

The Snowball Effect of Uncertainty in Potential Games

Maria-Florina Balcan, Florin Constantin, and Steven Ehrlich

College of Computing, Georgia Institute of Technology
{ninamf,florin,sehrlich}@cc.gatech.edu

Abstract. Uncertainty is present in different guises in many settings, in particular in environments with strategic interactions. However, most game-theoretic models assume that players can accurately observe interactions and their own costs. In this paper we quantify the effect on social costs of two different types of uncertainty: adversarial perturbations of small magnitude to costs (effect called the Price of Uncertainty (PoU) [3]) and the presence of several players with Byzantine, i.e. arbitrary, behavior (effect we call the Price of Byzantine behavior (PoB)). We provide lower and upper bounds on PoU and PoB in two well-studied classes of potential games: consensus games and set-covering games.

1 Introduction

Uncertainty, in many manifestations and to different degrees, arises naturally in applications modeled by games. In such settings, players can rarely observe accurately and assign a precise cost or value to a given action in a specific state. For example a player who shares costs for a service (e.g. usage of a supercomputer or of a lab facility) with others may not know the exact cost of this service. Furthermore, this cost may fluctuate over time due to unplanned expenses or auxiliary periodic costs associated with the service. In a large environment (e.g. the Senate or a social network), a player may only have an estimate of the behaviors of other players who are relevant to its own interests. Another type of uncertainty arises when some players are misbehaving, i.e., they are Byzantine.

The main contribution of this paper is to assess the long-term effect of small local uncertainty on cost-minimization *potential games* [7]. We show that uncertainty can have a strong *snowball effect*, analogous to the increase in size and destructive force of a snowball rolling down a snowy slope. Namely, we show that small perturbations of costs on a per-player basis or a handful of players with Byzantine (i.e. adversarial) behavior can cause a population of players to go from a good state (even a good equilibrium state) to a state of much higher cost. We complement these results highlighting the lack of robustness under uncertainty with guarantees of resilience to uncertainty. We assess the effects of uncertainty in two important classes of potential games using the framework introduced by [3]. The first class we analyze is that of *consensus games* [2, 6] for which relatively little was previously known on the effect of uncertainty. The second class we analyze is that of *set-covering games* [5], for which we improve on the previously known bounds of Balcan et al. [3]. We review in detail the uncertainty models and these classes of games, as well as our results below.

We consider both *improved-response* (IR) dynamics in which at each time step exactly one player may update strategy in order to lower his (apparent) cost and *best-response* (BR) dynamics in which the updating player chooses what appears to be the least costly strategy. Any state is assigned a *social cost*, which for most of our paper is defined as the sum of all players' costs in that state. We measure the effect of uncertainty as the maximum multiplicative increase in social cost when following these dynamics. We instantiate this measure to each type of uncertainty.

For the first uncertainty type, we assume adversarial perturbations of costs of magnitude at most $1 + \epsilon$ for $\epsilon > 0$ (a small quantity that may depend on game parameters). That is, a true cost of c may be perceived as any value within $[\frac{1}{1+\epsilon}c, (1 + \epsilon)c]$. Consider a game G and an initial state S_0 in G . We call a state S (ϵ, IR) -reachable from S_0 if there exists a valid ordering of updates in IR dynamics and corresponding perturbations (of magnitude at most ϵ) leading from S_0 to S . The *Price of Uncertainty* [3] (for IR dynamics) given ϵ of game G is defined as the ratio of the highest social cost of an (ϵ, IR) -reachable state S to the social cost of starting state S_0 .

$$PoU_{IR}(\epsilon, G) = \max \left\{ \frac{cost(S)}{cost(S_0)} : S_0; S (\epsilon, IR)\text{-reachable from } S_0 \right\}$$

For a class \mathcal{G} of games and $\epsilon > 0$ we define $PoU_{IR}(\epsilon, \mathcal{G}) = \sup_{G \in \mathcal{G}} PoU_{IR}(\epsilon, G)$ as the highest PoU of any game G in \mathcal{G} for ϵ . PoU_{BR} is defined analogously.

For the second uncertainty type, we assume B additional players with arbitrary, or Byzantine [8] behavior. We define the *Price of Byzantine behavior* ($PoB(B)$) as the effect of the B Byzantine players on social cost, namely the maximum ratio of the cost of a state reachable in the presence of B Byzantine agents to that of the starting state.

$$PoB(B, G) = \max \left\{ \frac{cost(S)}{cost(S_0)} : S_0; S B\text{-Byz-reachable from } S_0 \right\}$$

where state S of G is B -Byz-reachable from S_0 if some valid ordering of updates by players (including the B Byzantine ones) goes from S_0 to S . $PoB(B, \mathcal{G}) = \sup_{G \in \mathcal{G}} PoB(B, G)$ for class \mathcal{G} . PoB, like PoU, may depend on the dynamics¹.

