The Price of Uncertainty*

Maria-Florina Balcan† Avrim Blum‡ Yishay Mansour§

October 25, 2012

Abstract

In this work 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. We show that in certain cases, even extremely small fluctuations can have the ability to 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-cover, job-scheduling) even a single Byzantine player can cause best-response dynamics to transition from low-cost states 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.

Overall, our work can be viewed as analyzing the inherent resilience or *safety* of games to different kinds of imperfections in player behavior, player information, or in modeling assumptions made.

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 Morgenstern [29], de Finetti [13] and Savage [24]. In essence, these results connect between agents' 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 [27].) In this work we explore how small fluctuations or uncertainties about agents' own utilities, or a small number of agents whose utilities are not well-modeled by the game description, 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

*A preliminary version of these results appeared in *Proceedings of the ACM Conference on Electronic Commerce*, 2009.

†School of Computer Science, Georgia Institute of Technology. This work was conducted in part while the author was a Postdoctoral Fellow at Microsoft Research, New England. This work was supported in part by NSF grants CCF-0953192 and CCF-1101283, by ONR grant N00014-09-1-0751, and AFOSR grant FA9550-09-1-0538.

‡School of Computer Science, Carnegie Mellon University. Supported in part by the National Science Foundation under grants CCF-0830540 and CCF-1101215.

§The Blavatnik School of Computer Science, Tel Aviv University. Supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant from the Israel Science Foundation, by a grant from United States-Israel Binational Science Foundation (BSF) and by The Israeli Centers of Research Excellence (I-CORE) program, (Center No. 4/11). This publication reflects the authors' views only.

the load on each machine. An agent has to approximate this speed from its observations, for example by measuring output quantity per unit 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 [16].

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 delays that are some given function of the amount of traffic on them, but in reality delays may also depend on external environmental factors or transient conditions which have been abstracted out of the model. Therefore delays would not be exactly identical at two different times even if the agents chose the exact same routes. We can again view this similarly: at each time t, edge j has cost $c_j^t \in [c_j/(1+\epsilon), c_j(1+\epsilon)]$ where c_j is the "base" cost of that edge (not including the external factors that have been abstracted out of the model), 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, if the machines also have small maintenance tasks they are performing in the background that come and go and are not part of the game description.

The question we are interested in is: can these fluctuations—in perceived or actual costs—cause natural dynamics that otherwise would be well-behaved to in effect spin out of control and move from low-cost states to states of much worse social welfare? What values of ϵ can different well-studied games tolerate? We focus on the most basic myopic dynamics, in particular best-response and improved-response (also known as *better-response*) dynamics, and on potential games, where such dynamics and these questions are especially natural. In particular, potential games [21] have a potential function that ordinarily would decrease with each best-response or improved-response move, and therefore the maximum gap between potential and cost in such a game would ordinarily bound how much worse any state could get under these dynamics.[1] With fluctuations in real or perceived costs, however, this no longer holds, and we define a new measure we call the *Price of Uncertainty (PoU)*, that quantifies 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 perturbed cost function at time t. The PoU of a given game is then defined as the maximum, over all possible initial states S_0 and all possible states S_T that can be reached from S_0 via such a dynamics, of the ratio of the social cost of S_T to the social cost of S_0.[2] A small PoU implies that the system cannot significantly deteriorate due to such perturbations in costs, and provides a certain degree of robustness to the system. A large PoU means that, in contrast, deterioration due to small perturbations could be severe. Note that all the games we consider have only a small gap between potential and cost; so, as mentioned above, without any fluctuations these dynamics would never cause social welfare to deteriorate substantially. Furthermore, this small gap implies the games 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.

We also explore a different kind of robustness, which is a robustness to adversarial (Byzantine) players who may move arbitrarily. This can be viewed as a *fault-tolerance* of the system to a few misbehaved agents or players whose true motivations are not correctly modeled by the given game. For example, one player in a resource allocation game might, unknown to the system, be representing a series of different clients

[1] An (ordinal) potential game has a potential function Φ defined over states, such that if any player i has a move that reduces its own cost, then such a move reduces the potential as well. Note that if furthermore $\texttt{cost}(S) \leq \Phi(S) \leq c \cdot \texttt{cost}(S)$ for all states S (where "cost" represents social cost), then starting from some state S_0, any state S_T reachable by these dynamics must satisfy $\texttt{cost}(S_T) \leq \Phi(S_T) \leq \Phi(S_0) \leq c \cdot \texttt{cost}(S_0)$; i.e., cost can increase by at most a factor of c. See Section 2 for further discussion.

[2] The initial state S_0 could be an equilibrium, in which case the PoU can be viewed as studying the stability of the equilibrium. Alternatively, S_0 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 would start from a more arbitrary state.

over time who each have different needs, and therefore be constantly changing the resources it requests in an unpredictable manner. For our lower bounds, we concentrate on the case of a single Byzantine player, and measure to what degree such a player can cause the system's social welfare to deteriorate. For our positive results, we allow any number of Byzantine players, so long as they cannot themselves dominate the overall social cost of any *given* state.

Our Results: We analyze a number of well-studied potential games from this perspective, including cost-sharing [2], matroid games [1], job-scheduling [17, 12], and the class of β-nice games [4] (these games are formally defined in the sections where they are analyzed). We prove both upper and lower bounds on how resilient these games 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, a polynomially-small upper-bound on perturbations is sufficient to ensure good behavior, and finally some games 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 move from an equilibrium of cost $O(\mathbf{OPT})$ to one of cost $\Omega(n \cdot \mathbf{OPT})$, where n is the number of players, and $\mathbf{OPT}$ is the cost of the socially-optimal state. That is, the PoU is $\Omega(n)$, matching the Price of Anarchy of the game. For set-cover games—a natural special case of fair cost-sharing studied by [9] 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-cover games 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 lower bound of $\Omega(\epsilon\sqrt{n}/\log(1/\epsilon))$ on the PoU of best-response dynamics for set-cover games, showing that $\epsilon < \mathrm{polylog}(n)/\sqrt{n}$ 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 [4], 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. In addition, we present results for job scheduling on unrelated machines and consensus games.

We show 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 cause best-response dynamics to increase cost from constant to $\Omega(n)$, and that the same is true for consensus games. On the other hand, we show that for job scheduling on identical machines the effect of Byzantine players is very limited. Additionally, for β-nice games having the property that Byzantine players can have only limited impact on the total 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 will be at most $6\beta \cdot \mathbf{OPT} \cdot \mathrm{GAP}$, where GAP bounds the ratio between potential and cost for the game (e.g., for fair cost sharing games $\mathrm{GAP} = O(\log n)$).

Our work can be thought of as asking what kinds of fault-tolerance properties we can guarantee in multi-agent environments. For example, 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, low-order errors in our modeling assumptions, 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 [5, 22, 19]. However, the focus of that work has been on the effect of such players on the quality of the worst Nash or coarse-correlated equilibria. That is, the goal in the Price of Malice is to see how the worst equilibrium deteriorates as we increase the number of malicious agents. In contrast, our Byzantine model studies how much worse the system can get from a given initial configuration for a given *dynamics* (and in particular, for settings where without any perturbations or Byzantine players, behavior would never degrade substantially). As our results show, even a single Byzantine player may have the ability to produce a chain of events leading to an overall deterioration of the system, even if the cost of the worst equilibrium has not changed by much.

We also wish to relate our setting to that of Candogan, Ozdaglar, and Parrilo [10]. Candogan et al. consider the question: suppose a given game G' is *not* a potential game, but is close, in terms of its payoff structure, to some true potential game G. In that case, best response and improved response dynamics in G' may not converge to a pure Nash equilibrium. However, they show that such dynamics will converge to the set of ϵ-equilibria, where ϵ can be bounded in terms of the distance of G' from G [10]. In contrast, our model considers what can be thought of as a fluctuating series of games $G_1, G_2, G_3, \ldots$, each of which is a potential game and each of which is close to an initial game G. Given these fluctuations, our concern is not so much the extent to which the final state produced by best response or improved response dynamics is an approximate equilibrium, but rather the amount by which social cost can increase due to the effect of the fluctuations on these dynamics.

