the presence of a small number of misbehaving (Byzantine) players. Our analysis also points out a fragility in standard potential-function arguments in cases where the underlying model is not quite perfect.

Related Work: Recent work on the “Price of Malice” and related notions have also considered the effect that Byzantine players can have in several natural games [4, 18, 15]. The focus of that work has been on the effect of such players on the quality of the worst Nash or coarse-correlated equilibria. In contrast, we are interested in the effect of such players on an initial state that may be much better than the worst equilibrium, for a wide class of potential games with a low gap between potential and cost (so that without any perturbations or Byzantine players, behavior would never degrade substantially).

2. THE MODEL

A game is denoted by a tuple $G = \langle N, (\mathcal{S}_i), (\text{cost}_i) \rangle$ where N is a set of n players, \mathcal{S}_i is the finite action space of player $i \in N$, and cost_i is the cost function of player i . The joint action space of the players is $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$. For a joint action $S \in \mathcal{S}$ we denote by S_{-i} the actions of players $j \neq i$, i.e., $S_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. Furthermore we denote by $S \oplus s'_i$ the state $(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$. The cost function of player i maps a joint action $S \in \mathcal{S}$ to a real non-negative number, i.e., $\text{cost}_i : \mathcal{S} \rightarrow \mathbb{R}^+$. Every game has associated a social cost function $\text{cost} : \mathcal{S} \rightarrow \mathbb{R}$ that maps a joint action to a real value. In the cases discussed in this paper the social cost is a simple function of the costs of the players. In particular, we discuss the sum, i.e., $\text{cost}(S) = \sum_{i=1}^n \text{cost}_i(S)$, and the max, i.e., $\text{cost}(S) = \max_{i=1}^n \text{cost}_i(S)$. (In the context of load balancing games we call the maximum social function the *makespan* social cost function.) The optimal social cost of a game G is $\mathbf{OPT}(G) = \min_{S \in \mathcal{S}} \text{cost}(S)$. We sometimes overload notation and use \mathbf{OPT} for a joint action S that achieves cost $\mathbf{OPT}(G)$.

Given a joint action S , the *Best Response (BR)* of player i is the set of actions $BR_i(S)$ that minimizes its cost, given the other players actions S_{-i} , i.e., $BR_i(S_{-i}) = \arg \min_{s \in \mathcal{S}_i} \text{cost}_i(s, S_{-i})$. Given a joint action S , the *Improved Response (IR)* of player i is the set of actions $IR_i(S)$ that have lower cost than its current action, i.e., $IR_i(S) = \{s \in \mathcal{S}_i | \text{cost}_i(s, S_{-i}) \leq \text{cost}_i(S)\}$.

A joint action $S \in \mathcal{S}$ is a *pure Nash Equilibrium (NE)* if no player $i \in N$ can benefit from unilaterally deviating to another action, namely, every player is playing a best response action in S , i.e., $s_i \in BR_i(S_{-i})$ for every $i \in N$. A *best (improved) response dynamics* is a process in which at each time step, some

player which is not playing a best response switches its action to a best (improved) response action, given the current joint action. In this paper we will concentrate on games in which any best (improved) response dynamics converges to a pure Nash equilibrium (which are equivalent to the class of ordinal potential games [14, 17]). Throughout the paper we denote by $\text{GAP}(G, n)$ the maximum ratio between the cost of a given joint action and the value of the potential function for it, where G is a game of n players.

Let $\mathcal{N}(G)$ be the set of Nash equilibria of the game G . The *Price of Anarchy* (PoA) is defined as the ratio between the maximum cost of a Nash equilibrium and the social optimum, i.e., $(\max_{S \in \mathcal{N}(G)} \text{cost}(S)) / \text{OPT}(G)$. The *Price of Stability* (PoS) is the ratio between the minimum cost of a Nash equilibrium and the social optimum, i.e., $(\min_{S \in \mathcal{N}(G)} \text{cost}(S)) / \text{OPT}(G)$.

In this paper we introduce and study the *Price of Uncertainty* (PoU). We consider three different variations.

Adversarial Model: Consider a game G , from a given class of games \mathcal{G} , where the agents start at some given initial configuration S_0 (which might be a Nash equilibrium or not). Now we progress in phases, where in phase t the following occurs. First, at time t , the adversary perturbs the costs of G by a small multiplicative factor from their initial values, so that for any S and j we have $\text{cost}_j^t(S) \in [\text{cost}_j(S)/(1+\epsilon), (1+\epsilon)\text{cost}_j(S)]$.² Then the adversary picks an agent i who performs a best (improved) response, and the new configuration is S_t .

Our main concern is to upper bound $\text{cost}(S_t)/\text{cost}(S_0)$ as a function of ϵ and the class of games \mathcal{G} . More precisely, let us define $\text{PoU}_{BR}(\epsilon, G) = \max \text{cost}(S_t)/\text{cost}(S_0)$, where the max operator is over the initial configuration S_0 , the number of time steps t , and a dynamics of t time steps which includes the selection of a player $i \in N$ and the selection of its best response at each time step. For a class of games \mathcal{G} , let $\text{PoU}_{BR}(\epsilon, \mathcal{G})$ be $\max_{G \in \mathcal{G}} \text{PoU}_{BR}(\epsilon, G)$. We define similarly PoU_{IR} for improved response.

Random Order Model: This model is similar to the adversarial Model, except that at each time step a random agent $i \in N$ is selected. We now care about the expected cost at time t . We remark that although the player is selected at random, its action is selected as an adversarial best (improved) response to an adversarially selected perturbation of its cost. In this case we denote the price of uncertainty by PoU_{BR}^R .³

Byzantine Model: In this model, rather than perturbing costs, the adversary instead has control over a small number of Byzantine agents. At each time step t , the adversary moves the Byzantine agents arbitrarily, and then an agent $i \in N$ is selected (either adversarially or at random), who then performs a best-response. Thus, while in the other models the adversary can perturb all costs by a small amount, in this model the adversary can influence only a few players, but for those players it has full control. This implies that the adversary can typically influence the costs of only a few resources at a time (those used by the Byzantine players) by an amount that depends on the game and the current joint action of the players.

One can view the adversarial model as a directed graph, where the nodes are the possible joint actions. There is a directed edge from S to S' if they differ in the action of only one player i , and

²We require the adversary's perturbations to be consistent with the class \mathcal{G} . For congestion games (see Section 2.1) this means the adversary may perturb the costs of each resource by a $1 + \epsilon$ factor; for job-scheduling, the adversary may perturb the cost of the machines. Thus, at each time t , the resulting game G_t remains a potential game.

³For improved response, the adversarial and the random order models are identical, since the adversary can make the random players select the same action until his desired player is selected.

the adversary can cause player i to move from its action in S to its action in S' : for improved response dynamics, this means perturbing costs so that $\text{cost}_i^t(S') \leq \text{cost}_i^t(S)$, and for best response dynamics this means $\text{cost}_i^t(S')$ is the minimum cost of any state player i can reach unilaterally from S . Given this graph, for each joint action S let $V(S)$ be the set of nodes reachable from it. The PoU then bounds the ratio between the social cost in S and the maximum social cost of any joint action reachable from S , i.e., $\max_S \max_{v \in V(S)} \text{cost}(v) / \text{cost}(S)$.

2.1 Classes of Games Studied

In this paper we extensively study *congestion games* and important sub-classes of them. A congestion game G is defined by a tuple $(N, R, (S_i), (d_i))$, where N is the set of n players, R is the set of m resources, the action of player i is $S_i \subseteq 2^R$, and the goal of player i is to play a strategy S_i that minimizes its individual cost cost_i . The cost $\text{cost}_i(S)$ of player i is given by $\sum_{r \in S_i} d_r(n_r(S))$, where $n_r(S)$ is the number of players sharing resource r in state S and $d_r : N \rightarrow N$ is a delay function associated with resource r . Rosenthal [17] shows that every congestion game possesses at least one Nash equilibrium by considering the potential function $\Phi(S) = \sum_{r \in R} \sum_{i=1}^{n_r(S)} d_r(i)$. We call a congestion game *symmetric* if all the players share the same set of strategies, otherwise we call it asymmetric. Specific classes of congestion games that we study in this work are cost-sharing games, matroid congestion games, β -nice games, and consensus games. We define all these in their corresponding sections, namely Sections 4.1, 4.2, 5.1 and 4.3. For games with $\Phi(S) \geq \text{cost}(S)$ we define $\text{GAP}(G, n)$ as $\max_S \Phi(S) / \text{cost}(S)$. More generally, for games such that $\Phi(S) \in [c_1 \text{cost}(S), c_2 \text{cost}(S)]$ we define $\text{GAP}(G, n) = c_2 / c_1$, where we assume $c_1 \leq 1 \leq c_2$.