A low PoU or PoB shows resilience of a system to small errors by players in estimating costs or behavior of others. In the games we study, social costs cannot increase much without uncertainty (namely $PoU(0) = PoB(0)$ are small), yet modest instances of uncertainty (in terms of ϵ or B) can lead to significant increases in costs (i.e. large $PoU(\epsilon)$ and $PoB(B)$). We introduce in the following the classes of games we study and summarize our results.

Consensus games [6] model a basic strategic interaction: choosing one side or the other (e.g. in a debate) and incurring a positive cost for interacting with each agent that chose the other side. More formally, there are two colors (or strategies), white (w) and red (r), which each player may choose; hence IR and

¹ We omit parameters from PoU and PoB when they are clear from context.

BR dynamics coincide. Each player occupies a different vertex in an undirected graph G with vertices $\{1, \dots, n\}$ and edges $E(G)$ (without self-loops). A player's cost is defined as the number of neighbors with the other color. We establish $PoU(\epsilon) = \Omega(n^2\epsilon^3)$ for $\epsilon = \Omega(n^{-1/3})$ and $PoU(\epsilon) = O(n^2\epsilon)$ for any ϵ . These bounds are asymptotically tight for constant ϵ . We exactly quantify $PoB(B)$ as $\Theta(n\sqrt{nB})$ by exhibiting an instance with $\Theta(n\sqrt{nB})$ edges that is flippable (i.e. it can flip from one monochromatic state to the other given B Byzantine players) and then reducing any other consensus game to this instance.

Set-covering games [5] model many applications where all users of a resource share fairly its base cost. These natural games fall in the widely studied class of fair-cost sharing games [1]. In a set-covering game, there are m sets, each set j with its own fixed weight (i.e. base cost) w_j . Each of the n players must choose exactly one set j (different players may have access to different sets) and share its weight equally with its other users, i.e. incur cost $w_j/n_j(S)$ where $n_j(S)$ denotes the number of users of set j in state S . We prove $PoU_{IR}(\epsilon) = (1 + \epsilon)^{O(m^2)}O(\log m)$ for $\epsilon = O(\frac{1}{m})$ — this is small for a small number of resources even if there are many players. This improves over the previous bounds of [3], which had an additional dependence on the number of players n . We also improve the existing lower bound for these games (due to [3]) to $PoU_{IR}(\epsilon) = \Omega(\log^p m)$ for $\epsilon = \Theta(\frac{1}{m})$ and any constant $p > 0$. Our new lower bound is a subtle construction that uses an intricate “pump” gadget with finely tuned parameters. A pump replaces, in a non-trivial recursive manner with identical initial and final pump states, one chip of small cost with one chip of high cost. Finally, we show a lower bound of $PoU_{BR}(\epsilon) = \Omega(\epsilon n^{1/3}/\log n)$ for $\epsilon = \Omega(n^{-1/3})$ and $m = \Omega(n)$ which is valid even if an arbitrary ordering of player updates is specified a priori, unlike the existing lower bound of [3].

We note that our lower bounds use judiciously tuned gadgets that create the desired snowball effect of uncertainty. Most of them hold even if players must update in a specified order, e.g. round-robin (i.e. cyclically) or the player to update is the one with the largest reduction in (perceived) cost. Our upper bounds on PoU hold no matter which player updates at any given step.

Due to the lack of space we only provide sketches for most proofs in this paper. Full proofs appear in the long version of the paper [4].

2 Consensus Games

In this section, we provide lower and upper bounds regarding the effect of uncertainty on consensus games. Throughout the section, we call an edge *good* if its endpoints have the same color and *bad* if they have different colors. The social cost is the number of bad edges (i.e. half the sum of all players' costs) plus one, which coincides with the game's potential. Thus $PoU(0) = PoB(0) = 1$. Since the social cost is in $[1, n^2]$, $PoU(\epsilon) = O(n^2), \forall \epsilon$ and $PoB(B) = O(n^2), \forall B$.

2.1 Lower Bound and Upper Bound for Perturbation Model

Perturbation model. The natural uncertainty here is in the number of neighbors of each color that a vertex perceives. We assume that if a vertex i has n' neighbors of some color, then a perturbation may cause i to perceive instead an arbitrary integer in $[\frac{1}{1+\epsilon}n', (1+\epsilon)n']$. Since each action's cost is the number of neighbors of the other color, this is a cost perturbation model. In this model, only an $\epsilon = \Omega(\frac{1}{n})$ effectively introduces uncertainty.² We also assume $\epsilon \leq 1$ in this section.

Theorem 1. $PoU(\epsilon, \text{consensus}) = \Omega(n^2\epsilon^3)$ for $\epsilon = \Omega(n^{-1/3})$ even for arbitrary orderings of player updates.

Proof sketch. We sketch below the three components of our construction assuming that the adversary can choose which player updates at any step. We can show that the adversary can reduce an arbitrary ordering (that he cannot control) to his desired schedule by compelling any player other than the next one in his schedule not to move.