There is also a significant body of work on ways that certain types of random noise can *benefit* game dynamics, by smoothing over local optima. For example, in the noise model studied by Blume [7, 8], noise has the effect that the probability the current player chooses a given action depends exponentially on the gap between the cost of that action and the cost of its best-response action. The result is that the dynamics behaves much like simulated annealing on the global potential function, and so in the low-noise limit will stabilize at states of optimum global potential. Arieli and Young [3] consider noise of a form where players are probabilistically unable to observe the payoff of other action choices, and show how this can improve convergence properties in weakly-acyclic games. Our motivation is instead to examine the possible *dangers* of noise or other imperfections in the cost structure of the game, or in players' perceptions of this cost structure. For this reason, we examine a *much less structured* model of noise, bounding only its magnitude, as well as consider a different question: namely, asking how much worse things could get from a given state if agents follow a best-response or improved-response path under a legal noise sequence.

Organization of this paper: We describe our formal models in Section 2, and present a discussion of design decisions made in our framework in Section 2.3. In Sections 3, 4, and 5 we discuss our perturbation model: giving several useful preliminary properties in Section 3, presenting our results on adversarial-order dynamics in Section 4, and giving our results on random-order dynamics in Section 5. In Section 6 we analyze the case of Byzantine players. Section 7 describes subsequent work on these models and some open questions.

2 The Model

2.1 Background

A game is denoted by a tuple $G = \langle N, (\mathcal{S}_i), (\texttt{cost}_i)\rangle$ where N is a set of n players, $\mathcal{S}_i$ is the finite action space of player $i \in N$, and $\texttt{cost}_i$ is the cost function of player i. The joint action space of the players is $\mathcal{S} = \mathcal{S}_1 \times \ldots \times \mathcal{S}_n$. For a joint action $S \in \mathcal{S}$ we denote by S_{-i} the actions of players $j \neq i$, i.e., $S_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$. Furthermore we denote by $S \oplus s_i'$ the state $(s_1, ..., s_{i-1}, s_i', s_{i+1}, ..., s_n)$. The cost function of player i maps a joint action $S \in \mathcal{S}$ to a real non-negative number, i.e., $\texttt{cost}_i : \mathcal{S} \to \mathbb{R}^+$. Every game has associated a social cost function $\texttt{cost} : \mathcal{S} \to \mathbb{R}$ 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., $\texttt{cost}(S) = \sum_{i=1}^{n} \texttt{cost}_i(S)$, and the max, i.e., $\texttt{cost}(S) = \max_{i=1}^{n} \texttt{cost}_i(S)$. (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 $\mathbf{OPT}(G) = \min_{S \in \mathcal{S}} \texttt{cost}(S)$. 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 action, or set of actions, $BR_i(S)$ that minimizes its cost, given the other players actions S_{-i}. I.e., $BR_i(S) = \arg\min_{s'_i \in \mathcal{S}_i} \texttt{cost}_i\ (S \oplus s'_i)$. Given a joint action S, the *Improved Response (IR)* of player i is the set of actions $IR_i(S)$ that have cost less than or equal to the cost of its current action, i.e., $IR_i(S) = \{s'_i \in \mathcal{S}_i | \texttt{cost}_i(S \oplus s'_i) \leq \texttt{cost}(S)\}$.[3]

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., $s_i \in BR_i(S)$ for every $i \in N$. A *best (improved) response dynamics* is a process in which at each time step, some player who is not playing a best response switches its action to a best (improved) response action, given the current actions of all the other players. In this paper we will concentrate on potential games [18, 21] which have the property that any best response dynamics converges to a pure Nash equilibrium.

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., $(\max_{S \in \mathcal{N}(G)} \texttt{cost}(S))/\mathbf{OPT}(G)$. The *Price of Stability* (PoS) is the ratio between the minimum cost of a Nash equilibrium and the social optimum, i.e., $(\min_{S \in \mathcal{N}(G)} \texttt{cost}(S))/\mathbf{OPT}(G)$.

2.1.1 Congestion Games and Potential Games

Our primary focus is on *congestion games* and important sub-classes of them. A congestion game G is defined by a tuple $(N, \mathrm{R}, (\mathcal{S}_i), (d_r))$, where N is the set of n players, R is a set of m resources, the action set of player i is $\mathcal{S}_i \subseteq 2^{\mathrm{R}}$, and the goal of player i is to play a strategy $S_i \in \mathcal{S}_i$ that minimizes its individual cost $\texttt{cost}_i$. The cost $\texttt{cost}_i(S)$ of player i in state S 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 $d_r : \mathbb{N} \to \mathbb{R}$ is a "delay function" associated with resource r (d_r need not be a monotone function). Closely related to congestion games are potential functions and potential games [18]. A function $\Phi(S)$ is a *generalized ordinal potential function* for game G if for any states S, S' that differ in the action of just one player i, if $\texttt{cost}_i(S') < \texttt{cost}_i(S)$ then $\Phi(S') < \Phi(S)$. If game G has such a function, then G is called a generalized ordinal potential game. Note that for any such game, best response dynamics is guaranteed to reach a pure Nash equilibrium since each move strictly reduces the potential function (which in turn implies a pure Nash equilibrium must exist). A function $\Phi(S)$ is called an *exact* potential function if for any states S, S' that differ in the action of just one player i, $\Phi(S') - \Phi(S) = \texttt{cost}_i(S') - \texttt{cost}_i(S)$; a game having a potential function satisfying this more restrictive condition is called an exact potential game. Rosenthal [21] shows that every congestion game is also an exact potential game by showing that $\Phi(S) = \sum_{r \in \mathrm{R}} \sum_{i=1}^{n_r(S)} d_r(i)$ is an exact potential function for the game. In fact, the converse is true as well: every exact potential game can be written as a congestion game [18], so the notions of congestion game and exact potential game are equivalent.

A congestion game is *symmetric* if all the players share the same set of strategies, otherwise it is asymmetric. Specific classes of congestion games that we study in this work are cost-sharing games, matroid congestion games, consensus games and β-nice games. We define all these in their corresponding sections, namely Sections 4.1, 4.2, 4.3 and 5.1. For games with $\Phi(S) \geq \texttt{cost}(S)$ we define $\mathrm{GAP}(G)$ as $\max_S \Phi(S)/\texttt{cost}(S)$. More generally, for games such that $\Phi(S) \in [c_1\texttt{cost}(S), c_2\texttt{cost}(S)]$ we define $\mathrm{GAP}(G) = c_2/c_1$, where we assume $c_1 \leq 1 \leq c_2$.

Another class of games we study are *load balancing games* (see [20]) which we define in Section 4.4. These are not exact potential games but are generalized ordinal potential games.

[3]We define an Improved Response as cost less than or equal to its current cost, rather than using strict inequality, to allow players to make no change if they choose.

2.2 Price of Uncertainty

In this paper we define and study the *Price of Uncertainty (PoU)* in order to analyze the resilience of potential games to imperfections in player dynamics. We consider two settings: a *perturbation model* in which an adversary can make small perturbations to the cost function between moves, and a *Byzantine player model* in which an adversary can control the actions of a small number of players (i.e., no assumptions about those players' behavior may be made). Within each model, we further consider two variations: one in which the order in which players move is also controlled by an adversary, and one in which the order of movement is random.

Specifically, consider a game G, from a given class of games $\mathcal{G}$, where players start at some given initial configuration S_0 which might or might not be a Nash equilibrium. The game then progresses in phases, where we consider the following two general models for the evolving dynamics:

Perturbation model: In the perturbation model, in each phase t the following occurs. First, an adversary perturbs the costs of G by a small multiplicative factor from their initial values, so that for any S and j we have $\texttt{cost}_j^t(S) \in [\texttt{cost}_j(S)/(1+\epsilon), (1+\epsilon)\texttt{cost}_j(S)]$.[4] Then a player i is chosen who performs a best (improved) response to these perturbed costs, and the new configuration is S_t.