Another class of games we study are *load balancing games* (see [16]) which we define in Section 4.4.

3. PRELIMINARIES

We start with a few simple observations regarding the price of uncertainty in general, and for congestion games in particular. First, note that for $\epsilon = 0$ we get the “standard” best (improved) response dynamics. In this case the PoU is simply asking by how much can the social welfare deteriorate, assuming that all the players are implementing best (improved) response dynamics (even this can be higher than the PoA, see Theorem 3.1). Our first observation is that even for $\epsilon = 0$ the PoU is at least the Price of Stability. This follows since we can start at S_0 as the social optimal configuration, and any best response dynamics will reach some equilibrium S_t .

FACT 3.1. *For any $\epsilon \geq 0$ we have:*

$$\text{PoU}_{IR}(\epsilon, \mathcal{G}) \geq \text{PoU}_{BR}(\epsilon, \mathcal{G}) \geq \text{PoU}_{BR}(0, \mathcal{G}) \geq \text{PoS}(\mathcal{G}).$$

For fair cost-sharing games, Fact 3.1 implies an $\Omega(\log n)$ bound, due to the price of stability results [2]. Second, one should expect the ratio to also depend on the magnitude of ϵ . For example, for any given game, for sufficiently small ϵ , the perturbations of the adversary would have no real effect, and the agents would simply follow some best (improved) response dynamics from the initial configuration. More specifically:

FACT 3.2. *For any game class \mathcal{G} there is an $\epsilon_0 > 0$ such that for any $\epsilon < \epsilon_0$, $\text{PoU}_{IR}(\epsilon, \mathcal{G}) = \text{PoU}_{IR}(0, \mathcal{G})$ and $\text{PoU}_{BR}(\epsilon, \mathcal{G}) = \text{PoU}_{BR}(0, \mathcal{G})$*

Again, for fair cost sharing games, since the social cost of any configuration S is at most a logarithmic factor from the value of the potential function [2], this would give an $O(\log n)$ upper bound

for $\epsilon = 0$, i.e., $PoU_{IR}(0, FCSG) = O(\log n)$. In exact potential games, an immediate observation is:

FACT 3.3. *In any exact potential game after t steps the potential function increases by at most $(1 + \epsilon)^{2t}$.*

Clearly, Fact 3.3 implies that the potential function increases by at most $(1 + \epsilon)^{2L}$, where L is the number of configurations of players; for congestion games, since players choose subsets of resources, $L \leq 2^{mn}$.

It is interesting to note that there exists \mathcal{G} such that $PoU_{BR}(0, \mathcal{G})$ is larger by a multiplicative $\Omega(\log n)$ factor than $PoA(\mathcal{G})$. In particular we can show the following.

THEOREM 3.1. *Let \mathcal{G} be the class of market-sharing games. Then $PoA(\mathcal{G}) = 2$ while $PoU_{BR}(0, \mathcal{G}) = \Theta(\log n)$.*

PROOF. The bound on the PoA is from [22]. For the lower bound on PoU, consider players $\{2, \dots, n\}$ where player i can select between a dedicated site s_i with benefit $1/i$ and a shared site s_1 with benefit 1. Initially, each player i uses its dedicated site s_i , and the social welfare benefit is $\ln n$. Now let the players perform best response in an increasing order of the indices. Player i has a benefit of $1/i$, and the benefit from moving to s_1 is $1/(i-1)$, so it prefers to move to s_1 . This implies that at the end of the sequence the social welfare benefit is 1. The upper bound follows immediately from the fact that the gap between potential and cost in this game satisfies $GAP = O(\log n)$. \square

Note that in such cases we are willing to lose the $\log n$ factor, and we are interested in for what values of ϵ the value of $PoU_{BR}(\epsilon, \mathcal{G})$ is not much larger than $PoU_{BR}(0, \mathcal{G})$.

Finally, we point out a simple relationship between the perturbation and Byzantine models. Consider a class of games \mathcal{G} such that a single player cannot change the cost of any given state S to any player i by more than a factor of α , and whose effect is monotone (for example, for fair cost-sharing, a new player can reduce the cost of a given state S to a player by at most a factor of 2, and in any state it cannot increase the total cost of the other players). Then, an adversary with $\epsilon = \sqrt{\alpha} - 1$ can simulate the effect of a Byzantine player on best-response (or improved-response) dynamics. This implies that any lower bound for a single Byzantine player (such as in Theorems 6.2, 6.3 and 6.5) translates to a lower bound on $PoU_{BR}(\sqrt{\alpha} - 1, \mathcal{G})$.

4. ADVERSARIAL ORDER

In this section we present our results in the adversarial model and give upper and lower bound on PoU_{BR} and PoU_{IR} for a number of well-studied classes of games. We begin by considering set-cover games, a natural type of cost-sharing game studied in [6], showing that both best-response and improved-response dynamics are resilient to polynomially-small fluctuations (Theorem 4.1), but that even for best-response this resilience breaks down for $\epsilon \geq \sqrt{2/n}$ (Theorem 4.2). We then consider two generalizations of these games: fair cost sharing in general directed graphs [2], and matroid games [1]. In both cases, we show that even exponentially small fluctuations can cause improved-response dynamics to move to high cost states (Lemma 4.2 and Theorems 4.5 and 4.8); however, best-response dynamics remains resilient to polynomially-small fluctuations (Theorems 4.3, 4.4, and 4.7). We also present results for job-scheduling and consensus games.

4.1 Fair Cost Sharing Games

In this section we consider fair cost sharing games (FCSG), a class of congestion games defined as follows. We are given a graph

$G = (V, E)$, which can be directed or undirected, where each edge $e \in E$ has a nonnegative cost $w_e \geq 0$. There is a set $N = \{1, \dots, n\}$ of n players, where player i is associated with a source s_i and a sink t_i . The strategy set of player i is the set \mathcal{S}_i of $s_i - t_i$ paths. In an outcome of the game, each player i chooses a single path $P_i \in \mathcal{S}_i$. Given an outcome $S = (P_1, \dots, P_n)$, let $n_e(S)$ be the number of agents whose path contains edge e . In the fair cost sharing game the cost to agent i is $\text{cost}_i(S) = \sum_{e \in P_i} \frac{w_e}{n_e(S)}$ and the goal of each agent is to connect its terminals with minimum total cost. The social cost is defined to be $\sum_{e \in \cup_i P_i} w_e$.

It is well known that the price of anarchy in these games is $\Theta(n)$ while the price of stability is $H(n)$ [2], where $H(n) = \sum_{i=1}^n 1/i = \Theta(\log n)$. A well known characterization of the potential function [17] of these games [2] is the following.

LEMMA 4.1. *In fair cost sharing, for any joint action $S \in \mathcal{S}$, we have: $\text{cost}(S) \leq \Phi(S) \leq H(n) \cdot \text{cost}(S)$.*

4.1.1 Set Covering Games

Set-cover games (SCG) were considered in [6]. In a set-cover game, there are n players, and m subsets over the players. The cost associated with set i is w_i . Each player j has to choose one of the sets i that contains him and gets to split the cost of the set with other players who choose the same set. Set-cover games are a special case of fair cost sharing games. We begin with an upper bound for improved-response dynamics.

THEOREM 4.1. *In the set covering game, for any $\epsilon > 0$ we have, $PoU_{IR}(\epsilon, SCG) \leq (1 + \epsilon)^{2mn} \log n$. Therefore, for $\epsilon = O(\frac{1}{nm})$, we have $PoU_{IR}(\epsilon, SCG) = O(\log n)$.*

PROOF. Suppose the initial configuration S_0 has k_i players using set/edge i of cost w_i . Think of this as a stack of k_i chips, where we label each chip with the name of its initial set and its position in the stack. So the bottom chip for set i is labeled $(i, 1)$, then the one above it is labeled $(i, 2)$, and so on. We will give chip (i, j) a value of w_i/j . So, the sum of values of the chips equals the potential function of the initial configuration, which according to Lemma 4.1 is at most a factor of $\log n$ larger than the original cost, i.e, $\text{cost}(S_0) \cdot \log n$. Now, when a player moves from some set i_1 to some set i_2 , we move the top chip from stack i_1 to stack i_2 . The claim is that we maintain the invariant that if chip (i, j) is currently at some position j' on some stack i' , then it must be the case that $w_{i'}/j' \leq (w_i/j) \cdot (1 + \epsilon)^{2mn}$. This will imply what we want, because it means that we can pay for any new set that gets taken using the bottom chip on its stack. (We are using here the fact that a chip can only be in one stack.)