- The *output* component has $k_{out} = \Theta(\frac{1}{\epsilon}(\log n - 2\log \frac{1}{\epsilon}))$ levels, where any level $i \geq 0$ has $\frac{1}{\epsilon}(1+\epsilon)^i$ nodes. Each node on level $i \geq 1$ is connected to all the nodes on levels $i-1$ and $i+1$.
- The *input* component consists of two cliques, K_{red} and K_{white} , each of size $1/\epsilon^2$, and each of these nodes is connected to the first output level. The dynamics are “seeded” from the input component.
- There is an *initializer* gadget³ with $k_{init} = \frac{1}{\epsilon} \log \frac{2}{\epsilon}$ levels, where level i has $\frac{1}{\epsilon}(1+\epsilon)^i$ nodes. Again, each node on level $i \geq 1$ is connected to all the nodes on levels $i-1$ and $i+1$. Every node of the final level of the initializer is connected to all nodes in the clique K_{red} .

We require $\epsilon \geq n^{-1/3}$ so that the number $\Theta(\frac{1}{\epsilon^3})$ of initializer nodes is at most a constant fraction of the total n nodes. Then the number $\Theta(\frac{1}{\epsilon^2}(1+\epsilon)^{k_{out}})$ of output nodes will be a large fraction of the n nodes.

The initial coloring is: all output nodes white, K_{red} and K_{white} white, and all initializer nodes white except for the first two levels which are red. We thus initially have $\frac{(1+\epsilon)^3}{\epsilon^2}$ bad edges, all in the initializer. Throughout the dynamics, for both the initializer and output components, the nodes in each level have the same color, except for the level that is currently being updated. The schedule consists of two epochs. The first epoch is for all the initializer nodes and the clique K_{red}

² If $\epsilon < 1/3n$, consider a node with r red neighbors, w white neighbors, and $r > w$. Since r and w are both integers, $r \geq w+1$. For a cost increasing move to occur, we must have $(1+\epsilon)^2w \geq r \geq w+1$. This implies that $n \geq w \geq 1/(2\epsilon+\epsilon^2) \geq 1/3\epsilon > n$.

³ Without this initializer, we can get a worse lower bound $PoU(\epsilon, \text{consensus}) = \Omega(n^2\epsilon^4)$, for a wider range of $\epsilon = \Omega(n^{-1/2})$, again for an arbitrary ordering. The main difference is that K_{red} is initially red and the initial state has $\Theta(\frac{1}{\epsilon^3})$ bad edges. In the long version of our paper, we additionally show that when the adversary can control the ordering of updates to match its schedule, we can improve both this lower bound and that of Theorem 1 to $\Omega(n^2\epsilon^3)$ for the range $\epsilon = \Omega(n^{-1/2})$.

to change color. Thereafter they are left alone. At this point, the adversary can prevent the clique nodes from changing their color, and he can change the color of nodes in the first output level at will. Indeed, these nodes have $\frac{1}{\epsilon^2}$ neighbors of either color in each clique and $\Theta(\frac{1}{\epsilon})$ neighbors in the second output level, difference small enough to be overcome by a $(1 + \epsilon)$ -factor perturbation. The second epoch has a phase for each two consecutive output levels i and $i + 1$ in which these levels obtain their final color and then are never considered again. This is achieved by changing all prior levels to the intended color of i and $i + 1$.

In the final state we have the first two output levels colored red, the next two colored white, then the next two red, and so on. The final number of bad edges is $\Omega(n^2\epsilon)$. Since we started with only $\Theta(\frac{1}{\epsilon^2})$ bad edges the number of bad edges has increased by a factor of $\Omega(n^2\epsilon^3)$. Thus $PoU(\epsilon, consensus) = \Omega(n^2\epsilon^3)$. \square

We note that the previously known lower bound of $\Omega(1 + n\epsilon)$ due to Balcan et. al [3] was based on a much simpler construction. Our new bound is better by a factor of at least $n^{1/3}$ for $\epsilon = \Omega(1/n^3)$. We also note that since $PoU(\epsilon) = O(n^2)$ for any ϵ , it implies a tight PoU bound of $\Theta(n^2)$ for *any* constant ϵ . We also provide a PoU *upper bound* for consensus games that depends on ϵ . It implies that the existing $\Theta(n^2)$ lower bound cannot be replicated for any $\epsilon = o(1)$. The proof is based on comparing the numbers of good and bad edges at the first move that increases the social cost.

Theorem 2. $PoU(\epsilon, consensus) = O(n^2\epsilon)$.

2.2 Tight Bound for Byzantine Players

As described earlier, Byzantine players can choose their color ignoring their neighbors' colors (and therefore their own cost). Note however the Byzantine players cannot alter the graph⁴. In this section we show a tight bound on the effect of B Byzantine players, for any B : the effect of one Byzantine player is very high, of order $n\sqrt{n}$ and that the subsequent effect of $B \leq n$ Byzantine players is proportional to the square root of B . As was the case for PoU , the effect of uncertainty is decomposed multiplicatively into a power of n and a power of the extent of uncertainty (ϵ for PoU, B for PoB).