In the *adversarial-order* version of the model, the player i who moves is chosen by an adversary. Our main concern is to upper bound $\texttt{cost}(S_t)/\texttt{cost}(S_0)$ as a function of ϵ and the class of games $\mathcal{G}$. More precisely, let us define $PoU_{BR}(\epsilon, G) = \max \texttt{cost}(S_t)/\texttt{cost}(S_0)$, 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 $PoU_{BR}(\epsilon, \mathcal{G})$ be $\max_{G \in \mathcal{G}} PoU_{BR}(\epsilon, G)$. We define similarly $PoU_{IR}(\epsilon, G)$ for improved response dynamics.

In the *random-order* version, we instead assume that at each time step a *random* agent $i \in N$ is selected. We now care about the *expected* cost at time t. In particular, we define $PoU^R_{BR}(\epsilon, G) = \max_{t, S_0, \mathcal{A}} \mathbf{E}[\texttt{cost}(S_t)]/\texttt{cost}(S_0)$, where the max operator is over time steps t, start states S_0, and adversary strategies $\mathcal{A}$ that for any state S produce some legal perturbation to costs. The expectation is over the sequence of random choices of players to move next. We note that at each time t, first the adversarial perturbation is performed, then a random agent is selected to move next (and then if there are multiple best-responses, the adversary can choose which of those the agent takes). We also point out that for improved response, the adversarial and the random order models are identical, since the adversary can choose to make the selected random player not change its action until his desired player is selected. Thus, we will only consider best-response dynamics.

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 perturbation model 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-order version of these models 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,

[4] We require the adversary's perturbations to be consistent with the class $\mathcal{G}$. For congestion games (see Section 2.1.1) 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.

and 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 (or moving the Byzantine players) so that $\texttt{cost}_i^t(S') \leq \texttt{cost}_i^t(S)$, and for best response dynamics this means $\texttt{cost}_i^t(S')$ 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., $\max_S \max_{v \in V(S)} \texttt{cost}(v)/\texttt{cost}(S)$.

2.3 Discussion of Model and Design Choices

As discussed above, our goal is to understand how small fluctuations in (real or perceived) costs, or small imperfections in the way a game models its players' utilities, can impact the result of natural dynamics. In making an abstraction to study this issue, there are several (reasonable) alternatives one might consider. Here, we discuss a few of the design choices made in our definitions. First, in this work we define the Price of Uncertainty in terms of the worst *state* encountered during the execution of the dynamics, rather than the worst *equilibrium* encountered. The logic behind this choice is that we are interested in the question of safety properties of the system: whether or not the system can be guaranteed to remain in low-cost states. In many cases even a temporary move to a high-cost state might cause players to abandon the system (such as in a ride-sharing system modeled as a cost-sharing game) or to incur severe distress (such as a market speculator purchasing on margin, who might be forced to sell if prices drop sharply even temporarily). Thus, we care about the worst state encountered even if it is not an equilibrium.

A second issue, in the case of random-order dynamics, is over what we are taking the maximum. In this work we define the Price of Uncertainty in terms of the maximum over time periods t, of the expected cost at time t. Namely, we are taking the maximum of an expectation. An alternative, more adversarial, definition would be to take the expectation of the maximum over all t. Note, however, that if in the graph defined in Section 2.2, the set $V(S)$ of nodes reachable from the start state S is strongly-connected, then this would be equivalent to adversarial-order dynamics. That is because the adversary could cause any given state in $V(S)$ to occur eventually.

A third issue, also involving random order dynamics, is that while in these dynamics the players move in a random order, we take an adversarial view of the perturbations. We do this because we want to view our PoU upper bounds as providing strong safety properties, and thus wish to make as few assumptions as possible about any unmodeled motivations or sensing deficiencies of the agents. Random order is a fairly mild assumption as it can naturally be motivated by having exponential clocks, for instance, such as in [8, 26].[5] Studying random perturbations from some family of bounded magnitude, for games for which adversarial perturbations lead to large lower bounds, is an interesting problem that we leave open for future work (see Section 7).

A final issue concerns the focus on best response and improved response dynamics, even though there is an inaccuracy of the model as observed by the players. For example, one might say that players, if they are aware of this inaccuracy, ought to modify their actions only if their benefit is above the inaccuracy parameter. This would guarantee that any action leads to a decrease in potential. However, the flip side is that since this is a worst-case model, players might not want to obey such a discipline and deny themselves the better action just because of the possibility of an inaccuracy or bad series of events. Moreover, if we insist that players modify their actions only if they have a gain large enough, then we might be blocking the players from reaching good states. After all, the search for better states is inherent in the motivation for the existence of dynamics.

[5] In this view, each player has an alarm clock distributed according to a mean-1 exponential distribution, and has a turn to make a move whenever its alarm goes off. Formally, we have an infinite set $\{X_{it}\}$ of independent mean-1 exponentially-distributed random variables, and player i makes its Tth move at time $\sum_{t=1}^{T} X_{it}$.

3 Perturbation model: 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 $\epsilon = 0$, the PoU is simply asking how much the social welfare can deteriorate from some initial state, assuming that all the players are implementing "standard" best (improved) response dynamics. That is, even without fluctuations, by how much could myopic behavior increase the social cost of the system? Note that 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 *In any potential game, for any $\epsilon \geq 0$ we have: $PoU_{IR}(\epsilon, \mathcal{G}) \geq PoU_{BR}(\epsilon, \mathcal{G}) \geq PoU_{BR}(0, \mathcal{G}) \geq PoS(\mathcal{G})$.*

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 (perhaps exponentially small in the game description), 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 G, there is an $\epsilon_0 > 0$ such that for any $\epsilon < \epsilon_0$,*

$$PoU_{IR}(\epsilon, G) = PoU_{IR}(0, G) \text{ and } PoU_{BR}(\epsilon, G) = PoU_{BR}(0, G).$$

Proof: Consider a directed graph $\mathsf{G}_{0,IR}$ (resp., $\mathsf{G}_{0,BR}$) with one vertex for each state S, and a directed edge from state S to state S' if they differ in the action of only one player i and it is an improved response (resp., best response) for player i to move from its action in S to its action in S' even without any perturbation. That is, this is the graph described at the end of Section 2.2 for $\epsilon = 0$. Now, for every pair of states S, S' that does *not* have an edge in G_0 and that differ in the action of a single player i, let $\epsilon_{S,S',IR}$ (resp., $\epsilon_{S,S',BR}$) be the minimum perturbation value that would be needed for it to be an improved response (resp., best response) for player i to move from its action in S to its action in S' (or infinity if no such perturbation is possible). Let $\epsilon_{0,IR} = \min_{S,S'}[\epsilon_{S,S',IR}]$ and $\epsilon_{0,BR} = \min_{S,S'}[\epsilon_{S,S',BR}]$. Then, for any perturbation $\epsilon < \min(\epsilon_{0,IR}, \epsilon_{0,BR})$, the graph of possible moves is identical to case $\epsilon = 0$ and therefore $PoU_{IR}(\epsilon, G) = PoU_{IR}(0, G)$ (resp., $PoU_{BR}(\epsilon, G) = PoU_{BR}(0, G)$). ■

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 on the price of uncertainty for $\epsilon = 0$. In exact potential games, an immediate observation is:

Fact 3.3 *In any exact potential game, in T steps the potential function increases by a multiplicative factor of at most $(1 + \epsilon)^{2T}$.*

Proof: In each step t, if there is a change from state S to state S' via a move of player i, then it must be the case that $\texttt{cost}_i^t(S') \leq \texttt{cost}_i^t(S)$. This implies $\texttt{cost}_i(S') \leq \texttt{cost}_i(S)(1 + \epsilon)^2$. ■

Clearly, Fact 3.3 implies that the potential function increases by at most a multiplicative factor $(1 + \epsilon)^{2L}$, 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 *There exists a class of games $\mathcal{G}$ such that $PoA(\mathcal{G}) = 2$ while $PoU_{BR}(0, \mathcal{G}) = \Theta(\log n)$.*

Proof: See Appendix A. ■

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 the addition of a single player cannot change the cost incurred by any other player by more than a factor α. Specifically, suppose there exist c_1, c_2 with $c_2/c_1 \leq \alpha$ such that for any state $S = (s_1, \ldots, s_n)$ and any s_{n+1}, under state $S' = (s_1, \ldots, s_n, s_{n+1})$ we have for all $1 \leq i \leq n$ that $\texttt{cost}_i(S') \in [c_1\texttt{cost}_i(S), c_2\texttt{cost}_i(S)]$. (For example, in fair cost-sharing games, we would have $c_1 = 1/2$ and $c_2 = 1$: the addition of a new player cannot increase the cost incurred by any player and can help by at most a factor of 2.) 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 6.1, 6.3 and 6.5) translates to a lower bound on $PoU_{BR}(\sqrt{\alpha} - 1, \mathcal{G})$.