The argument for the invariant is that there are at most $m \cdot n$ different positions a given chip can be in (m stacks, n positions per stack) so if you consider the path a chip takes from its initial location to its current location, this path has length at most $m \cdot n$ (you can remove loops in this configuration space). Since players follow an improved response dynamics, each step in this move causes the ratio of cost of the current stack to position in the stack to increase by at most a factor $(1 + \epsilon)^2$. So, overall, the total increase is at most a factor of $(1 + \epsilon)^{2mn}$. So, $\text{cost}(S_t) \leq (1 + \epsilon)^{2mn} \text{cost}(S_0) \log n$ for all t , as desired. \square

We now give a lower bound, showing that for $\epsilon \gg \log(n)/\sqrt{n}$, the price of uncertainty can get large even for best-response dynamics.

THEOREM 4.2. *In the set covering game we have $PoU_{BR}(\epsilon, SCG) = \Omega(\frac{\epsilon\sqrt{n}}{\log(1/\epsilon)})$.*

The proof builds on a construction in Section 6 giving a lower bound in the presence of a single Byzantine player, but adding in an

extra “gadget” of $\Theta(\sqrt{n}/\epsilon+1/\epsilon^2)$ players that allows the adversary to simulate the effect of a Byzantine player even via quite small fluctuations. The construction in this lower bound is fairly intricate so we defer the proof to Appendix A.

4.1.2 Fair Cost Sharing Games in General Graphs

We now consider fair cost sharing games in general directed graphs. We show here that so long as the number of players n_i of each type (i.e., associated to each (s_i, t_i) pair) is large, then the game is stable even for large values of ϵ under best-response dynamics. Specifically, so long as $n_i = \Omega(m)$ for all i , we have a constant price of uncertainty. On the other hand, for improved-response dynamics, then costs can grow exponentially even for exponentially small values of ϵ , even for the symmetric (single source, single sink) case.

THEOREM 4.3. *For fair cost sharing games, we have, $PoU_{BR}(\epsilon, FCSSG) \leq \left(1 + \frac{(1+\epsilon)^2}{n_{min}}\right)^m$, where $n_{min} = \min_i n_i$. This implies that for $n_{min} = \Omega(m)$, we have $PoU_{BR}(\epsilon, FCSSG) = O(1)$.*

PROOF. Call an edge “marked” if it is ever used throughout the best-response process, including those used in the initial state S_0 , and let c^* be the total cost of all marked edges. So, c^* is an upper bound on the cost of the current state. Any time a best-response path for some (s_i, t_i) pair uses an unmarked edge, the total cost of the unmarked edges used is at most $(c^*/n_i) \cdot (1 + \epsilon)^2$, because $(c^*/n_i)(1 + \epsilon)$ is an upper bound on the perturbed average cost of players of type (s_i, t_i) and therefore is an upper-bound on the perturbed cost of the best-response path for any such player. This in turn is within a $(1 + \epsilon)$ factor of the actual cost of this path. Thus, c^* increases by at most a multiplicative $(1 + (1 + \epsilon)^2/n_i)$ factor. We can mark new edges at most m times, so the final cost is at most $\text{cost}(S_0)(1 + (1 + \epsilon)^2/n_{min})^m$, where $n_{min} = \min_i n_i$. This implies that as long as $n_{min} = \Omega(m)$ we have $\text{cost}(S_t) = O(\text{cost}(S_0))$, for all t , as desired. \square

For symmetric fair cost sharing games (SCFCSSG) we can get a low price of uncertainty even when the number of players is much smaller than the number of edges, i.e., $n \ll m$.

THEOREM 4.4. *For symmetric fair cost sharing games, where the edge costs are in the range $[w_{min}, w_{max}]$, we have $PoU_{BR}(\epsilon, SFCSG) = O(\log n)$, for $\epsilon < \frac{w_{min}}{4w_{max}} \frac{1}{mn(n-1) \log n}$.*

PROOF. We start with some notation. We say that at time t , the best-response player j_t moves from path P_t to path P'_t , creating state S_t . We will say that a time t is “interesting” if $P_{t+1} \neq P'_t$: that is, if the next player moves from a path different from the one the current player moved to. Let us index the interesting times as t_1, t_2, \dots . The argument now proceeds in two steps: we first show an upper bound on the number of interesting time steps of $U = O(\frac{1}{\alpha} \log n)$ for $\alpha = \frac{w_{min}}{8w_{max}} \frac{1}{n(n-1)m}$. We then prove that the potential of the final state S_T satisfies $\Phi(S_T) \leq (1 + \epsilon)^{2U} \Phi(S_0)$. Using the fact that $\epsilon < \frac{w_{min}}{4w_{max}} \frac{1}{mn(n-1) \log n}$ and the $O(\log n)$ gap between potential and cost in these games, we get the desired result.

Let R_k denote the true (unperturbed) cost of the path P'_k at time t_k ; that is, $R_k = \text{cost}_{j_{t_k}}(S_{t_k})$. We now claim that

$$R_k \leq \left(R_{k-1} - \frac{w_{min}}{n(n-1)}\right)(1 + \epsilon)^2. \quad (4.1)$$

Specifically, note that because $P'_{t_{k-1}}$ and $P_{t_{k-1}+1}$ differ in at least one edge of cost at least w_{min} , and because $P'_{t_{k-1}+1} = P_{t_{k-1}+2}, \dots, P'_{t_k-1} = P_{t_k}$, any of the players $j_{t_{k-1}+1}, \dots, j_{t_k}$ could have

chosen to move to path $P'_{t_{k-1}}$ for a true (unperturbed) cost at most $R_{k-1} - \frac{w_{min}}{n(n-1)}$. In particular, $\frac{w_{min}}{n(n-1)}$ is a lower bound on the savings produced by having one more player on the edges in which $P'_{t_{k-1}}$ and $P_{t_{k-1}+1}$ differ (which implies the desired statement for $t_k = t_{k-1} + 1$) and each player $j_{t_{k-1}+2}, \dots, j_{t_k}$ could have moved to path $P_{t_{k-1}+1}$ reverting the system to state $S_{t_{k-1}}$ (which extends the statement to the case $t_k > t_{k-1} + 1$). Therefore, since player j_{t_k} is performing best response to the perturbed costs, the true cost R_k of P'_k is at most a factor $(1 + \epsilon)^2$ larger than $R_{k-1} - \frac{w_{min}}{n(n-1)}$.

For our given values of α and ϵ , and using the fact that $R_{k-1} \leq w_{max} \cdot m$, inequality (4.1) implies that $R_k \leq R_{k-1}(1 + \alpha)$. Since $R_1 \leq \mathbf{OPT}(1 + \epsilon)^2$ and by definition of \mathbf{OPT} it must be that $R_t \geq \mathbf{OPT}/n$, we get that the number of interesting time steps is at most $U = O(\frac{1}{\alpha} \log n)$.

We now bound the potential in terms of the number of interesting time steps. Specifically, note that player j_{t_k} could have moved to path $P_{t_{k-1}+1}$, which would revert the system to state $S_{t_{k-1}}$ because $P'_{t_{k-1}+1} = P_{t_{k-1}+2}, \dots, P'_{t_k-1} = P_{t_k}$. Because player j_{t_k} chose path P'_k instead, which was best-response to the perturbed costs, this means $\Phi(S_{t_k}) \leq \Phi(S_{t_{k-1}})(1 + \epsilon)^2$. Therefore, the final state S_T satisfies $\Phi(S_T) \leq \Phi(S_0)(1 + \epsilon)^{2U}$, completing the argument. \square

Improved response The above results give upper bounds for best response in fair cost sharing games. In contrast, we now show that for improved-response dynamics, the price of uncertainty is exponentially large even for exponentially-small values of ϵ , even for symmetric fair cost sharing games.

LEMMA 4.2. *For symmetric fair cost sharing, for a single player (i.e., $n = 1$), the price of uncertainty for improved-response dynamics satisfies*

$$PoU_{IR}(\epsilon, SFCSG) \geq 1 + 2(2^{m/2} - 1)\epsilon/m.$$

PROOF. The graph consists of a line of parallel edges arranged as follows. We have two parallel edges from $s = v_0$ to vertex v_1 of cost 1 and $1 + \epsilon$ respectively, then two parallel edges from v_1 to vertex v_2 of costs 1 and $1 + 2\epsilon$, then two parallel edges from v_2 to v_3 of costs 1 and $1 + 4\epsilon$, and in general from v_i to v_{i+1} of costs 1 and $1 + 2^i\epsilon$. Finally we let sink $t = v_{m/2}$ so we have a total of m edges. The player begins on the cheapest path, of cost $m/2$.