Theorem 3. $PoB(B, consensus) = \Theta(n\sqrt{n \cdot B})$.

The proof of the $O(n\sqrt{n \cdot B})$ upper bound follows from Lemmas 1, 2 and 3 below. The key to this bound is the notion of a flippable graph. For any consensus game, let S_{red} be the configuration where all nodes are red, and similarly let S_{white} be the configuration where all nodes are white.

⁴ For the lower bound, we assume that a player will break ties in our favor when he chooses between two actions of equal cost. With one more Byzantine player the same bound holds even if players break ties in the worst possible way for us. For the upper bound, we assume worst possible players' moves from the social cost point of view.

Definition 1 (*B-Flippable graph*). Consider graph G on n vertices of which B are designated special nodes and the other $n - B$ nodes are called normal. We say G is *B-flippable* (or just *flippable* when B is clear from context) if in the consensus game defined on G where the special nodes are the Byzantine agents, the state S_{white} is *B-Byz-reachable* from S_{red} .

We now describe the concept of a *conversion dynamics* in a consensus game which we use in several of our proofs. In such a dynamics, we start in a state where all vertices are red and have Byzantine players change their color to white. Then all normal nodes are allowed in a repeated round-robin fashion to update, so long as they are currently red. This ends when either every vertex is white or no vertex will update its color.

We note that in a flippable graph the conversion dynamics induces an ordering of the normal vertices: nodes are indexed by how many other white nodes are present in total at the state when they change color to white. We note that there may be more than one valid ordering. In the following with each *B-flippable* graph, we shall arbitrarily fix a canonical ordering (by running the conversion dynamics). Where there is sufficient context, we shall use v a vertex interchangeably with its index in this ordering. Using this ordering we induce a canonical orientation by orienting edge uv from u to v if and only if $u < v$. We also orient all edges away from the B special nodes. To simplify notation, we shall write $v_{in} = |\delta^-(v)|$ and $v_{out} = |\delta^+(v)|$ for a vertex v 's in-degree and out-degree respectively. We note that by construction, for a flippable graph we have $v_{in} \geq v_{out}$. One can easily show the following:

Claim 1. *A graph is B-flippable if and only if there exists an ordering on the $n - B$ normal vertices of the graph such that, in the canonical orientation of the edges, every normal vertex v has $v_{in} \geq v_{out}$. A graph is B-flippable if and only if for every pair of states S, S' , S is B-Byz-reachable from S' .*

Lemma 1. *Fix a game G on n vertices, B of which are Byzantine, and a pair of configurations S_0 and S_T such that S_T is B-Byz-reachable from S_0 . If $\text{cost}(S_0) \leq n$, then there exists a B-flippable graph F with at most $3n$ nodes and at least $\text{cost}(S_T)$ edges (in total).*

Proof Sketch. The proof has two stages. In the first stage, we construct a consensus game G' and configuration S'_T such that $\text{cost}(S'_T) \geq k$, S'_T is B-Byz-reachable from S_{red} , but $V(G') \leq 3n$. In the second stage, we delete some edges of G' to create a graph G'' , showing that S_{white} is B-Byz-reachable from S_{red} in G'' (thus G'' is flippable) while ensuring that $E(G'') \geq k$.

We first describe the construction of G' . We separate the nodes of G into two sets I_r and I_w based on their color in the initial configuration S_0 . For each edge that is bad in S_0 we introduce a ‘mirror’ gadget. A mirror consists of a single node whose color the adversary can easily control and a helper node to change the color. The nodes of G' are all the nodes of G and at most $2n$ nodes from mirror gadgets. The edges of G' are all edges that are good in state S_0 of G , and at most $5n$ edges introduced by the mirrors.

In G' , the nodes of I_r and I_w interact with each other only indirectly, via mirrors that are controlled by the adversary. Using this fact, the adversary can simulate the dynamics from S_0 to S_T on I_r . For any state S , let \bar{S} be the state in which every node has the opposite color from in S . The adversary can also simulate the dynamics over states in which the color of every node has been reversed. Thus the adversary can simulate dynamics from \bar{S}_0 to \bar{S}_T in I_w .

At the end of this process, every edge that was bad in S_T is also bad in this final state S'_T . Note that the initial state is one in which all nodes in I_r are red, and all nodes in I_w are red (since they are red in \bar{S}_0). Thus the dynamics lead to S'_T from S_{red} . Thus we have created G' and a state S_T where $cost(S'_T) \geq k$ and G' is flippable, but $|V(G')| \leq 3n$.