4 Perturbation model: 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 [9], 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 $\epsilon \geq \log(n)/\sqrt{n}$ (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 (Theorems 4.5, 4.6, 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 $\mathsf{G} = (V, E)$, which can be directed or undirected, on m edges where each edge $e \in E$ has a nonnegative cost $w_e \geq 0$. There is a set $N = \{1, ..., n\}$ 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 and the *resources* in the game are the edges. In an outcome of the game, each player i chooses a single path $P_i \in \mathcal{S}_i$. Given an outcome $S = (P_1, ..., P_n)$, let $n_e(S)$ be the number of players whose path contains edge e. In the fair cost sharing game, the cost to player i is $\texttt{cost}_i(S) = \sum_{e \in P_i} \frac{w_e}{n_e(S)}$. That is, player i gets to share the cost of each edge e on its path with all other players who are also using that edge. The goal of each player is to connect its terminals with minimum total cost. The social cost is defined to be $\sum_{e \in \cup_i P_i} w_e$, which is the sum of costs of all the players.

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 [21] of these games [2] is the following.

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

4.1.1 Set Covering Games

Set-cover games (SCG) were considered in [9]. 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 w_i of the set with other players who choose the same set. For example, sets could represent different possible locations for placing wireless printers, where players using the same printer agree to split its cost. It is not hard to see that set-cover games are a special case of fair cost sharing games (think of a 3-layer graph where players all have the same common source, each set is an edge in the middle layer, and players each have distinct destinations connected only to edges corresponding to the sets they belong to). We begin with an upper bound for improved-response dynamics.

Theorem 4.1 *In the set-cover 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 i of cost w_i. We view this as a stack of k_i chips, where we label each chip with the name of its initial set and its initial 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 say that position j on stack i has price $d_i(j) = w_i/j$, and will give chip (i,j) a *value* v_{ij} equal to the price of its initial position. So, $\sum_{i,j} v_{ij} = \sum_i \sum_{j=1}^{k_i} w_i/j = \Phi(S_0)$, the potential function of the initial configuration. Furthermore, recall that according to Lemma 4.1, $\Phi(S_0) \leq \texttt{cost}(S_0) \cdot \log n$. 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 $d_{i'}(j') \leq v_{ij} \cdot (1+\epsilon)^{2mn}$. This will yield the theorem, because it means that the potential of the final configuration is at most a factor $(1+\epsilon)^{2mn}$ larger than the potential of the initial configuration; that is, $\Phi(S_t) = \sum_{i'} \sum_{j'=1}^{k_{i'}^t} d_{i'}(j') \leq \sum_{i,j} v_{ij}(1+\epsilon)^{2mn} \leq \Phi(S_0)(1+\epsilon)^{2mn}$ (using $k_{i'}^t$ to denote the number of chips in stack i' at time t).

The argument for the invariant is as follows. First, there are at most mn different positions a given chip can be in (m stacks, n positions per stack) so if you examine the path a chip takes from its initial location to its current location, removing loops, this path has length at most mn. Second, since players follow an improved response dynamics, each step in this move, say from position j_1 in stack i_1 to position j_2 in stack i_2, has the property that $w_{i_2}/j_2 \leq (1+\epsilon)^2 w_{i_1}/j_1$. So, overall, the total increase is at most a factor of $(1+\epsilon)^{2mn}$. This means that

$$\texttt{cost}(S_t) \leq \Phi(S_t) \leq (1+\epsilon)^{2mn}\Phi(S_0) \leq (1+\epsilon)^{2mn}\texttt{cost}(S_0)\log n$$

for all t, as desired. ■

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. The proof is fairly intricate so we give just a sketch of the main ideas here, with a full proof in Appendix A.

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

Proof Sketch: The high-level idea of the construction is that we will have two types of players, Type-I and Type-II, where the perturbations cause the Type-II players to slowly lure the Type-I players from their initial configuration onto substantially more expensive sets. In the case that $\epsilon \geq \sqrt{2}-1$, this can be done in a fairly direct way as follows. Assume $n = 2N-2$. We have N Type-I players and $N-2$ Type-II players. Type-II player k (for $k = 2, 3, \ldots, N-1$) has a personal set of cost N/k that only it can use. Additionally, there are a collection of "public" sets of cost $N-1$ usable by any of the Type-II players. Notice that perturbations of magnitude $\sqrt{2}-1$ are sufficient to bring the Type-II players, one at a time, from their personal sets to any desired one of the public sets (beginning with Type-II player $k=2$, then $k=3$, and so on). Furthermore, even if they are joined on that public set by one additional (Type-I) player, these perturbations are sufficient to bring the Type-II players (in reverse order) back to their personal sets. Using this, the idea of the construction for $\epsilon \geq \sqrt{2}-1$ is that initially, all Type-I players begin on some common set of cost N, and then, one at a time, the Type-II players lure each of them onto different public sets, causing a factor $\Omega(N)$ increase in their cost.

For the case of small perturbations, the construction is more involved. We add an additional set of "helper" players who, via their movements, help to amplify the effect of small perturbations on the personal sets of the Type-II players, though we must be careful to do so in such a way that they do not themselves cause the initial state to be too expensive. In addition, there is a further complication that we want the Type-II players to join onto a *new* "public" set in each round (different from the ones that already have a Type-I player). This could be enforced by the perturbations in the large-perturbation case but can no longer be made a best-response in the above construction when perturbations are smaller. In order to address this issue, we add additional sets and additional restricted Type-II players, as well as change the starting configuration of the Type-I players. The net effect is that Type-I players are brought out from their common sets in waves, and after being lured onto the public sets are then guided onto final "storage" sets before the next wave begins. The full proof appears in 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 (Theorems 4.3 and 4.4). On the other hand, for improved-response dynamics, the costs can grow exponentially even for exponentially small values of ϵ, even for the symmetric (single source, single sink) case (Theorems 4.5 and 4.6).

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

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 in that path 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 any time an unmarked edge is used. We can mark new edges at most m times, so the final cost is at most $\texttt{cost}(S_0)(1 + (1 + \epsilon)^2/n_{min})^m$, where $n_{min} = \min_i n_i$. This implies that as long as $n_{min} = \Omega(m)$ we have $\texttt{cost}(S_t) = O(\texttt{cost}(S_0))$, for all t, as desired. ■

For symmetric fair cost sharing games (SFCSG) 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 $t_1, t_2, \ldots$. 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'_{t_k} at time t_k; that is, $R_k = \mathtt{cost}_{j_{t_k}}(S_{t_k})$. We now claim that

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

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 $P'_{t_{k-1}+1} = P_{t_{k-1}+2}, \ldots, P'_{t_k-1} = P_{t_k}$, any of the players $j_{t_{k-1}+1}, \ldots, j_{t_k}$ 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 $j_{t_{k-1}+2}, \ldots, j_{t_k}$ 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'_{t_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 $P'_{t_{k-1}+1} = P_{t_{k-1}+2}, \ldots, P'_{t_k-1} = P_{t_k}$. Because player j_{t_k} chose path P'_{t_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. ■

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.

Theorem 4.5 *For symmetric fair cost sharing, for any $n \geq 1$, the price of uncertainty for improved-response dynamics satisfies*

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

Proof: We consider the case $n = 1$. The generalization for multiple players is straightforward.

The graph consists of a line with 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. We then have 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} two edges 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 $b = b_{m/2-1}\ldots b_2 b_1 b_0$, 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). ■

We can use Theorem 4.5 to imply a bound also for routing games [20] 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 [25].

