Improved Equilibria via Public Service Advertising

Maria-Florina Balcan* Avrim Blum† Yishay Mansour‡

Abstract

Many natural games have both high and low cost Nash equilibria: their Price of Anarchy is high and yet their Price of Stability is low. In such cases, one could hope to move behavior from a high cost equilibrium to a low cost one by a "public service advertising campaign" encouraging players to follow the low-cost equilibrium, and if every player follows the advice then we are done. However, the assumption that *everyone* follows instructions is unrealistic. A more natural assumption is that some players will follow them, while other players will not. In this paper we consider the question of to what extent can such an advertising campaign cause behavior to switch from a bad equilibrium to a good one even if only a fraction of people actually follow the given advice, and do so only temporarily. Unlike the "value of altruism" model, we assume everyone will ultimately act in their own interest.

We analyze this question for several important and widely studied classes of games including network design with fair cost sharing, scheduling with unrelated machines, and party affiliation games (which include consensus and cut games). We show that for some of these games (such as fair cost sharing), a random α fraction of the population following the given advice is sufficient to get a guarantee within an $O(1/\alpha)$ factor of the price of stability for any $\alpha > 0$. For other games (such as party affiliation games), there is a strict threshold (in this case, $\alpha < 1/2$ yields almost no benefit, yet $\alpha > 1/2$ is enough to reach near-optimal behavior). Finally, for some games, such as scheduling, no value $\alpha < 1$ is sufficient. We also consider a "viral marketing" model in which certain players are specifically targeted, and analyze the ability of such targeting to influence behavior using a much smaller number of targeted players.

*School of Computer Science, Carnegie Mellon University. Supported in part by NSF grant CCF-0514922, by an IBM Graduate Fellowship, and by a Google Research Grant.

†School of Computer Science, Carnegie Mellon University. Supported in part by NSF grant CCF-0514922 and by a Google Research Grant

‡School of Computer Science, Tel Aviv University and Google Research. This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant no. 1079/04 from the Israel Science Foundation and by a grant from United States-Israel Binational Science Foundation (BSF). This publication reflects the authors' views only.

1 Introduction

Understanding the *quality* of Nash equilibria in a game has been a major focus of algorithmic game theory. The main motivation is to understand what is the additional cost that is incurred when we assume that agents are behaving in a strategic manner, rather than performing a global optimization.

Koutsoupias and Papadimitriou [14] proposed the notion of *Price of Anarchy (PoA)* as the ratio of the cost of the worst Nash equilibrium to the social optimum [16]. The PoA has been studied for a large variety of games, including routing, network design with cost-sharing, job scheduling, network creation etc. (see [17, 9, 12, 2, 10, 13]). While the PoA takes a worst-case view regarding the resulting Nash equilibrium, the *Price of Stability (PoS)* [4, 8] takes the reverse view, and considers the ratio of the lowest-cost Nash equilibrium to the social optimum. In fact, for many natural games the Price of Anarchy may be very large while the Price of Stability is quite low. For example, in job scheduling on unrelated machines the Price of Anarchy is unbounded while there always is a Nash equilibrium which is socially optimal (and hence the Price of Stability is 1). Another example is fair cost-sharing games with n players, where the Price of Anarchy is $\Theta(n)$ and the Price of Stability is $\Theta(\log n)$.

In cases where there are both high and low cost Nash equilibria, a central authority could hope to "move" behavior from a high cost equilibrium to a low cost one by running a public service advertising campaign promoting the better behavior. If indeed *everyone* follows the given advice, then we reach the desired equilibrium. This has motivated much of the research on price of stability [4]. The starting point of our work, however, is the realization that it is quite a lot to assume that everyone will follow any given piece of advice, even if the behavior is optimal if everyone else follows it as well. A more realistic assumption is that we might advertise a certain recommended behavior and then some fraction $0 < \alpha < 1$ of players will actually go along. In this paper, we consider the question of what can be done in such a scenario, where a public authority would like to encourage behavior to move from a bad equilibrium to a good one, but can only count on a fraction of players acting in the desired manner. Is affecting a small constant fraction of players sufficient to cause the rest of the players to head towards a low-cost equilibrium, or on the other hand