We can describe a path from s to t by a binary number $b = b_{m/2-1} \dots b_2 b_1 b_0$, where bit $b_i = 0$ if the path uses the edge of cost 1 from v_i to v_{i+1} and $b_i = 1$ if the path instead uses the edge of cost $1 + 2^i\epsilon$. Thus, path b has cost exactly $m/2 + b\epsilon$, and the player begins at path 0.

We now claim that using a series of perturbations and improved-response moves, one can cause the player to repeatedly increment, moving from path b to path $b + 1$ until the player finally reaches path $2^{m/2} - 1$, achieving the desired bound. Specifically, since the difference in true cost between path $b + 1$ and path b is exactly ϵ , it is sufficient to choose some arbitrary edge in path b that is not in path $b + 1$ and increase its cost by a multiplicative factor $1 + \epsilon$ to cause $b + 1$ to be an improvement over b (and we can similarly decrease the cost of an edge in $b + 1$ that is not in b to make it a strict improvement). \square

The generalization for multiple players is straightforward.

THEOREM 4.5. *For symmetric for cost sharing, for any number of players n , the price of uncertainty for improved-response dynamics satisfies $PoU_{IR}(\epsilon, SFCSG) \geq 1 + 2(2^{m/2} - 1)\epsilon/m$.*

We can use Lemma 4.2 to imply a bound also for routing games [16] with linear (or even constant) latency functions, since for the case of a single player these games are identical.

THEOREM 4.6. *For routing with linear latency functions, the price of uncertainty for improved-response dynamics satisfies $PoU_{IR}(\epsilon, ROUTING) \geq 1 + 2(2^{m/2} - 1)\epsilon/m$.*

4.2 Matroid Games

We now analyze matroid congestion games, a broad class of games considered in [1]. Before we give a formal definition of such games, we briefly introduce a few standard facts about matroids; for a detailed discussion, we refer the reader to [20].

DEFINITION 4.1. *A tuple $M := (R, \mathcal{I})$ is a matroid if R is a finite set of resources and \mathcal{I} is a nonempty family of subsets of R such that if $I \in \mathcal{I}$ and $J \subseteq I$, then $J \in \mathcal{I}$, and if $I, J \in \mathcal{I}$ and $|J| < |I|$, then there exists an $i \in \mathcal{I}$ such that $J \cup \{i\} \in \mathcal{I}$.*

Let $M := (R, \mathcal{I})$ be a matroid. Let $I \subset R$; if $I \in \mathcal{I}$ then we call I independent, otherwise we call it dependent. It is well known that all maximal independent sets of \mathcal{I} have the same size, which is denoted by the rank $\text{rk}(M)$ of the matroid. A maximal independent set of M is called a basis of M . It is well known that such a basis can be found by a greedy algorithm. In the following we state two additional useful properties of the matroids. We denote by \mathcal{B} the set of bases of a matroid, and assume that $B_1, B_2 \in \mathcal{B}$.

LEMMA 4.3. *Let $r_2 \in B_2 \setminus B_1$, then there exists $r_1 \in B_1 \setminus B_2$ such that $B_1 \cup \{r_2\} \setminus \{r_1\} \in \mathcal{B}$.*

We denote by $G(B_1 \triangle B_2)$ the bipartite graph (V, E) with $V = (B_1 \setminus B_2) \cup (B_2 \setminus B_1)$ and $E = \{\{r_1, r_2\} | r_1 \in B_1 \setminus B_2, r_2 \in B_2 \setminus B_1, B_1 \cup \{r_2\} \setminus \{r_1\} \in \mathcal{B}\}$. Then it is known that [20].

LEMMA 4.4. *There exists a perfect matching for $G(B_1 \triangle B_2)$.*

We are now ready to define matroid congestion games. A congestion game is a matroid congestion game if for every player $i \in N$ we have that $M_i := (R, \mathcal{I}_i)$ with $\mathcal{I}_i = \{I \subseteq S | S \in \mathcal{S}_i\}$ is a matroid and \mathcal{S}_i is the set of bases of M_i . We denote by $\text{rk}(M) = \max_{i \in N} \text{rk}(M_i)$ the rank of the matroid congestion game M . For example, set-cover games are matroid games of rank 1 and market-sharing games with uniform costs are matroid games [11] (though even symmetric fair cost sharing need not be a matroid game). We now show that for best-response dynamics, matroid games have similar resilience to fluctuations as set-cover games; however, for improved response we give an exponential lower bound.

THEOREM 4.7. *In a matroid game, $PoU_{BR}(\epsilon, \text{Matroid}) \leq (1 + \epsilon)^{2mn} \text{GAP}(\text{Matroid}, n)$. This implies that for $\epsilon = O(1/(n \cdot m))$, we have $PoU_{BR}(\epsilon, \text{Matroid}) = O(\text{GAP}(\text{Matroid}, n))$.*

PROOF. The proof proceeds as in Theorem 4.1. However, we initially have $\sum_{i \in N} \text{rk}(M_i) \leq n \cdot \text{rk}(M)$ chips and the $\text{cost}(S_0)$ is within a $\text{GAP}(\text{Matroid}, n)$ factor from the sum of the values of the chips.

Let S be a state of the matroid congestion game M and let s_i^* be a best response of player i to S according to the perturbed cost \tilde{d}_r . Consider the bipartite graph $G(s_i^* \triangle s_i)$ which contains a perfect matching P_M according to Lemma 4.4. Let $S^* = S \oplus s_i^*$ and observe that for every edge $(r, r^*) \in P_M$ with $r^* \in s_i^* \setminus s_i$ and $r \in s_i \setminus s_i^*$, $\tilde{d}_{r^*}(n_{r^*}(S^*)) \leq \tilde{d}_r(n_r(S^*) + 1) \leq \tilde{d}_r(n_r(S))$ since otherwise s_i^* is not a best response of player i with respect to \tilde{d}_r .

When a player does a best response we now move $\text{rk}(G)$ chips (corresponding to at most $\text{rk}(G)$ resources), and each movement sets up an inequality of the type $d_{i'}(j') \leq d_i(j) \cdot (1 + \epsilon)^2$. The claim is that we maintain the invariant that if chip (i, j) is currently at some position j' on some stack i' , then it must be the case that

$d_{i'}(j') \leq d_i(j) \cdot (1 + \epsilon)^{2mn}$. The argument is the same as in Theorem 4.1: there are at most $m \cdot n$ different positions a given chip can be in (m stacks, n positions per stack) so if you look at the path a chip takes from its initial location to its current location, this path has length at most $m \cdot n$ (you can remove loops in this configuration space).⁴ So, for all t we have $\Phi(S_t) \leq (1 + \epsilon)^{2mn} \Phi(S_0)$, which implies $\text{cost}(S_t) \leq (1 + \epsilon)^{2mn} \text{cost}(S_0) \cdot \text{GAP}(\text{Matroid}, n)$, which completes the proof. \square

Note: As opposed to the set-covering result (Theorem 4.1), this result holds for *best* response dynamics only. We can in fact show that improved response is not sufficient in these games, even if ϵ is exponentially small in the rank of the matroid. In particular, even though symmetric cost-sharing is not a matroid game, the proof of Lemma 4.2 applies equally well to improved-response dynamics if we replace the graph structure with a uniform matroid having m resources and with rank $r = m/2$. We therefore have the following:

THEOREM 4.8. *There exists a matroid game of rank $r = m/2$ such that the price of uncertainty for improved-response dynamics satisfies $PoU_{IR}(\epsilon, \text{Matroid}) \geq 1 + 2(2^{m/2} - 1)\epsilon/m$.*

4.3 Consensus Games

Consensus games [7] are played by users viewed as vertices in a connected, undirected simple graph $G = (N, E)$ with n vertices, where $N = \{1, \dots, n\}$. Each player i has two actions r or b , i.e., $\mathcal{S}_i = \{r, b\}$. A player has cost 1 for each incident edge on which he disagrees with his neighbor. $\text{cost}_i(S) = \sum_{(i,j) \in E} \mathbf{I}_{(s_i \neq s_j)}$. The overall social cost is the sum of the costs of all the users, plus 1, i.e., $\text{cost}(S) = 1 + \sum_{i \in N} \text{cost}_i(S)$. It is straightforward to show that these games are congestion games [14] and that the potential function can be rewritten as $\Phi(S) = (\text{cost}(S) - 1)/2$. The two social optimum solutions in a consensus game are “all blue” and “all red”, both of which are also a Nash equilibrium (so the Price of Stability is 1). On the other hand, there are Nash equilibria with cost $\Omega(n^2)$.