Definition 4.1 *A tuple $M := (\mathrm{R}, \mathcal{I})$ 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 $M := (\mathrm{R}, \mathcal{I})$ be a matroid. Let $I \subset \mathrm{R}$; 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 $\mathrm{rk}(M)$ of the matroid. A maximal independent set of M is called a basis of M. In the following we state two additional useful properties of matroids. We denote by $\mathcal{B}$ the set of bases of a matroid, and assume that $B_1, B_2 \in \mathcal{B}$.

Lemma 4.2 *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 $\mathsf{G}(B_1 \bigtriangleup B_2)$ 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 [25]:

Lemma 4.3 *There exists a perfect matching for $\mathsf{G}(B_1 \bigtriangleup B_2)$.*

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 $M_i := (\mathrm{R}, \mathcal{I}_i)$ 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 $\mathrm{rk}(M) = \max_{i \in N} \mathrm{rk}(M_i)$ 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 [15] (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 any matroid game G, we have $PoU_{BR}(\epsilon, G) \leq (1+\epsilon)^{2mn}\mathrm{GAP}(G)$. This implies that for $\epsilon = O(1/(n \cdot m))$, we have $PoU_{BR}(\epsilon, G) = O(\mathrm{GAP}(G))$.*

Proof: The proof proceeds as in Theorem 4.1. However, we initially have $\sum_{i \in N} \mathrm{rk}(M_i) \leq n \cdot \mathrm{rk}(M)$ chips and the $\mathtt{cost}(S_0)$ is within a $\mathrm{GAP}(G)$ 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 $\tilde{d}$. Consider the bipartite graph $\mathsf{G}(s_i^* \bigtriangleup s_i)$ which contains a perfect matching P_M according to Lemma 4.3. 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 $\mathrm{rk}(G)$ chips (corresponding to at most $\mathrm{rk}(G)$ 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).[6] So, for all t we have

$$\Phi(S_t) \leq (1+\epsilon)^{2mn}\Phi(S_0),$$

[6]Note that the dynamics can last for a long time, it's just that we can shortcut them in the argument, because we are simply computing a chain of inequalities.

which implies $\texttt{cost}(S_t) \leq (1+\epsilon)^{2mn}\texttt{cost}(S_0) \cdot \mathrm{GAP}(G)$, which completes the proof. ■

Note: As opposed to the set-cover game 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 Theorem 4.5 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 G of rank $r = m/2$ and $\mathrm{GAP}(G) = O(\log n)$ such that the price of uncertainty for improved-response dynamics satisfies $PoU_{IR}(\epsilon, G) \geq 1 + 2(2^{m/2} - 1)\epsilon/m$.*

4.3 Consensus Games

Consensus games [11] are played by users viewed as vertices in a connected, undirected simple graph $\mathsf{G} = (N, E)$ with n vertices, where $N = \{1, ..., n\}$. Each player i has two actions R or B (red or blue), i.e., $\mathcal{S}_i = \{R, B\}$. A player has cost 1 for each incident edge on which he disagrees with his neighbor: $\texttt{cost}_i(S) = \sum_{(i,j)\in E} \mathbf{I}_{(\mathbf{s_i} \neq \mathbf{s_j})}$. The overall social cost is $\texttt{cost}(S) = \sum_{i\in N} \texttt{cost}_i(S)$. It is straightforward to show that these games are congestion games [18] and that the potential function can be rewritten as be $\Phi(S) = \texttt{cost}(S)/2$. 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 (note: these bounds have subsequently been improved [6], see Section 7.1).

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

Proof: Let $k = 1/\epsilon$. Consider the following graph $\mathsf{G}(V, E)$ over $V = \{v_1, \ldots v_n\}$. Node v_i is connected to node v_{i+1} with k parallel edges. In addition, all the nodes are connected to node v_n with an additional edge. This is a consensus game, so each node needs to select a color $\{B, R\}$. The initial configuration S_0 has all the nodes v_j choose action $s_j = B$ except node v_1 which chooses $s_1 = R$. The cost of the initial configuration is $2(k + 1)$.

At time step $i > 1$ the adversary will make node v_i flip its action to $s_i = R$. Specifically, at this time node v_i has two neighbors v_{i-1} with $s_{i-1} = R$ and v_{i+1} with $s_{i+1} = B$, and also a direct edge to v_n (where $s_n = B$). The adversary increases the weights of the edges to v_{i-1} by a factor of $\gamma = (k + 1)/k$ and decreases the weight of the edges to v_{i+1} by a factor of $k/(k + 1)$. (Recall that $\epsilon = 1/k$.) Now the adversary lets node v_i do a best response, and it changes its action to $s_i = R$. After time step $n - 1$ we have nodes $\{v_1, ...v_{n-1}\}$ using action R and node v_n has $s_n = B$. The cost of S_{n-1} is $2(n - 1 + k)$, since each edge adjacent to v_n adds 1 to the cost of both endpoints. The price of uncertainty is therefore at least $\frac{n+k-1}{k+1} = \frac{(n-1)\epsilon+1}{1+\epsilon} = \Omega(1 + n\epsilon)$.

For $\epsilon > \sqrt{2} - 1$ we use a different construction. Note that now, if a node has k neighbors voting R and $2k$ voting B we can switch its vote from B to R. We consider the following network G. The network is built from $\log(n/4)$ levels, and levels $2i$ and $2i + 1$ have 2^i nodes. Between two adjacent level we have full connectivity. (Namely, the number of nodes in the levels are: $1, 2, 2, 4, 4, 8, 8, 16, 16, ..., n/4, n/4$ and full connectivity between levels)

Start with the single node at level 1 being B and all the others R. The cost is exactly 2, At each step we can make another level flip values from R to B. We can do this since the nodes in level $2i$ have 2^{i-1} neighbors in level $2i - 1$ and $2^i = 2 \cdot 2^{i-1}$ neighbors in level $2i + 1$. (Similarly, nodes in level $2i + 1$ have 2^i neighbors in level $2i$ and 2^{i+1} neighbors in level $2i + 2$.) At the end we have all the edges in the last two levels connecting $n/4$ to $n/4$ nodes, which is $\Omega(n^2)$ edges. Since one layer has B and the other has R the cost is $\Omega(n^2)$. ■

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 $PoU_{BR}(\epsilon, WCG) \geq (1+\epsilon)^{n-2}$.*

Proof: Consider a line graph $\mathsf{G}(V,E)$ over $V = \{v_1, \ldots v_n\}$, where v_i is connected to node v_{i+1} by a single edge of weight $(1+\epsilon)^i$. The initial configuration S_0 has all the nodes have $s_j = B$ except node v_1 which has $s_1 = R$. The cost of the initial configuration is $2(1+\epsilon)$.

At time step $i \geq 2$ the adversary makes node v_i flip its action to $s_i = R$. At time i node v_i has neighbor v_{i-1} with $s_{i-1} = R$ and v_{i+1} with $s_{i+1} = B$. The adversary increases the weights of the edge to v_{i-1} by a factor of $1+\epsilon$ and decreases the weights of the edges to v_{i+1} by a factor of $1+\epsilon$. Now the adversary lets node v_i do a best response, and its changes its action to $s_i = R$. After time step $n-1$ we have nodes $\{v_0, ... v_{n-1}\}$ using action R and node v_n has $s_n = B$. The cost of S_{n-1} is $2(1+\epsilon)^{n-1}$. Therefore, the price of uncertainty is at least $(1+\epsilon)^{n-2}$. ■

4.4 Job Scheduling

In this section we consider *job scheduling on unrelated machines (JSUM)* (see [20]) 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 we 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., $L_i(S) = \sum_{j \in B_i(S)} c_{i,j}$ 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., $\texttt{cost}_j(s) = L_{s_j}(S)$. For the social cost we use the *makespan*, which is the load on the most loaded machine, i.e., $\texttt{cost}(S) = \max_i L_i(S)$. The price of stability in these games is 1, since there is always a pure Nash equilibrium which is also socially optimal [14].