is even a small constant fraction *not* paying attention enough to cause the whole thing to unravel? What we show is that certain well-studied games are quite resilient to this setting, resulting in low-cost equilibria for any constant $\alpha > 0$; some have a threshold property (e.g., producing good equilibria for $\alpha > 1/2$ but poor equilibria for $\alpha < 1/2$) and some perform poorly for all constant $\alpha < 1$.

As a motivating example consider a traffic control setting. Suppose that currently we have reached some equilibrium which has a high average latency. The authorities would like to change route selections in order to improve the social welfare. One possible way of changing a driver's preferred route is by posting advertisements (e.g., road signs) suggesting an alternative (for example: *The best route to reach the bridge is ...*). Some drivers would follow the advertisements (road signs) while others would ignore them and maybe even try to minimize their driving time given that only some drivers follow the advertisement and switch routes. We would like to assume that eventually every driver would act to minimize driving time conditioned on the behavior of the others (and thus the advertisements have a only limited time effect). The hope is that by having some drivers switch routes, we would converge eventually to a better overall equilibrium. A similar motivation can be given to the fair cost-sharing games which we discuss in this paper. The cost of an edge in this case can be viewed as the maintenance cost of that road segment, and the player's cost can be viewed as paying road tolls to cover the maintenance cost of the road segments the player uses.

To model this type of scenario we introduce the following *advertisement model*. In this model, the authority first suggests to each player a proposed action, and each player accepts the proposal with some (constant) probability α. The players that accept the new action are called the *receptive* players (since they are receptive to the advertising campaign), and they stay with the new action while the *non-receptive* players move to some arbitrary Nash equilibrium for themseleves (given that the receptive players' actions are fixed). Then all players follow a best response dynamics and converge to some pure Nash equilibrium for the entire game (we will only consider games where the best response dynamics is guaranteed to converge to a pure Nash equilibrium).

Our Results: We consider three classes of games: *fair cost-sharing* games, *job scheduling on unrelated machine* games, and *party affiliation* games. For simplicity assume we advertise the socially optimal solution (although all our results apply if we start with any approximation of the optimal solution, including the best Nash equilibrium, and the guarantees degrade proportionally to the approximation factor). Our main results are the following:

- For fair cost sharing games (Section 3) we show that for any α, the expected cost of the new equilibrium is $O((\log n)/\alpha)$ from the optimal solution. Recall that the price if stability is $\Theta(\log n)$ for this game, so the difference in guarantees is only $O(1/\alpha)$, whereas the Price of Anarchy is $\Theta(n)$. Our positive result regarding fair cost-sharing extends to the case where we add to the cost model a linear latency function which depends on the load observed on the edges; the proof for this case involves analyzing a related "shadow" game.[1]

- For job scheduling on unrelated machines (Section 4) we show a negative result, showing that for $m = n$ machines, even if we pick $n - 2$ of the n jobs, there is still a possibility of reaching a pure Nash equilibrium whose cost is unbounded compared to the social optimum. We also show for two machines and n jobs, that if we pick $n/2 - 1$ jobs, there is still a possibility of reaching a pure Nash equilibrium whose cost is unbounded compared to the social optimum. We complement those results showing that there is always a set of $n - n/m$ jobs, such that any pure Nash equilibrium we reach is socially optimal.

- For party affiliation games (Section 5) where players have degree $\omega(\log n)$ we show a threshold property: any value of $\alpha < 1/2$ is not sufficient to improve the equilibrium beyond the $\Omega(n^2)$ price of anarchy, while any value $\alpha > 1/2$ is sufficient to produce behavior within an $O(1)$ factor of optimal. In the case of players with low degree, we show that there is a set of players of size βn, for some constant $1 > \beta > 1/2$, such that if they switch their action, the dynamics will result in a good equilibrium.