The above describes unweighted consensus games; in weighted-consensus, the edges have non-negative weights and the cost to a player is the total weight of edges on which it disagrees with its neighbors. In our model, we can show the following (for proofs, see the full version [5]):

THEOREM 4.9. *For any unweighted consensus game (UCG), for any ϵ , we have $PoU_{BR}(\epsilon, \text{UCG}) \geq \frac{(n-1)\epsilon+1}{1+\epsilon}$. For $\epsilon > \sqrt{2} - 1$ we can show $PoU_{BR}(\epsilon, \text{UCG}) = \Omega(n^2)$.*

For a weighted consensus game we can show an exponential lower bound.

THEOREM 4.10. *For any weighted consensus game (WCG), for any ϵ , we have $PoU_{BR}(\epsilon, \text{WCG}) \geq (1 + \epsilon)^n$.*

4.4 Job Scheduling

In this section we consider *job scheduling on unrelated machines (JSUM)* (see [16]) defined by (N, M, c) as follows. The set N is a set of n jobs, and M is the set of m machines. Each player is associated with a job, so have n players. Every job can impose a load on one of the machines, so for every player j its set of feasible actions is to assign job j to some machine $i \in M$. Each job $j \in N$ has associated a cost $c_{i,j}$ for running on machine $i \in M$. Given an assignment of jobs to machines, the load of machine i is the sum of the costs of the jobs that are assigned to that machine, i.e.,

⁴Note that the dynamics can last for a long time, it’s just that we can shortcut them in the argument.

$L_i(S) = \sum_{j \in B_i(S)} c_{i,j}$ where $B_i(S)$ is the set of jobs assigned to machine i , i.e., $B_i(S) = \{j : S_j = i\}$. The cost of a player j is the load on the machine that player j selected, i.e., $\text{cost}_j(S) = L_{s_j}(S)$. For the social cost we use the *makespan*, which is the load on the most loaded machine, i.e., $\text{cost}(S) = \max_i L_i(S)$. The price of stability in these games is 1, since there is always a pure Nash equilibrium which is also socially optimal [10].

The Price of Uncertainty in these games can be exponentially large, even for two machines, when ϵ is large compared to $1/n$. (For proofs, see the full version [5]):

THEOREM 4.11. *For job scheduling on unrelated machines, for $M = 2$ machines and any $\epsilon > 2/n$, we have $\text{PoUBR}(\epsilon, JSUM) \geq (1 + \epsilon)^{n/2} \left[1 - \frac{2}{\epsilon n}\right] + \frac{2}{\epsilon n} = \Omega((1 + \epsilon)^{n/2})$.*

For job scheduling on identical machines (JSIM) we have a simple upper bound, even for large perturbations.

THEOREM 4.12. *For job scheduling on identical machines, $\text{PoUIR}(JSIM, \epsilon) \leq 2(1 + \epsilon)$*

5. RANDOM ORDER

We now consider the effect of perturbations on *random* order best-response dynamics (for improved-response dynamics, random and adversarial order are equivalent since the adversary can simply choose not to cause a player to move). Our main result is that for the broad class of β -nice games introduced by [3], which for constant β includes congestion games with linear (or constant-degree polynomial) latency functions, market-sharing games, and many others, these dynamics are resilient to fluctuations even for constant $\epsilon > 0$. On the other hand, we give lower bounds showing that job-scheduling and consensus games can still behave poorly.

5.1 β -nice games

Let us consider an exact potential game. Let S be a profile of the players and let S^i denote the configuration produced by a best-response move by player i according to cost_i . For each player i define $\Delta_i(S) = \text{cost}_i(S) - \text{cost}_i(S^i)$ and $\Delta(S) = \sum_i \Delta_i(S)$.

DEFINITION 5.1. *An exact potential game with a potential function Φ is β -nice iff for any state S we have $\text{cost}(S) \leq \beta \text{OPT} + 2\Delta(S)$.*

As shown in [3] number of important games are β -nice, for β equal to the price of anarchy of the game.

Here we show that β -nice games additionally have the property that the expected price of uncertainty in the random order model is only $O(\beta \cdot \text{GAP})$ even for constant $\epsilon > 0$. We start by showing that if the true (unperturbed) cost of the current configuration S is greater than $2\beta \cdot \text{OPT}$, then no matter how the adversary adjusts the costs, the expected drop in potential is at least $\text{cost}(S)(1/4 - 4\epsilon)/n$. For $\epsilon < 1/16$, this is $\Omega(\text{cost}(S)/n)$. That is, the adversary may make the cost exceed $2\beta \text{cost}(S_0)$ but only temporarily.⁵

LEMMA 5.1. *For $\epsilon < 1/32$, if $\text{cost}(S_t) \geq 2\beta \text{OPT}$ then $\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] \leq -\text{cost}(S_t)/(8n)$.*

PROOF. As above, let S^i denote the configuration produced by a best-response move by player i according to cost_i , and let \tilde{S}^i denote the configuration produced by a best-response move by player i according to the perturbed cost function cost_i^t . So, $\text{cost}_i(S^i) \leq \text{cost}_i(\tilde{S}^i)$ and $\text{cost}_i^t(\tilde{S}^i) \leq \text{cost}_i^t(S^i)$. Recall that $\Delta_i(S) = \text{cost}_i(S) - \text{cost}_i(S^i)$ and $\Delta(S) = \sum_i \Delta_i(S)$. We will also need the following two quantities:

⁵This implies that with high probability the cost will drop to below $2\beta \text{OPT}$ within a polynomial number of steps.

1. $\tilde{\Delta}_i(S) = \text{cost}_i^t(S) - \text{cost}_i^t(\tilde{S}^i)$ is the improvement in perturbed cost of player i due to a best-response by player i in the perturbed game, with $\tilde{\Delta}(S) = \sum_i \tilde{\Delta}_i(S)$, and
2. $\hat{\Delta}_i(S) = \text{cost}_i(S) - \text{cost}_i(S^i)$ is the improvement in unperturbed cost of player i due to a best-response by player i in the perturbed game, with $\hat{\Delta}(S) = \sum_i \hat{\Delta}_i(S)$.

Now, suppose $\text{cost}(S) > 2\beta \text{OPT}$. Then by definition of β -nice we have $\Delta(S) > \text{cost}(S)/4$. Now we want to use this to show that $\hat{\Delta}(S)$ must be large as well. Specifically, for each i , since the improvement in perturbed cost of the best-response to the perturbed game is at least the improvement in perturbed costs of the best response to the unperturbed game, we have:

$$\begin{aligned} \tilde{\Delta}_i(S) &\geq \text{cost}_i^t(S) - \text{cost}_i^t(\tilde{S}^i) \\ &\geq (1 - \epsilon)\text{cost}_i(S) - (1 + \epsilon)\text{cost}_i(S^i) \\ &\geq \Delta_i(S) - 2\epsilon \text{cost}_i(S). \end{aligned} \quad (5.2)$$

Now, summing over all i we have:

$$\tilde{\Delta}(S) \geq \Delta(S) - 2\epsilon \text{cost}(S) \geq \text{cost}(S)(1/4 - 2\epsilon).$$

This means that when a random player moves to his “best perturbed response”, we have

$$E_i[\tilde{\Delta}_i(S)] \geq \text{cost}(S)[1/4 - 2\epsilon]/n.$$

Now, by a similar argument to that equation (5.2) we have:

$$\begin{aligned} \hat{\Delta}_i(S) &= \text{cost}_i(S) - \text{cost}_i(S^i) \\ &\geq \text{cost}_i^t(S) - \epsilon \text{cost}_i(S) - \text{cost}_i^t(\tilde{S}^i) - \epsilon \text{cost}_i(\tilde{S}^i) \\ &= \tilde{\Delta}_i(S) - 2\epsilon \text{cost}_i(S) + \epsilon \hat{\Delta}_i(S). \end{aligned}$$

So, $\hat{\Delta}_i(S) \geq (\tilde{\Delta}_i(S) - 2\epsilon \text{cost}_i(S))/(1 - \epsilon)$. Putting this together with the above and using the fact that $E_i[\text{cost}_i(S)] = \text{cost}(S)/n$ gives us

$$E_i[\hat{\Delta}_i(S)] \geq \text{cost}(S)[1/4 - 4\epsilon]/n$$

which is the expected drop in the potential Φ for the unperturbed game caused by a random best-response move in the perturbed game. If $\epsilon < 1/32$, we then get the desired result. \square

So, Lemma 5.1 shows if the true (unperturbed) cost of current configuration S is greater than $2\beta \cdot \text{OPT}$, then no matter how the adversary adjusts the costs, the expected drop in potential is at least $\text{cost}(S)(1/4 - 4\epsilon)/n$. A Chernoff bound argument can then be used to say that with high probability the sum of drops in potential will be close to their expectation. Note that we do not show that once cost is low it will necessarily stay there forever – just that if the adversary is ever able to make the cost go above $2\beta \text{OPT}$ then with high probability it will have to drop back below it in a small number of steps.