In the second stage, we identify a set of edges in G' , and delete them to form G'' . This set of edges are precisely those whose endpoints both remain red in the conversion dynamics. We show that these edges are not bad in any state, hence none of these edges are bad in S'_T , and secondly in G'' , S_{white} is B -Byz-reachable from S_{red} . The first statement implies that $cost(S'_T) \leq |E(G'')|$, and the second statement implies that G'' is flippable, which was what we wanted. \square

Definition 2 ($F_{seq}(n, B)$). *Let $F_{seq}(n, B)$ be the B -flippable graph with $n - B$ normal nodes with labels $\{1, 2, \dots, n - B\}$. There is an edge from each special node to each normal node. Every normal node v satisfies $v_{out} = \min(v_{in}, (n - B) - v)$, and v is connected to the nodes of $\{v + 1, \dots, v + v_{out}\}$. This is called the no-gap property. In general, if $k = \min(v_{in}, n - v)$ then v has out-arc set $\{v + 1, \dots, v + k\}$.*

By claim 1 we immediately get that F_{seq} is B -flippable. Our upper bound follows by showing $|E(F)| \leq |E(F_{seq}(n, B))|$ for any flippable graph F on n vertices. For this, we take a generic flippable graph and transform it into F_{seq} without reducing the number of edges. We say there is a *gap*(a, b, c) for $a < b < c$ if vertex a does not have an edge to b but does have an edge to c . Note that this is defined in terms of an ordering on the vertices; we use the conversion ordering for each graph.

Lemma 2. *A flippable graph on n vertices has at most as many edges as $F_{seq}(n, B)$.*

Proof sketch. We prove this by inducting on the lexicographically minimal gap of flippable graphs. If a gap is present, then we can either add or move edges to create a lexicographically greater gap. Eventually this eliminates all gaps without reducing the number of edges. Since a graph with no gaps is a subgraph of $F_{seq}(n, B)$, we have bounded the number of edges in any flippable graph. \square

Our last lemma tightly counts the number of edges in $F_{seq}(n, B)$ via an inductive argument and thus, by Lemma 2, it also upper bounds the number of edges in any flippable graph.

Lemma 3. *If $B \leq \frac{n}{2}$, the flippable graph $F_{seq}(n, B)$ has $\Theta(n\sqrt{nB})$ edges.*

Proof sketch. We count the edges by counting the number of in-edges to a given node. By induction, we show that the first node to have jB in-edges has index $\binom{j+1}{2}B + 1$. This implies that any node $k \in [n]$ has $\Theta(\sqrt{kB})$ in-edges. Summing over all n nodes, we find there are $\Theta(n\sqrt{nB})$ edges in the graph in total. \square

Proof of Theorem 3. We first argue that the $PoB(B, \text{consensus}) = O(n\sqrt{nB})$. Consider a consensus graph G on n nodes, and a pair of configurations S_0 and S_T B -Byz-reachable from S_0 . If $B \geq n/2$, then the statement is trivial, so we may assume that $B < n/2$. We assume $\text{cost}(S_0) < n$: if $\text{cost}(S_0) \geq n$, since G has fewer than n^2 edges, we get $PoB(B, G) \leq n^2/n = n$. Denote by $k := \text{cost}(S_T) - 1$ the number of bad edges in S_T . By Lemma 1, we demonstrate a flippable graph F on fewer than $3n$ nodes, with at least k edges. By Lemma 2, F has at most as many edges as $F_{seq}(3n, B)$, which has only $O(n\sqrt{nB})$ edges by Lemma 3. We get $PoB(B) = O(n\sqrt{nB})$.

It will now be enough to prove that $PoB(B, F_{seq}(n, B)) = \Omega(n\sqrt{nB})$. We claim now that if G is a flippable graph with m edges, then $PoB(B, G) \geq \frac{m}{2}$. We get this via the following probabilistic argument using the fact that the adversary can color G arbitrarily (by claim 1). Consider a random coloring of the graph, where each node is colored white independently with probability $1/2$. The probability an edge is bad is $1/2$, so in expectation, there are $m/2$ bad edges. Thus some state has at least $m/2$ bad edges and it is reachable via dynamics from any other state (claim 1) since G is a flippable graph. This establishes $PoB(B, G) \geq \frac{m}{2}$. Since $F_{seq}(n, B)$ is flippable and it has $m = \Theta(n\sqrt{nB})$ edges, we get $PoB(B, F_{seq}(n, B)) = \Omega(n\sqrt{nB})$. \square

In contrast to the existing bound $PoB(1) = \Omega(n)$, our bound is parametrized by B , sharper (by a $\Theta(\sqrt{n})$ factor for $B = 1$) and asymptotically tight.

3 Set-Covering Games and Extensions

Set-covering games (SCG) are a basic model for fair division of costs, and have wide applicability, ranging e.g. from a rental car to advanced military equipment shared by allied combatants. A set-covering game admits the potential function $\Phi(S) = \sum_{j=1}^m \sum_{i=1}^{n_j(S)} \frac{w_j}{i} = \sum_{j=1}^m \Phi^j(S)$ where $\Phi^j(S) = \sum_{i=1}^{n_j(S)} \frac{w_j}{i}$. $\Phi^j(S)$ has an intuitive representation as a stack of $n_j(S)$ *chips*, where the i -th chip from the bottom has a cost of w_j/i . When a player i moves from set j to j' one can simply move the topmost chip for set j to the top of stack j' . This tracks the change in i 's costs, which equals by definition the change in potential Φ . We will only retain the global state (number of players using each set) and discard player identities. This representation has been introduced for an existing PoU_{IR} upper bound of [3]; we refine it for our improved upper bound.