The Price of Uncertainty in these games can be exponentially large, even for two machines, when ϵ is large compared to $1/n$.

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

Proof: Assume we have $n/2$ jobs of type 1 with loads $(1, X_i)$, and $n/2$ jobs of type 2 with loads $(X_i, 1)$, where $i \in \{1, \ldots, n/2\}$. We define a sequence of numbers L_i as follows: $L_0 = n/2$ and $L_k = (1+\epsilon)L_{k-1} - 1 = (1+\epsilon)^k L_0 - \sum_{j=0}^{k-1}(1+\epsilon)^j$. We set $X_k = \epsilon L_{k-1}$.

The initial configuration S_0 is the optimal configuration, i.e., type 1 jobs are on machine M_1 and type 2 on machine M_2. The cost of the initial configuration is $n/2$, i.e., $\texttt{cost}(S_0) = n/2 = L_0$.

We will consider an adversary that divides the time to $n/2$ phases. At the end of phase k it will maintain that the load on both machines is identical, and equal to $\texttt{cost}(S_k) = L_k$. During a phase k the adversary does the following. First the adversary increases the cost of each job on machine M_1 by a factor of $(1+\epsilon)$, and divides the costs on machine M_2 by $(1+\epsilon)$. The new load on machine M_1 is $(1+\epsilon)L_{k-1}$ while the load on machine M_2 is $L_{k-1}/(1+\epsilon)$. Now the adversary lets job k of type 1 move from machine M_1 to machine M_2, since it is a best response for it. (Since $(1+\epsilon)L_{k-1} = L_{k-1} + X_k$.) Next, the adversary returns the costs on both machines to their original weights. The new load on machine M_2 is $(1+\epsilon)L_{k-1}$ while the load on machine M_1 is $L_{k-1} - 1$. Now the adversary lets job k of type 2 move from machine M_2 to machine M_1, since it is a best response for it. At the end of phase k the load on each machine is $L_k = L_{k-1} - 1 + X_k = (1+\epsilon)L_{k-1} - 1$.

After phase $n/2$ we have that $\texttt{cost}(S_{n/2}) = L_{n/2} = (1+\epsilon)^{n/2}(n/2) - \frac{(1+\epsilon)^{n/2}-1}{\epsilon}$. Recall that the price of uncertainty is at least $\texttt{cost}(S_{n/2})/\texttt{cost}(S_0)$. Since $1 > \frac{2}{\epsilon n}$ the theorem follows. ■

For job scheduling on identical machines (JSIM)—that is, if there exist job weights w_j such that $c_{ij} = w_j$ for all i, j—we have a simple upper bound, even for large perturbations.

Theorem 4.12 *For job scheduling on identical machines, $PoU_{IR}(\epsilon, JSIM) \leq 2(1+\epsilon)^2$.*

Proof: Let $W = \sum_j w_j$ be the sum of the weights of all the jobs. In any configuration, any job j has a best response whose cost is at most $W/m + w_j$. Therefore, a perturbation cannot cause job j to move to a machine whose cost is greater than $[W/n + w_j](1+\epsilon)^2$. Note that $OPT \geq \max\{W/m, w_j\} \geq \frac{1}{2}[W/n + w_j]$. So, the PoU is at most $2(1+\epsilon)^2$. ■

5 Perturbation model: Random Order

We now consider the effect of perturbations on *random* order best-response dynamics. Our main result is that for the broad class of β-nice games introduced by Awerbuch et al. [4], 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 $\epsilon > 0$. In particular, we show $PoU^R_{BR}(\frac{1}{32}, G) \leq 5\beta \cdot \mathrm{GAP}(G)$ for any β-nice game G. 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 $\texttt{cost}_i$. For each player i define $\Delta_i(S) = \texttt{cost}_i(S) - \texttt{cost}_i(S^i)$ and $\Delta(S) = \sum_i \Delta_i(S)$.

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

By rewriting Definition 5.1 as $\Delta(S) \geq \frac{1}{2}[\texttt{cost}(S) - \beta\mathbf{OPT}(G)]$ it is clear that the Price of Anarchy for a β-nice game is at *most* β. Awerbuch et al. [4] show that in fact a number of important games are β-nice for Price of Anarchy *equal* to β. This includes 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 ($\beta = 2$). This is of interest because, as shown by [4], certain natural dynamics will quickly converge in β-nice games to states of cost not much more than $\beta\mathbf{OPT}(G)$.

Here we show that β-nice games additionally have the property that the expected price of uncertainty in the random order model is only $O(\beta \cdot \mathrm{GAP}(G))$ even for constant $\epsilon > 0$. We start by showing that if the true (unperturbed) cost of the current configuration S is greater than $2\beta \cdot \mathbf{OPT}(G)$, then no matter how the adversary adjusts the costs, the expected drop in potential is at least $\texttt{cost}(S)(1/4 - 4\epsilon)/n$. For $\epsilon < 1/16$, this is $\Omega(\texttt{cost}(S)/n)$. That is, the adversary may make the cost exceed $2\beta\texttt{cost}(S_0)$ but only temporarily.[7] We next use this to bound $PoU^R_{BR}(\epsilon, G)$ for such games G.

Lemma 5.1 *For any β-nice game G, for $\epsilon \leq 1/32$, if $\texttt{cost}(S_t) \geq 2\beta\mathbf{OPT}(G)$ then $\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] \leq -\texttt{cost}(S_t)/(8n)$.*

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

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

1. $\tilde{\Delta}_i(S) = \texttt{cost}_i^t(S) - \texttt{cost}_i^t(\tilde{S}^i)$ 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. $\hat{\Delta}_i(S) = \texttt{cost}_i(S) - \texttt{cost}_i(\tilde{S}^i)$ 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 $\texttt{cost}(S) > 2\beta\mathbf{OPT}$. Then by definition of β-nice we have $\Delta(S) > \texttt{cost}(S)/4$. 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:

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

Now, summing over all i we have:

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

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

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

Now, since $\texttt{cost}_i^t(S) \leq \texttt{cost}_i(S)(1+\epsilon)$ and $\texttt{cost}_i^t(\tilde{S}^i) \geq \texttt{cost}_i(\tilde{S}^i)(1-\epsilon)$, we have:

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

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

$$\begin{aligned}
E_i[\hat{\Delta}_i(S)] &\geq (\texttt{cost}(S)[1/4 - 2\epsilon]/n - 2\epsilon\texttt{cost}(S)/n)/(1-\epsilon) \\
&\geq \texttt{cost}(S)[1/4 - 4\epsilon]/n,
\end{aligned}$$

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 $\epsilon \leq 1/32$, we then get the desired result. ■

So, Lemma 5.1 shows if the true (unperturbed) cost of current configuration S is greater than $2\beta \cdot \mathbf{OPT}$, then no matter how the adversary adjusts the costs, the expected drop in potential is at least $\texttt{cost}(S)(1/4 - 4\epsilon)/n$. 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 this does not necessarily imply that once cost is low it will stay there forever – just that if the adversary is ever able to make the cost go above $2\beta\mathbf{OPT}(G)$ 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* $\texttt{cost}(S_t)$, $\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] \leq 2\epsilon\texttt{cost}(S)/(n(1-\epsilon))$.

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

Assume $\epsilon \leq 1/32$. 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 $O(\beta \cdot \mathrm{GAP}(G))$ even for constant $\epsilon > 0$. Recall that we define $\mathrm{GAP}(G) = c_2/c_1$ where $c_1 \leq 1 \leq c_2$ are values such that for any state S we have $\Phi(S) \in [c_1\mathtt{cost}(S), c_2\mathtt{cost}(S)]$.

Theorem 5.1 *For any β-nice game G, any $\epsilon \leq 1/32$, and any $t > 0$, we have*

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

Therefore, $PoU^R_{BR}(\epsilon, G) \leq 5\beta \cdot \mathrm{GAP}(G)$.