In the following we show a bound on the expectation that holds for all time steps. To do so, we use the following additional lemma:

LEMMA 5.2. *For any value of $\text{cost}(S_t)$, $\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] \leq 2\epsilon \text{cost}(S)/(n(1 - \epsilon))$.*

PROOF. This just follows from the statement that $\hat{\Delta}_i(S) \geq (\tilde{\Delta}_i(S) - 2\epsilon \text{cost}_i(S))/(1 - \epsilon)$, and using the fact that $\tilde{\Delta}_i(S)$ is always non-negative. \square

Assume $\epsilon < 1/32$. We can now use these lemmas to prove that for β -nice games the expected price of uncertainty in the random order model is only $O(\beta \cdot \text{GAP})$ even for constant $\epsilon > 0$. Recall that we define $\text{GAP} = c_2/c_1$ where $c_1 \leq 1 \leq c_2$ are values such that for any state S we have $\Phi(S) \in [c_1 \text{cost}(S), c_2 \text{cost}(S)]$.

THEOREM 5.1. *For any $t > 0$, we have*

$$\mathbf{E}[\Phi(S_t)] \leq \max[5c_2\beta\mathbf{OPT}, c_2\mathbf{cost}(S_0)] \leq 5c_2\beta\mathbf{cost}(S_0).$$

Therefore, $\mathbf{E}[\mathbf{cost}(S_t)] \leq 5\beta\mathbf{cost}(S_0) \cdot \text{GAP}$.

PROOF. We will show that if $\mathbf{E}[\mathbf{cost}(S_t)] \geq 4\beta\mathbf{OPT}$ then $\mathbf{E}[\Phi(S_{t+1})] \leq \mathbf{E}[\Phi(S_t)]$. This will be sufficient because Lemma 5.2 implies that the expectation can never increase by too much. In particular, even if $\mathbf{E}[\mathbf{cost}(S_t)] \leq 4\beta\mathbf{OPT}$, by Lemma 5.2 we still have

$$\begin{aligned} \mathbf{E}[\Phi(S_{t+1})] - \mathbf{E}[\Phi(S_t)] &\leq 4\epsilon\beta\mathbf{OPT}/(n(1-\epsilon)) \\ &< \beta\mathbf{OPT} \leq c_2\beta\mathbf{OPT}. \end{aligned}$$

Specifically, suppose $\mathbf{E}[\mathbf{cost}(S_t)] \geq 4\beta\mathbf{OPT}$. Let p_t be the probability that $\mathbf{cost}(S_t) \geq 2\beta\mathbf{OPT}$. Therefore, we have:

$$\begin{aligned} \mathbf{E}[\mathbf{cost}(S_t)] &= p_t\mathbf{E}[\mathbf{cost}(S_t)|\mathbf{cost}(S_t) \geq 2\beta\mathbf{OPT}] \\ &\quad + (1-p_t)\mathbf{E}[\mathbf{cost}(S_t)|\mathbf{cost}(S_t) < 2\beta\mathbf{OPT}] \\ &\leq p_t\mathbf{E}[\mathbf{cost}(S_t)|\mathbf{cost}(S_t) \geq 2\beta\mathbf{OPT}] + 2\beta\mathbf{OPT}, \end{aligned}$$

so we have $\mathbf{E}[\mathbf{cost}(S_t)|\mathbf{cost}(S_t) \geq 2\beta\mathbf{OPT}] \geq 2\beta\mathbf{OPT}/p_t$. Now, using Lemmas 5.1 and 5.2, we can write:

$$\begin{aligned} \mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] &\leq \left(-\frac{p_t}{8n}\right)\mathbf{E}[\mathbf{cost}(S_t)|\mathbf{cost}(S_t) \geq 2\beta\mathbf{OPT}] + \\ &\quad \frac{2\epsilon(1-p_t)}{n(1-\epsilon)}\mathbf{E}[\mathbf{cost}(S_t)|\mathbf{cost}(S_t) < 2\beta\mathbf{OPT}] \\ &\leq 2\beta\mathbf{OPT}/(-8n) + 2\beta\mathbf{OPT}\frac{2\epsilon}{n(1-\epsilon)} \leq 0. \end{aligned}$$

Thus, as desired, if $\mathbf{E}[\mathbf{cost}(S_t)] \geq 4\beta\mathbf{OPT}$ then $\mathbf{E}[\Phi(S_{t+1})] \leq \mathbf{E}[\Phi(S_t)]$, proving the claim. \square

As shown in [3], a number of common games are β -nice for constant β , including congestion games with linear latency functions, both unweighted ($\beta = 2.5$) and weighted ($\beta \approx 2.618$), congestion games with polynomial latency functions of constant degree d ($\beta = d^{d(1-o(1))}$), and market-sharing games ($\beta = 2$).

Note: Interestingly, the guarantee in Lemma 5.1 breaks down in the adversarial-order setting: for example, for market-sharing games, which are β -nice for $\beta = 2$, we have price of uncertainty $\Omega(\log n)$ even for $\epsilon = 0$, as shown in Theorem 3.1.

5.2 Job Scheduling and Consensus games

For job scheduling we can show (see full version [5] for proofs):

THEOREM 5.2. *For $M = 2$ machines, for any $\epsilon > 2/n$, we have that $PoU_{BR}^R(\epsilon, JSUM) \geq \epsilon n/8$.*

We can similarly adapt our lower bounds for consensus games to the random order model. For a weighted consensus game we can show an exponential lower bound.

THEOREM 5.3. *For any weighted consensus game (WCG), for any ϵ , we have $PoU_{BR}^R(\epsilon, WCG) \geq (1+\epsilon)^{n/2-1}$.*

6. BYZANTINE PLAYERS

We now consider the case that, rather than perturbing weights, the adversary instead controls a certain number of Byzantine players who can move arbitrarily between best-response moves by the ordinary (non-Byzantine) players. Our main results in this model are an upper bound for β -nice games, showing that random-order

best-response dynamics is resilient to Byzantine players, and a lower bound for set-covering games, showing that in these games an adversary can increase the cost of the normal players by a factor of $\Omega(n)$ even with just *one* Byzantine player. We also give results for job-scheduling and consensus games as well.

6.1 β -nice games

Earlier, we showed that β -nice games are resilient to cost perturbations in the random order model. Here we show they are also resilient to the addition of Byzantine players. For this, we make two additional reasonable assumptions about the game and the number of Byzantine players:

Assumption 1 (monotonicity): We assume that adding new players into the game can only increase the cost incurred by any given player (e.g., as in linear congestion games).

Assumption 2 (low direct impact of Byzantine players): configuration S , the social cost of S with Byzantine players removed is at least $7/8$ of the social cost of S with Byzantine players included. In other words, the Byzantine players cannot change the cost of any given state by more than a small constant factor.

We will consider random best-response dynamics. Recall that in this model, Byzantine players may move arbitrarily between two moves of the normal (non-Byzantine) players. The key to the analysis is that we will track the cost and potential of the configuration minus the Byzantine players, viewing the Byzantine players as merely perturbations to the perceived costs of the normal players, causing them to act in an unusual way. We then will follow the main steps of the analysis of β -nice games in Section 5.1. However, note that now the Byzantine players can affect the perceived cost of any given normal player substantially, even though by Assumption 2 they cannot change the aggregate cost by too much.

Specifically, let players $1, \dots, n$ be the normal players, and we will index the Byzantine players as $n+1, \dots, n+k$. Given a configuration S at time t , define $\mathbf{cost}(S)$ to be the social cost of S with Byzantine players removed, and define $\mathbf{cost}^t(S)$ to be the social cost of S with Byzantine players included. Similarly, define $\mathbf{cost}_i(S)$ and $\mathbf{cost}_i^t(S)$ to be the cost incurred by player i with Byzantine players removed or included, respectively. So, $\mathbf{cost}(S) = \sum_{i=1}^n \mathbf{cost}_i(S)$ and $\mathbf{cost}^t(S) = \sum_{i=1}^{n+k} \mathbf{cost}_i^t(S)$. Also, by Assumptions 1 and 2 we have $\mathbf{cost}^t(S) \geq \mathbf{cost}(S) \geq \frac{7}{8}\mathbf{cost}^t(S)$.

Define the potential $\Phi(S)$ to be the standard potential function for configuration S but with Byzantine players removed, and let S_t denote the state at time t (counting each move of a non-Byzantine player as one time step). We now prove the following lemma.