Related Work: It is worthwhile to compare and contrast our model with that of Value of Altruism [18] and Strong Price of Anarchy [3]. The Value of Altruism [18] assumes that the authorities control some fraction of the players (or flow) and this part never behaves in a strategic way. In contrast we assume that the receptive players return (eventually) to play strategically, and hence the dynamics always converge to a pure Nash equilibrium. The k-Strong Price of Anarchy [3] focuses on those equilibria such that no subset of at most k players can deviate and *all* strictly benefit; thus it is like a model of stability when players can intelligently form coalitions. In contrast, we consider players that are more myopic in the usual Nash sense, except some fraction are willing to give the advertised behavior a try. The two solution concepts are incomparable in their final guarantees. For job scheduling on two unrelated machines, the 2-Strong Price of Anarchy is $O(n)$ [3], while in our model we show that even if we pick $n/2 - 1$ jobs, the ratio of the cost of the equilibrium

[1] We remark that in both cases, after modifying a constant fraction of the players' actions the social cost can still be quite high, so the result can not be derived by the standard potential-function argument.

produced to OPT can still be unbounded. On the other hand, one can also show a reverse example, of a high cost strong equilibrium [3] where modifying the action of any single job would lead to an optimal outcome.

Charikar et al. [5] consider fair cost sharing for the case that the graph is *undirected* and all players have a common sink, and show that good equilibria can be reached by a process in which players enter one at a time and then undergo best-response dynamics. However, for directed graphs, it is easy to construct simple examples where such a process fails and players reach an equilibrium that is $\Omega(n)$ from optimal. This further motivates our work as it shows that very bad equilibria can be reached by natural dynamics in natural games.

2 The Model

A game is denoted by a tuple $\mathcal{G} =< N, (\mathcal{S}_i), (\texttt{cost}_i) >$ 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., $s_{-i} = (s_1, ..., s_{i-1}, s_{i-1}, ..., s_n)$. The cost function of player i maps a joint action $s \in \mathcal{S}$ to a real non-negative number, i.e., $\texttt{cost}_i : \mathcal{S} \rightarrow \mathbb{R}^+$. Every game has associated a social cost function $\texttt{cost} : \mathcal{S} \rightarrow \mathbb{R}^+$ that maps a joint action to a non-negative real value. In the cases discussed in this paper the social cost is a simple function of the costs of the players. For example, we discuss the sum, i.e., $\texttt{cost}(s) = \sum_{i=1}^{n} \texttt{cost}_i(s)$, and the maximum, 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 is $\mathbf{OPT}(\mathcal{G}) = \min_{s \in \mathcal{S}} \texttt{cost}(s)$. We sometimes overload notation and use $\mathbf{OPT}$ for a joint action s that achieves cost $\mathbf{OPT}(\mathcal{G})$.

Given a joint action s, the *Best Response (BR)* of player i is the set of actions $BR_i(s)$ that minimizes its cost, given the other players actions s_{-i}, i.e., $BR_i(s_{-i}) = \arg\min_{a \in \mathcal{S}_i} \texttt{cost}_i(a, s_{-i})$.

A joint action $s \in \mathcal{S}$ is a *pure Nash Equilibrium (NE)* if no player $i \in N$ can benefit from unilaterally deviating to another action, namely, every player is playing a best response action in s, i.e., $s_i \in BR_i(s_{-i})$ for every $i \in N$. A best response dynamics is a process in which at each time step, some player that is not playing a best response switches its action to a best response action, given the current joint action. In this paper we will concentrate on games in which any best response dynamics converges to a pure Nash equilibrium (which are equivalent to the class of ordinal potential games [15]).