Proof: We will show that if $\mathbf{E}[\mathtt{cost}(S_t)] \geq 4\beta\mathbf{OPT}$ 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 $\mathbf{E}[\mathtt{cost}(S_t)] \leq 4\beta\mathbf{OPT}$, 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 $\mathbf{E}[\mathtt{cost}(S_t)] \geq 4\beta\mathbf{OPT}$. Let p_t be the probability that $\mathtt{cost}(S_t) \geq 2\beta\mathbf{OPT}$. Therefore, we have:

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

Since we assume that $\mathbf{E}[\mathtt{cost}(S_t)] \geq 4\beta\mathbf{OPT}$ we have $\mathbf{E}[\mathtt{cost}(S_t)|\mathtt{cost}(S_t) \geq 2\beta\mathbf{OPT}] \geq 2\beta\mathbf{OPT}/p_t$. Now, using Lemmas 5.1 and 5.2, we can write:

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

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

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:

Theorem 5.2 *For $M = 2$ machines, for any $\epsilon > 2/n$, we have that $PoU^R_{BR}(\epsilon, JSUM) = \Omega(\epsilon n)$.*

Proof: We use a similar construction to Theorem 4.11: we have $n/2$ jobs of type 1 with loads $(1, X_i)$, and $n/2$ jobs of type 2 with loads $(X_i, 1)$, but we now set all $X_i = X = \epsilon n/2$. The initial configuration S_0 is the optimal configuration, i.e., type 1 jobs are on machine M_1 and type 2 on machine M_2, and the cost of S_0 is $n/2$.

The adversary will now cause states to transition according to the following Markov chain. Let $even_k$ denote the state in which both machines have k jobs of load X and $n/2 - k$ jobs of load 1, and let odd_{k+1} denote the state in which one machine has $k + 1$ jobs of load X and $n/2 - k$ jobs of load 1 while the other machine has k jobs of load X and $n/2 - k - 1$ jobs of load 1. So, the initial state of the system is $even_0$.

If the system is in state $even_k$, then the adversary increases the cost of each job on machine M_1 by a factor $(1 + \epsilon)$. Now, if the random job chosen to move next is on machine M_2, it will prefer to stay on its current machine and we remain in state $even_k$. However, if the job chosen is on the higher-load machine M_1, then it will move no matter which type of job it is. Thus, overall we have a $\frac{1}{2}$ chance of staying put, a $\frac{k}{n}$ chance of transitioning to odd_k, and a $\frac{1}{2} - \frac{k}{n}$ chance of transitioning to odd_{k+1}. On the other hand, if the system is in state odd_{k+1}, then the adversary does no adjustment to the loads, so the difference between the loads on the two machines is $X + 1$. Therefore if the random job chosen to move next is on the less loaded machine it will stay put, but if it is on the more highly loaded machine then it will move, no matter which type of job it is. Thus, overall we have a $\frac{1}{2} - \frac{1}{n}$ chance of staying put, a $\frac{k+1}{n}$ chance of transitioning to $even_k$, and a $\frac{1}{2} - \frac{k}{n}$ chance of transitioning to $even_{k+1}$.

Thus, the behavior of the system follows a random walk on the line $even_0, odd_1, even_1, odd_2, \ldots$, with greater probability of transitioning to the right than left, up until the state $even_{n/4}$ whose left and right transition probabilities are equal. In particular, for $k \leq n/8$, we have probability at least $3/8$ in each step of moving to the right, and probability at most $1/8$ of moving to the left. This implies that for $t = n/4$, the expected value of k at time t is $\Omega(n)$. Since the makespan for a state of index k is at least $k\epsilon n/2$, this in turn implies that for $t = n/4$, the expected cost is $\Omega(\epsilon n^2)$, yielding a randomized price of uncertainty of $\Omega(\epsilon n)$. ■

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 $PoU_{BR}^{R}(\epsilon, WCG) \geq (1+\epsilon)^{n/2-2}$.*

Proof: We use the construction of Theorem 4.10 but with edge weights of $(1 + \epsilon)^{i/2}$. Specifically, at round i, by keeping the edge weight between v_{i-2} and v_{i-1} high, the adversary can ensure that v_i is the only node that has incentive to switch its action. In addition, we begin with nodes v_1 and v_2 both red, rather than just v_1. Finally, after round $n - 2$ (so v_{n-1} and v_n are now the only two blue nodes) the adversary adjusts weights to keep this an equilibrium. Thus we have $PoU_{BR}^{R}(\epsilon, WCG) \geq (1 + \epsilon)^{n/2-2}$. ■

6 The Byzantine Model

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 a lower bound for set-cover 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, and an upper bound for β-nice games, showing that random-order best-response dynamics are resilient to even a large number of Byzantine players. We also give results for job-scheduling and consensus games.

6.1 Set-cover 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 $O(\mathbf{OPT})$ to an equilibrium of cost $\Omega(n \cdot \mathbf{OPT})$.

Note that this is the largest gap possible since the Price of Anarchy for this game is n.

Theorem 6.1 *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: Let $N = (n+1)/2$. 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 $k \in \{2, \ldots, N-1\}$. In addition, we have one Byzantine player who may choose any of the sets, for a total of $2N-1 = n$ 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 $f_{N-1}, f_{N-2}, \ldots, f_2$, 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 type II player k, for $k = 2, 3, \ldots, N-2$, 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. ∎

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

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

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

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

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. ■

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

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

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

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

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

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

Theorem 6.2 *For any $t > 0$, $\mathbf{E}[\Phi(S_t)] \le \max[6\beta\mathbf{OPT}, \mathtt{cost}(S_0)] \le 6\beta\mathtt{cost}(S_0)$.*

Proof: We will show that if $\mathbf{E}[\Phi(S_t)] \ge 5\beta\mathbf{OPT}$ then $\mathbf{E}[\Phi(S_{t+1})] \le \mathbf{E}[\Phi(S_t)]$. Specifically, suppose $\mathbf{E}[\Phi(S_t)] \ge 5\beta\mathbf{OPT}$, which implies that $\mathbf{E}[\mathtt{cost}^t(S_t)] \ge 5\beta\mathbf{OPT}$ since $\mathtt{cost}^t(S_t) \ge \mathtt{cost}(S_t) \ge \Phi(S_t)$. Let p_t be the probability that $\mathtt{cost}^t(S_t) \ge 2\beta\mathbf{OPT}$. Note that

$$\begin{aligned}
\mathbf{E}[\mathtt{cost}^t(S_t)] &= p_t\mathbf{E}[\mathtt{cost}^t(S_t)|\mathtt{cost}^t(S_t) \ge 2\beta\mathbf{OPT}] + (1-p_t)\mathbf{E}[\mathtt{cost}^t(S_t)|\mathtt{cost}^t(S_t) \le 2\beta\mathbf{OPT}] \\
&\le p_t\mathbf{E}[\mathtt{cost}^t(S_t)|\mathtt{cost}^t(S_t) \ge 2\beta\mathbf{OPT}] + 2\beta\mathbf{OPT},
\end{aligned}$$

so we have $\mathbf{E}[\mathtt{cost}^t(S_t)|\mathtt{cost}^t(S_t) \ge 2\beta\mathbf{OPT}] \ge 3\beta\mathbf{OPT}/p_t$. Now, using Lemmas 6.1 and 6.2, we can write:

$$\begin{aligned}
\mathbf{E}[\Phi(S_{t+1}) - \Phi(S_t)] &\le p_t\left(\frac{-1}{8n}\right)\mathbf{E}[\mathtt{cost}^t(S_t)|\mathtt{cost}^t(S_t) \ge 2\beta\mathbf{OPT}] \\
&\quad +(1-p_t)\frac{1}{8n}\mathbf{E}[\mathtt{cost}^t(S_t)|\mathtt{cost}^t(S_t) < 2\beta\mathbf{OPT}] \\
&\le \frac{-3\beta\mathbf{OPT}}{8n} + \frac{2\beta\mathbf{OPT}}{8n} \\
&< 0.
\end{aligned}$$

■

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

6.3 Job scheduling and consensus games

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

Theorem 6.3 *For two machines and one Byzantine job and all individual costs at most 1, the cost can increase from 2 to $\Omega(n)$, even in the random-order model.*

Proof: Consider two machines and $2n+1$ good jobs, $n+1$ 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, jobs of type II are on machine 2, and the Byzantine player is on machine 2; this is the optimal assignment. Now, the Byzantine job goes to machine 1, then a job of type I moves to machine 2, since this is its best response. Next, 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 continues until all the good jobs have switched places and 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 $\Omega(n)$. ■

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 $2OPT + W_b$, 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 $OPT \geq \max\{W/m, w_i\}$ and thus the result follows. ■

We end with a simple observation that for unweighted consensus, a single Byzantine player can cause cost to increase from 0 to a cost $\Omega(n)$.