SCGs have quite a small gap between potential and cost [1]: $\text{cost}(S) \leq \Phi(S) \leq \text{cost}(S)\Theta(\log n)$, $\forall S$. Hence without uncertainty, the social cost can only increase by a logarithmic factor: $PoU(0) = PoB(0) = \Theta(\log n)$.

3.1 Upper Bound for Improved-Response

We start with an upper bound on PoU_{IR} in set-covering games that only depends on the number m of sets.

Theorem 4. $PoU_{IR}(\epsilon, \text{set-covering}) = (1 + \epsilon)^{O(m^2)} O(\log m)$ for $\epsilon = O(\frac{1}{m})$.

In particular for $\epsilon = O(\frac{1}{m^2})$ we obtain a logarithmic $PoU_{IR}(\epsilon)$.

Proof of Theorem 4. We let J_0 denote the sets initially occupied and $W_0 = \text{cost}(S_0) = \sum_{j \in J_0} w_j$ be their total weight. We discard any set not used during the dynamics.

With each possible location of a chip at some height i (from bottom) in some stack j , we assign a *position of value*⁵ w_j/i . Thus a chip's cost equals the value of its position in the current state. We will bound the cost of the m most expensive chips by bounding costs of expensive positions and moves among them.

It is easy to see that any set has weight at most $W_0(1 + \epsilon)^{2(m-1)}$ (clearly the case for sets in J_0). Indeed, whenever a player moves to a previously unoccupied set j' from a set j , the weight of j' is at most $(1 + \epsilon)^2$ times the weight of j ; one can trace back each set to an initial set using at most $m - 1$ steps (there are m sets in all). We also claim that at most $mi(1 + \epsilon)^{2m}$ positions have value at least $\frac{W_0}{i}$, $\forall i$: indeed positions of height $i(1 + \epsilon)^{2m}$ or more on any set have value less than $\frac{W_0}{i}$ since any set has weight at most $W_0(1 + \epsilon)^{2(m-1)}$.

Fix a constant $C > (1 + \epsilon)^{2m}$ (recall $\epsilon = O(\frac{1}{m})$). Note that any chip on a position of value less than $\frac{W_0}{m}$ in S_0 never achieves a cost greater than $\frac{W_0}{m}(1 + \epsilon)^{2Cm^2}$. Indeed, by the reasoning above for $i = m$, there are at most $m \cdot m \cdot (1 + \epsilon)^{2m} \leq Cm^2$ positions of greater value. Thus the chip's cost never exceeds $\frac{W_0}{m}(1 + \epsilon)^{2Cm^2}$ as it can increase at most Cm^2 times (by an $(1 + \epsilon)^2$ factor).

We upper bound the total cost of the *final* m most expensive chips, as it is no less than the final social cost: for a set, its weight equals the cost of its most expensive chip. We reason based on chips' initial costs. Namely, we claim $h(i) \leq \frac{W_0}{i-1} \cdot (1 + \epsilon)^{2Cm^2}$, $\forall i \in [m]$, where $h(i)$ denotes the cost of i th most expensive chip in the final configuration. If this chip's initial cost is less than $\frac{W_0}{m}$ then the bound follows from the claim above. Now consider all chips with an initial cost at least $\frac{W_0}{m}$. As argued above, at most Cm^2 positions have value $\frac{W_0}{m}$ or more, and any of these chips increased in cost by at most $(1 + \epsilon)^{2Cm^2}$. A simple counting argument⁶ shows that for any i , there are at most i chips of initial cost at least $\frac{W_0}{i}$ and thus $h(i) \leq \frac{W_0}{i-1} \cdot (1 + \epsilon)^{2Cm^2}$, $\forall i$.

⁵ We refer to the weight of a set, the cost of a chip and the value of a position

⁶ We claim that for any k , there are at most k chips of initial cost at least $\frac{W_0}{k}$. Let J_0 , be the set of initially used resources. For each $j \in J_0$, let r_j be set j 's fraction of the initial weight (i.e. $w_j = r_j W_0$), and let p_j be the number of initial positions with value greater than $\frac{W_0}{k}$ in set j . We have $\frac{w_j}{p_j} = \frac{r_j W_0}{p_j} \geq \frac{W_0}{k}$, implying $p_j \leq kr_j$. Counting the number of initial positions with sufficient value yields $\sum_j p_j \leq \sum_j kr_j = k \sum_j r_j = k$ since $\sum_j r_j = 1$.