LEMMA 6.1. *If $\mathbf{cost}^t(S_t) \geq 2\beta\mathbf{OPT}$ then we have $\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] \leq -\mathbf{cost}^t(S_t)/(8n)$.*

PROOF. Given configuration S , let S^i denote the configuration resulting from player i performing best-response to the perceived costs (i.e., with Byzantine players included). Let $\hat{\Delta} = \sum_{i=1}^n [\mathbf{cost}_i(S) - \mathbf{cost}_i(S^i)]$. In other words, $\hat{\Delta}/n$ is the expected drop in the potential Φ caused by a random non-Byzantine player performing best-response to the costs with Byzantine players included.

Let $\tilde{\Delta} = \sum_{i=1}^{n+k} [\mathbf{cost}_i^t(S) - \mathbf{cost}_i^t(S^i)] = \mathbf{cost}^t(S) - \sum_{i=1}^{n+k} \mathbf{cost}_i^t(S^i)$. This is a somewhat strange quantity since the Byzantine players are not actually performing best response. Nonetheless, by the definition of β -nice, if $\mathbf{cost}^t(S) \geq 2\beta\mathbf{OPT}$ then we have $\tilde{\Delta}(S) \geq \mathbf{cost}^t(S)/4$. This implies $\sum_{i=1}^{n+k} \mathbf{cost}_i^t(S^i) \leq \frac{3}{4}\mathbf{cost}^t(S)$, and therefore surely $\sum_{i=1}^n \mathbf{cost}_i^t(S^i) \leq \frac{3}{4}\mathbf{cost}^t(S)$

as well. Putting this together, we now have:

$$\begin{aligned} \hat{\Delta}(S) &= \text{cost}(S) - \sum_{i=1}^n \text{cost}_i(S^i) \geq \text{cost}(S) - \sum_{i=1}^n \text{cost}_i^t(S^i) \\ &\geq \text{cost}(S) - \frac{3}{4} \text{cost}^t(S) \geq \frac{7}{8} \text{cost}^t(S) - \frac{3}{4} \text{cost}^t(S) \\ &\geq \text{cost}^t(S)/8, \end{aligned}$$

where the first inequality follows by monotonicity and the second to last follows by Assumption 2. Since $\hat{\Delta}/n$ is the expected drop in Φ , this concludes the proof. \square

To analyze the expected costs, we now need an analog of Lemma 5.2, showing that even if $\text{cost}^t(S_t)$ is low, the expected value of the potential will not increase too quickly.

LEMMA 6.2. *For any value of $\text{cost}^t(S_t)$, we have $\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] \leq \text{cost}^t(S_t)/(8n)$.*

PROOF. Let $S = S_t$. Using the notation from the proof of Lemma 6.1 we have

$$\begin{aligned} \hat{\Delta}(S) &= \text{cost}(S) - \sum_{i=1}^n \text{cost}_i(S^i) \geq \frac{7}{8} \text{cost}^t(S) - \sum_{i=1}^n \text{cost}_i(S^i) \\ &\geq \frac{7}{8} \text{cost}^t(S) - \sum_{i=1}^{n+k} \text{cost}_i^t(S^i) = \tilde{\Delta}(S) - \frac{1}{8} \text{cost}^t(S). \end{aligned}$$

This is at least $-\frac{1}{8} \text{cost}^t(S)$ as desired. \square

Putting these together we can now show the following analog of Theorem 5.1, for β -nice games satisfying Assumptions 1 and 2.

THEOREM 6.1. *For any $t > 0$, $\mathbf{E}[\Phi(S_t)] \leq \max[6\beta \text{OPT}, \text{cost}(S_0)] \leq 6\beta \text{cost}(S_0)$.*

Finally, note that if there is a bounded value $\text{GAP} = \max_S [\text{cost}^t(S)/\Phi(S)]$ then the above result implies that for all $t > 0$, $\mathbf{E}[\text{cost}^t(S_t)] \leq \max[6\beta \text{OPT}, \text{cost}^t(S_0)] \cdot \text{GAP}$.

6.2 Set-covering games

We now consider set-cover games and give a construction showing that just one Byzantine player can cause best-response dynamics to move from an equilibrium of cost $O(\text{OPT})$ to an equilibrium of cost $\Omega(n \cdot \text{OPT})$. Note that this is the largest gap possible since the Price of Anarchy for this game is n .

THEOREM 6.2. *For set-cover games, a single Byzantine player can cause best-response dynamics to move from a Nash equilibrium of cost $O(n)$ to a Nash equilibrium of cost $\Omega(n^2)$.*

PROOF. Consider n players of type I, where each player i has two sets to choose from: a common set s^* of cost n , and a set s_i of cost $n - 1$. There are additionally $n - 2$ players of type II, such that player k of type II may either choose any of the sets s_i or else its own set f_k of cost n/k , for $k \in \{2, \dots, n - 1\}$. In addition, we have one Byzantine player who may choose any of the sets, for a total of $2n - 1$ players total. The initial state is all players of type I in set s^* and all players of type II in set s_1 , for a total cost $O(n)$.

The Byzantine player and type-II players will now slowly lure all type-I players into the sets s_i , increasing the cost of the system to $n(n - 1)$. First the Byzantine player moves to set s_1 causing player 1 of type I to move from s^* (whose cost to the player is 1) to s_1 (whose cost to the player is $(n - 1)/n$). The Byzantine player then sequentially moves to each set $f_{n-1}, f_{n-2}, \dots, f_2$, causing

the players of type II to move to their sets f_k in that order. Specifically, at the time player k of type II moves, the set s_1 has a cost to it of $(n - 1)/k$, whereas set f_k has cost (with the Byzantine player) of $n/(2k)$.

Now the Byzantine player moves to set s_2 , causing the players k of type II for $k = 2, 3, \dots, n - 2$ to move one after the other to s_2 as well. Specifically, at the time player k moves, set s_2 has cost $(n - 1)/k$ which is lower than the cost n/k of f_k . At the end of this step we have the same configuration of type-II players as in the initial state, except with s_2 rather than s_1 . The entire process then repeats for player 2 of type I, and so on, until each player i of type I is on its own set s_i . Finally, since s^* is now empty, none of the type-I players wish to move so we are at an equilibrium. \square

As pointed out in Section 3, Theorem 6.2 immediately implies that in the perturbation model we have $\text{PoUBR}(\sqrt{2} - 1, \text{SCG}) = \Omega(n)$. Note that in Theorem 4.2 we extend this construction to the more delicate case of small values of ϵ as low as $\sqrt{2/n}$.

6.3 Job scheduling and consensus games

We begin with a simple lower bound for job scheduling in the presence of a single Byzantine job.

THEOREM 6.3. *For two machines and one Byzantine job, the cost can increase from 2 to $\Omega(n)$, even in the random-order model.*

PROOF. Consider two machines and $2n + 1$ good jobs, n of type I with cost $(1/n, 1)$ and n of type II with cost $(1, 1/n)$, and one Byzantine job with cost $(1, 1)$. Initially, jobs of type I are on machine 1 and jobs of type 2 are on machine 2, this is the optimal assignment for the good jobs. The Byzantine job goes to machine 1, then a job of type I moves to machine 2, since this is its best response. Now, the Byzantine job moves to machine 2 (note that the Byzantine job increased its load). Then, a job of type II moves from 2 to 1. This way we increase the cost from 2 to $n + 1$.

As in Theorem 5.2, this construction extends immediately to the random order model, with just a constant factor loss in the ratio, by analyzing the Markov chain produced as a result of the above adversary strategy. Specifically, so long as the system has more jobs with cost $1/n$ on their current machine than jobs with cost 1 on their current machine, the system is more likely to transition in the forward direction (increasing the number of high-cost jobs) than in the reverse direction. Thus, in $O(n)$ steps, with high probability the system reaches a state of cost $\Omega(n)$. \square

On the other hand, for job scheduling on identical machines, unless the Byzantine players by themselves have substantial weight, they cannot cause the system to reach a high-cost state.

THEOREM 6.4. *For job scheduling on identical machines, the makespan is at most $2\text{OPT} + W_b$, where W_b is the sum of the weights of the Byzantine jobs.*

PROOF. Let W_b be the weight of the Byzantine players and W_g the weight of the good players. Each time a good player i moves it has a best response whose cost is at most $(W_g + W_b)/m + w_i$, where w_i is the good job cost. Note that $\text{OPT} \geq \max\{W/m, w_i\}$ and thus the result follows. \square

We end with a simple observation that for unweighted consensus, a single Byzantine player can cause cost to increase by a factor $\Omega(n)$.