Theorem 6.5 *For the unweighted consensus game, a single Byzantine player can increase cost from* 0 *to* $\Omega(n)$.

Proof: The network is simply a line network $v_1, \ldots, v_n$. 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 $v_1, \ldots, v_k$ alternating between R and B, such that v_k plays B. During phase k, first v_{k+1} switches to B, and then players $v_k, \ldots v_1$ switch their action. At the end we have all players alternating between R and B, at cost $\Omega(n)$. ■

7 Conclusion

7.1 Subsequent Work

Subsequent to the initial conference publication of this work, Balcan et al. [6] have strengthened several of the bounds presented in this paper, in particular for consensus and set-cover games.

For unweighted consensus games (UGC), they establish $PoU_{BR}(\varepsilon, UCG) = \Omega(n^2\varepsilon^3)$ for $\varepsilon = \Omega(n^{-1/3})$ and $PoU_{IR}(\varepsilon, UCG) = O(n^2\varepsilon)$ for any ε.[8] This shows among other things that the Price of Uncertainty is quadratic in n for any constant $\epsilon > 0$ (Theorem 4.9 had only a linear lower bound for general ϵ). In the Byzantine-player model, they tightly quantify the effect of b Byzantine players, showing tight upper and lower bounds of $\Theta(n\sqrt{nb})$ on the effect such players can have. Note that even for the case of $b = 1$, this improves on Theorem 6.5 by a $\Theta(\sqrt{n})$ factor.

For set-cover games, [6] improve on the bound of Theorem 4.1 for the case of $m \ll n$ (many more players than sets) giving an upper bound that depends only on the number m of sets. Specifically, they show $PoU_{IR}(\epsilon) = (1+\epsilon)^{O(m^2)}O(\log m)$ for $\epsilon = O(\frac{1}{m})$. They also provide a lower bound, building on the construction of Theorem 4.2, showing that there exist set-cover game instances such that even for a *friendly* ordering of players (any ordering at all in which no players experience starvation) an adversary can schedule perturbations causing the initial state to increase in cost by a factor $\Omega(\epsilon n^{1/3}/\log n)$, for any $\epsilon = \Omega(n^{-1/3})$.

Balcan et al. [6] also show how results presented here for the class of β-nice games can be extended to the class of (λ, μ)-smooth games of Roughgarden [23].

7.2 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?

Another natural question is to address the *computational* problem of computing the Price of Uncertainty for a given game instance. That is, given an instance of a game (e.g., a weighted graph along with a collection

[8] Since in UCG each player has only two actions, the IR and BR models are equivalent

of (s_i, t_i) pairs in the case of cost-sharing games), given a value of ϵ (or a number of Byzantine players), and perhaps also given a start state S, compute how much worse a state S' could be reached via perturbed best-response or improved-response dynamics. This would be analogous to questions asked in the distributed systems and verification literature about whether a given system can reach an undesirable configuration.

Acknowledgements: We would like to thank Heiko Röglin for a number of helpful discussions about matroid congestion games. We also thank Florin Constantin and Steven Ehrlich for pointing out a bug in an earlier construction for Theorem 4.2. Finally, we would like to thank the anonymous referees for their helpful comments.

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] I. Arieli and P. Young. Fast convergence in population games. Technical Report 570, Oxford, 2011.

[4] 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.

[5] M. Babaioff, R. Kleinberg, and C.H. Papadimitriou. Congestion games with malicious players. In *ACM EC*, 2007.

[6] M.-F. Balcan, F. Constantin, and S. Ehrlich. The snowball effect of uncertainty in potential games. In *WINE*, 2011.

[7] L. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.

[8] L. Blume. How noise matters. *Games and Economic Behavior*, 44:251–271, 2003.

[9] N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative cost sharing games via subsidies. In *SAGT*, pages 337–349, 2008.

[10] O. Candogan, A. Ozdaglar, and P.A. Parrilo. Dynamics in near-potential games. Technical Report 2872, LIDS, 2011.

[11] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *Proc. 23rd STACS*, 2006.

[12] A. Czumaj and B. Voecking. Tight bounds for worst-case equilibria. In *Proc. 13th ACM-SIAM SODA*, 2002.

[13] B. de Finetti. La prevision: Ses lois logiques, ses sources subjectives. *Annales de lInstitut Henri Poincare*, 7:1–68, 1937.

[14] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3), 2007.

[15] 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.

[16] M. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6), 1998.

[17] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proc. 16th STACS*, pages 404–413, 1999.

[18] D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[19] T. Moscibroda, S. Schmid, and R.Wattenhofer. When selfish meets evil: byzantine players in a virus inoculation game. In *PODC*, 2006.

[20] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge, 2007.

[21] R.W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *Int. Journal Game Theory*, 2(1), 1973.

[22] A. Roth. The price of malice in linear congestion games. In *WINE*, pages 118–125, 2008.

[23] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, 2009.

[24] L. J. Savage. *The Foundations of Statistics*. John Wiley and Sons, 1954.

[25] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2002.

[26] D. Shah and J. Shin. Dynamics in congestion games. In *SIGMETRICS*, pages 107–118, 2010.

[27] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.

[28] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS*, 2002.

[29] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

A Additional Proofs

Theorem 3.1 *There exists a class of games $\mathcal{G}$ such that $PoA(\mathcal{G}) = 2$ while $PoU_{BR}(0, \mathcal{G}) = \Theta(\log n)$.*

Proof: We specifically consider the class of market-sharing games [28]. A market-sharing game is specified by a bipartite graph with nodes on the left representing the players and nodes on the right representing different markets. Each market j has some given benefit value b_j. Every player may choose one market j from among its neighbors, and receives benefit b_j/n_j, where n_j is the total number of players (including itself) who also chose market j.

For these market-sharing games, the bound of $PoA = 2$ is from [28]. For the lower bound on PoU, consider players $\{2, \ldots, n\}$ where player i can select between a dedicated market r_i with benefit $1/i$ and a shared market r_1 with benefit 1. Initially, each player i uses its dedicated market r_i, and the social welfare is $\ln n$. Now let the players perform best response in an increasing order of the indices. Player i has a benefit of $1/i$ from its dedicated market r_i, and the benefit of $1/(i-1)$ from the shared market r_1, so it prefers to move to r_1. This implies that at the end of the sequence the social welfare is 1. (Note that this is not an equilibrium.) The upper bound follows immediately from the fact that the gap between potential and cost in this game satisfies $\text{GAP} = O(\log n)$. ■

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

Proof: As pointed out in Section 3, the bound for the case $\epsilon \geq \sqrt{2}-1$ follows immediately from Theorem 6.1. Thus it suffices to focus on the case $\epsilon < \sqrt{2}-1$.

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 $k = 3, \ldots, N$, 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 $\frac{1}{\epsilon}N/k$; 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 $f_{k,j}$).

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 $i = 1, 2, \ldots, \sqrt{N}$, the following occurs. First, for each $j = 1, 2, \ldots, \sqrt{N}$ 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 $k = 3, 4, \ldots$ are made to follow along to $s_{active,j}$. In the case of $k = 3, \ldots, 1/\epsilon$, 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 can easily cause these players to move). After the above process has been completed for all $j = 1, 2, \ldots, \sqrt{N}$ (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 jth helper-player has a private set of cost $\frac{N}{k(1-j\epsilon)}$. By perturbing costs, adversary can cause these players for $j = 1, 2, \ldots, 1/(2\epsilon)$ to move in order to their private sets. Specifically, at the time the jth 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 $j = 1/(2\epsilon), \ldots, 2, 1$ (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. ∎