As the i^{th} most expensive chip has cost at most $\frac{W_0}{i-1}(1+\epsilon)^{2Cm^2}$ (at most $i-1$ chips have higher final cost),

$$\begin{aligned}\sum_{i=1}^m h(i) &= h(1) + \sum_{i=2}^m h(i) \leq h(1) + \sum_{i=2}^m \frac{W_0}{i-1}(1+\epsilon)^{2Cm^2} \\ &= O(W_0(1+\epsilon)^{2m} + W_0(1+\epsilon)^{2Cm^2} \log m) = W_0 \cdot (1+\epsilon)^{O(m^2)} O(\log m)\end{aligned}$$

As desired, $PoU_{IR}(\epsilon, \text{set-covering}) = (1+\epsilon)^{O(m^2)} O(\log m)$ as the final social cost is at most $\sum_{i=1}^m h(i)$. \square

The existing bound [3] is $PoU_{IR}(\epsilon) = O((1+\epsilon)^{2mn} \log n)$. Unlike our bound, it depends on n (exponentially) and it does not guarantee a small $PoU_{IR}(\epsilon)$ for $\epsilon = \Theta(\frac{1}{m^2})$ and $m = o(n)$. This bound uses chips in a less sophisticated way, noting that any chip can increase its cost (by $(1+\epsilon)^2$) at most mn times.

Our technique also yields a bound of $PoU_{BR}(\epsilon) = (1+\epsilon)^{O(m^2)} O(\log m)$ for $\epsilon = O(\frac{1}{m})$ in matroid congestion games [3] – see the full version for details [4]. These games are important in that they precisely characterize congestion games for which arbitrary BR dynamics (without uncertainty) converge to a Nash equilibrium in polynomial time.

3.2 Lower Bound for Improved-Response

Our upper bound showed that $PoU_{IR}(\epsilon)$ is logarithmic for $\epsilon = O(\frac{1}{m^2})$. A basic example (one player hopping along sets of cost $1, (1+\epsilon)^2, \dots, (1+\epsilon)^{2(m-1)}$), applicable to many classes of games, yields the lower bound $(1+\epsilon)^{2(m-1)} \leq PoU_{IR}(\epsilon, \text{set-covering})$. In fact, this immediate lower bound was the best known on $PoU_{IR}(\epsilon)$. For $\epsilon = \omega(\frac{1}{m})$, we get that $PoU_{IR}(\epsilon)$ is large. An intriguing question is what happens in the range $[\omega(\frac{1}{m^2}), \Theta(\frac{1}{m})]$, in particular for natural uncertainty magnitudes such as $\epsilon = \Theta(\frac{1}{m})$ or $\epsilon = \Theta(\frac{1}{n})$.

In this section we show that for $\epsilon = \Theta(\frac{1}{\min(m,n)})$, PoU can be as high as polylogarithmic. We provide a construction that repeatedly uses the snowball effect to locally increase one chip's cost, without other changes to the state. Our main gadget is a *pump*, which is used as a black box in the proof. A pump increases a chip's cost by $\alpha = \log n'$, where $n' = \min(m, n)$. We use p pumps to increase each chip's cost by a $\Omega(\log^p n')$ factor. As pumps are “small”, the total cost increase is $\Omega(\log^p n')$.

Theorem 5. $PoU_{IR}(\epsilon, \text{set-covering}) = \Omega(\log^p \min(m, n))$, for $\epsilon = \Theta(\frac{1}{\min(m, n)})$ and constant $p > 0$.

Before providing a sketch of Theorem 5, we provide the formal definition of a pump. An (α, W) -pump uses $O(\frac{1}{\epsilon})$ sets and $O(2^\alpha)$ players to increase, one by one, an arbitrary number of chip costs by an α factor from W/α to W . For ease of exposition, we assume $m = \Theta(n)$ and we only treat $p = 2$, i.e. how to achieve $PoU_{IR}(\frac{1}{n}) = \Omega(\log^2 n)$. For general p , we use p pump gadgets instead of two.

Definition 3 (Pump). An (α, W) -pump P is an instance of a set-covering game specified as follows:

- The number m_P of sets used is $O(\frac{1}{\epsilon})$. For our choice of ϵ , $m_P = O(n)$. The total weight W_P of all sets in P that are initially used is in $(2^\alpha W, e2^\alpha W)$. The number of players used is $n_P = 2^{\alpha+1} - 2$.
- Within $O(n^3)$ moves of IR dynamics contained within the pump, and with a final state identical to its initial state, a pump can consume any chip of cost at least W/α to produce a chip of cost W .

Proof sketch of Theorem 5. Let $N := \alpha^2 2^\alpha$. The number of players will be $n := N + n_{P_1} + n_{P_2}$. Thus $\alpha = \Theta(\log n)$. Note that each player can use any set.

We use two pumps, an $(\alpha, 1/\alpha)$ pump P_1 , and a $(\alpha, 1)$ pump P_2 . Aside from the pumps, we have Type-I, Type-II and Type-III sets, each with a weight of $1/\alpha^2$, $1/\alpha$ and 1 respectively. At any state of the dynamics, each such set will be used by no player or exactly one player. In the latter case, we call the set *occupied*. We have N Type-I sets, 1 Type-II set, and N Type-III sets, i.e. $m := 2N + 1 + m_{P_1} + m_{P_2} = \Theta(n)$ sets in all.