THEOREM 6.5. *For the unweighted consensus game, a single Byzantine player can increase cost from 1 to $\Omega(n)$.*

PROOF. The network is simply a line network v_1, \dots, v_n . The Byzantine player is the player v_1 at one end of the line. Assume

we start with all players being R and then the Byzantine player switches to B . Player v_2 is indifferent between R and B so it switches to B , and then player v_1 switches back to R . Then, player v_3 switches to B , player v_2 switches to R and player v_1 switches to B . In phase k , we start with v_1, \dots, v_k alternating between R and B , such that v_k plays B . During phase k , first v_{k+1} switches to B , and then players v_k, \dots, v_1 switch their action. At the end we have all players alternating between R and B , at cost $\Omega(n)$. \square

7. OPEN QUESTIONS

In terms of specific open questions, it would be interesting to close some of the gaps that remain in the adversarial order model. For example, can one extend the upper bound of Theorem 4.4 to non-symmetric cost-sharing games or extend the lower bound of Theorem 4.2 to symmetric cost sharing games? In the Byzantine model, can one get better upper-bounds for set-cover and fair cost-sharing games if we assume random order dynamics? More generally, for all the classes of games studied, can one get better upper bounds in the random order model in the case where the perturbations are not completely adversarial, but instead chosen from some distribution of bounded magnitude?

Acknowledgements: We would like to thank Heiko Röglin for a number of helpful discussions about matroid congestion games.

8. REFERENCES

- [1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In *Proc. 47th FOCS*, 2006.
- [2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *FOCS*, 2004.
- [3] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *ACM Conference on Electronic Commerce (EC)*, 2008.
- [4] M. Babaioff, R. Kleinberg, and C.H. Papadimitriou. Congestion games with malicious players. In *ACM EC*, 2007.
- [5] M.-F. Balcan, A. Blum, and Y. Mansour. The price of uncertainty (full version). <http://www.cs.cmu.edu/~ninamf/papers/pou-full.pdf>.
- [6] N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative cost sharing games via subsidies. In *SAGT*, pages 337–349, 2008.
- [7] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *Proc. 23rd STACS*, 2006.
- [8] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. 13th ACM-SIAM SODA*, 2002.
- [9] B. de Finetti. La prevision: Ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937.
- [10] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3), 2007.
- [11] M.X. Goemans, L. Li, V.S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. *IEEE Journal on Selected Areas in Communication*, 24:1020–1033, 2006.
- [12] M. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6), 1998.
- [13] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proc. 16th STACS*, pages 404–413, 1999.
- [14] D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [15] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: byzantine players in a virus inoculation game. In *PODC*, 2006.
- [16] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge, 2007.
- [17] R.W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *Int. Journal Game Theory*, 2(1), 1973.
- [18] A. Roth. The price of malice in linear congestion games. In *WINE*, pages 118–125, 2008.

- [19] L. J. Savage. *The Foundations of Statistics*. John Wiley and Sons, 1954.
- [20] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2002.
- [21] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.
- [22] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS*, 2002.
- [23] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

APPENDIX

A. SET COVERING GAMES

Theorem 4.2 In the set covering game we have $PoU_{BR}(\epsilon, SCG) = \Omega(\frac{\epsilon\sqrt{n}}{\log(1/\epsilon)})$ for $\epsilon \geq 2/\sqrt{n}$.

PROOF. The construction⁶ builds on that in the proof of Theorem 6.2. Let $N = n/4$. Consider N players of Type I, indexed by pairs (i, j) for $1 \leq i, j \leq \sqrt{n}$. Type-I player (i, j) has three sets to choose from: a set s_i^* of cost N , a set $s_{active,j}$ of cost N , and a private set $s_{i,j}$ of cost N . Initially, all Type-I players begin on the sets s_i^* for a total cost of $N\sqrt{N}$. We also have $N + \sqrt{N} - 2$ players of Type II as follows. For $k = 3, \dots, N$, we have a Type-II player who may choose any of the sets $s_{active,j}$ (for $1 \leq j \leq \sqrt{N}$) or else its own set f_k . For $k = 2$ we have \sqrt{N} Type-II players indexed by pairs $(2, j)$ for $1 \leq j \leq \sqrt{N}$ who may choose only the set $s_{active,j}$ or their own set $f_{2,j}$. The sets f_k and $f_{k,j}$ have costs as follows: for $k \in [2, 1/\epsilon]$, the set has cost $\frac{1}{k}N/k$; for $k > 1/\epsilon$, the set has cost N/k . However, for $k \in [2, 1/\epsilon]$, we have $1/\epsilon - 1$ “helper” players in set f_k (or set $f_{k,j}$) whose alternative options will be described in more detail below. Thus, with helper-players included, cost of each set f_k (or $f_{k,j}$) to type-II player k (or (k, j)) is N/k . The Type-II players begin in their sets f_k (or $f_{k,j}$).

As in the proof of Theorem 6.2, the Type-II players will now slowly lure all Type-I players into the private sets $s_{i,j}$, increasing their overall cost from $N\sqrt{N}$ to N^2 . Specifically, for $i = 1, 2, \dots, \sqrt{N}$, the following occurs. First, for each $j = 1, 2, \dots, \sqrt{N}$ in sequence, Type-I player (i, j) is lured onto $s_{active,j}$ as follows. First, Type-II player $(2, j)$ is moved to set $s_{active,j}$ by having its helper-players temporarily raise the effective cost of $f_{2,j}$ from $N/2$ to N (using a process described in the paragraph below) so that the adversary can cause it to move to $s_{active,j}$ with an arbitrarily small additional perturbation. Next, Type-II players $k = 3, 4, \dots$ are made to follow along to $s_{active,j}$. In the case of $k = 3, \dots, 1/\epsilon$, this is done by having the helper-players again temporarily raise the effective cost of f_k from N/k to $2N/k$, making the player prefer $s_{active,j}$ to f_k . In the case of $k > 1/\epsilon$, this can be done without helper players, since the ratio of the cost of $s_{active,j}$ to the cost of f_k is $k/(k-1) \leq 1 + \epsilon$, so perturbations are sufficient. (Note that Type-II player $(2, j)$ would have preferred $s_{active,j'}$ for $j' < j$ to $s_{active,j}$ because that set already has a Type-I player on it, but that is not one of its allowed sets; Type-II player $k = 3$ is indifferent (so can be made to move as desired with arbitrarily small perturbations) and Type-II players $k > 3$ will strictly prefer $s_{active,j}$ to $s_{active,j'}$ for $j' < j$.) Now, player (i, j) of Type I moves from s_i^* (whose cost to the player is at least \sqrt{N} to $s_{active,j}$ (whose cost to the player is 1). Finally, the players of Type II, in order from $k = N$ down to 2, sequentially move back to their sets f_k (or $f_{k,j}$). In particular, at the time player k of type II moves, the set $s_{active,j}$ has a cost to it of N/k , which is equal to the cost of set f_k (so with arbitrarily small perturbations, the adversary

⁶We thank Florin Constantin and Steven Ehrlich for pointing out a bug in an earlier construction.

can easily cause these players to move). After the above process has been completed for all $j = 1, 2, \dots, \sqrt{N}$ (so that s_i^* is now empty for the current value of i), the Type-I players (i, j) are now indifferent between all sets to which they are eligible. So, they can each be made to move to their private sets $s_{i,j}$ via arbitrarily small perturbations. We then increment i and repeat the entire process above.

To finish the argument, we need to describe how the helper-players raise the effective cost of f_k (or $f_{k,j}$). This proceeds as follows. For each k , the j th helper-player has a private set of cost $\frac{N}{k(1-j\epsilon)}$. By perturbing costs, adversary can cause these players for $j = 1, 2, \dots, 1/(2\epsilon)$ to move in order to their private sets. Specifically, at the time the j th player is to move, the ratio of the cost of its private set to the cost of f_k is $\frac{N}{k(1-j\epsilon)} \cdot \frac{k\epsilon(1/\epsilon-j+1)}{N} = \frac{1-(j-1)\epsilon}{1-j\epsilon} \leq 1 + 2\epsilon < (1 + \epsilon)^2$. This then raises the cost of f_k as desired. Once player k of type II has moved off of set f_k , the helper players return back to f_k in the order $j = 1/(2\epsilon), \dots, 2, 1$ (they are now indifferent between the two sets, so the adversary can cause them to move via arbitrarily small perturbations) bringing f_k back to its initial state. This completes the construction.

The total number of players is upper bounded by $2N + \frac{1}{\epsilon}(\sqrt{N} + \frac{1}{\epsilon})$. For $\epsilon \geq 1/\sqrt{N}$ and $N = n/4$, this is at most n . The total cost of the initial state is $O(N^{3/2} + N^{3/2}/\epsilon + (N/\epsilon) \log(1/\epsilon))$ and the final state has cost $\Omega(N^2)$, giving the ratio desired. \square