Let $\text{cfg}(i, j, k)$ refer to the configuration with i Type-I sets occupied, j Type-II sets occupied, and k Type-III sets occupied. We shall use $2N + 1$ intermediate states, denoted state_i . Our initial state is $\text{state}_0 = \text{cfg}(N, 0, 0)$, and our final configuration will be $\text{state}_{2N} = \text{cfg}(0, 0, N)$. In general, $\text{state}_{2i} = \text{cfg}(N - i, 0, i)$ and $\text{state}_{2i+1} = \text{cfg}(N - i - 1, 1, i)$. Thus we want to move each player on a Type-I set (initially) to a corresponding Type-III set, an α^2 increase in cost. To this purpose, we will pass each such player through the first pump and move it on the Type-II set. This achieves the transition from state_{2i} to state_{2i+1} . Since the player's cost is increased by an α factor (from $\frac{1}{\alpha^2}$ to $\frac{1}{\alpha}$), we can pass it through the second pump and then move it on the Type-III set. This achieves the transition from state_{2i+1} to state_{2i+2} .

The social cost of our initial configuration is $W_0 = N \cdot \frac{1}{\alpha^2} + W_{P_1} + W_{P_2} \leq 2^\alpha + e \cdot \frac{1}{\alpha} 2^\alpha + e \cdot 2^\alpha \leq 7 \cdot 2^\alpha$. The final social cost (excluding the pumps) is at least $N = \alpha^2 2^\alpha$. Thus $\text{PoU} = \Omega(\alpha^2)$, and $\alpha = \Theta(\log n)$.

Finally, we note that the pump is constructed using $1 + 1/\epsilon$ sets $s_0, \dots, s_{\frac{1}{\epsilon}}$ where set s_i has weight $W(1 + \epsilon)^i$. Additionally there are α ‘storage’ sets t_j with weight $2W/j$. In the initial configuration, the sets s_1, \dots, s_{2^α} are occupied, with set s_i having $\max(0, \alpha - \lceil \log_2 1 + i \rceil)$ players on it, for $i \geq 1$. s_0 also has $\alpha - 1$ players. The pump is activated by a chip of cost W/α moving onto set s_0 . The chips then advance into a configuration where set s_i has precisely $\max(0, \alpha - \lceil \frac{1}{\epsilon} - i + 1 \rceil)$ chips. Note that this roughly doubles the cost of each chip. The chips on set $s_{\frac{1}{\epsilon}}$ then use the storage sets to capture their current cost. The chip in storage with cost $2W$ exits the pump, and the other chips in storage fill up set s_0 . All the other chips can return to the initial configuration by making only cost decreasing moves. Note that the chip to leave the pump is not the same chip that entered. \square

In the full version of the paper, we show how our pump gadget can be tweaked to provide a polylogarithmic lower bound on PoU for generalized set-covering games with *increasing* delay functions, as long as they have bounded jumps,

i.e. if an additional user of a resource cannot increase its cost by more than a constant factor.

3.3 Lower Bound for Best-Response

We also show that a significant increase in costs is possible for a large range of ϵ even if players follow best-response dynamics with arbitrary orderings. This construction will use more sets than players, and so will not contradict Theorem 4

Theorem 6. $PoU_{BR}(\epsilon, \text{set-covering}) = \Omega(\epsilon n^{1/3} / \log n)$, for any $\epsilon = \Omega(n^{-1/3})$. This holds for any arbitrary ordering of the dynamics, i.e. no matter which player is given the opportunity to update at any time step.

Proof sketch. The proof has two stages. Our first step provides a construction which shows that in a set-covering game with best-response dynamics, the adversary can compel an increase of social cost that is polynomial (of fractional degree) in n . With our new construction, we then separately show that the adversary can cause this increase even when it cannot control which players update – we show that the adversary can cause only the relevant players to update. \square

We note that previous work provided a stronger lower bound of $\Omega(\epsilon n^{1/2} / \log n)$ [3], but which only works for a specific ordering of the updates.

4 Open Questions

It would be interesting to close our gap on PoU for consensus games. It would also be interesting to study a model where the perturbations are not completely adversarial, but instead chosen from some distribution of bounded magnitude.

Acknowledgements. This work was supported by NSF grants CCF-0953192 and CCF-1101215, by ONR grant N00014-09-1-0751, and by a Microsoft Research Faculty Fellowship.

References

1. E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS*, pages 295–304, 2004.
2. B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *EC*, 2008.
3. M.-F. Balcan, A. Blum, and Y. Mansour. The price of uncertainty. In *EC*, 2009.
4. M.-F. Balcan, F. Constantin, and S. Ehrlich. The snowball effect of uncertainty in potential games. www.cc.gatech.edu/~ninamf/papers/snowball-long.pdf, 2011.
5. N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative cost sharing games via subsidies. In *SAGT*, pages 337–349, 2008.
6. G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *STACS*, 2006.
7. D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
8. T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: byzantine players in a virus inoculation game. In *PODC*, pages 35–44, 2006.