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

The main model we introduce and study in this paper is one that we call the *advertising model*. In this model, the authority first suggests to each player an alternative action, and each player accepts the proposed action with some (constant) probability. The players that accept the new action are called receptive players, and they stay with the new action while the non-receptive players move to an arbitrary Nash equilibrium for themselves (given that the receptive players actions are fixed). Then all players follow an arbitrary best response dynamics and converge to some pure Nash equilibrium (we will only consider games where the best response dynamics is guaranteed to converge to a pure Nash equilibrium). We define this model formally in Figure 1.

Figure 1 Advertising model

Input: Game $\mathcal{G}$, parameter $\alpha > 0$.
Initially players are playing some joint action $s^{ini} \in \mathcal{S}$.

1. We use an advertising campaign and propose an action to each player. Formally, let $s^{ad} = \texttt{advertise}(\mathcal{G})$ be the advertised behavior. (Note: s^{ad} is selected only based on the description of the game $\mathcal{G}$ and independent of the initial actions s^{ini}.)

2. Each player i independently decides to follow or not to follow the proposed action s_i^{ad} by flipping a coin of bias α. Let R be set of players who decide to follow the proposal - we will call them the *receptive players*. Each player $i \in R$ now switches to playing s_i^{ad}.

3. The non-receptive players (players in $N \setminus R$) settle on a Nash equilibrium s^{nr} for themselves, given that the receptive players play s^{ad}. Namely, for each player $j \in N \setminus R$ we have that $s_j^{nr} \in BR_j(s_{-j}^{nr}; s_R^{ad})$. The equilibrium s^{nr} for players in $N \setminus R$ might be *adversarially selected*. Let $s^{med} = (s^{nr}; s_R^{ad})$ be the behavior at this point.

4. The players in R stop following the advertising campaign, and the entire set of players N undergoes a best response dynamics until a Nash equilibrium s^f for the whole game $\mathcal{G}$ is reached. (We will discuss only games where the best response dynamics is guaranteed to converge to a pure Nash equilibrium.)

When we refer later to an *advertising strategy* we mean a joint action $s^{ad} = \texttt{advertise}(\mathcal{G})$. Given a set of receptive players R, there is a set $U(s^{ad}, R)$ which includes all the equilibria s^f that the dynamics could reach having

an advertisement s^{ad} and a set of receptive players R. The *expected cost of the final equilibrium* given s^{ad} is $E_R[\max_{s^f \in U(s^{ad},R)} \texttt{cost}(s^f)]$. When we say that *for game $\mathcal{G}$ there exists a strategy for the advertising model which has an expected cost of the final equilibrium at most X* we mean that there exists a joint action s^{ad} for $\mathcal{G}$ such that $E_R[\max_{s^f \in U(s^{ad},R)} \texttt{cost}(s^f)] \leq X$.

Another natural model we study in this paper is the *viral marketing model*, where the only difference is that the set R is not random, but specifically selected with the proposed actions. Formally, we select R and s_i^{ad} for $i \in R$ based only on the description of the game $\mathcal{G}$ but independently of the initial joint action s^{ini}. (Each player $i \in R$ is assumed to switch to action s_i^{ad}.) Otherwise the viral model is identical to the advertising model.

The notation and definitions for the specific games addressed in the paper are provided in Sections 3, 4 and 5, where the appropriate games are studied.

3 Cost Sharing Games

In this section we consider fair cost sharing games defined as follows. We are given a graph $G = (V, E)$, which can be directed or undirected, where each edge $e \in E$ has a nonnegative cost $c_e \geq 0$. 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. In an outcome of the game, each player i chooses a single path $P_i \in \mathcal{S}_i$. A cost-sharing method assigns nonnegative cost shares to the players, as a function of the set of players that choose a path that contains the edge e. The social cost of an outcome $s = (P_1, ..., P_n)$ is defined to be $\texttt{cost}(P_1, ..., P_n) = \sum_{e \in \cup_i P_i} c_e$. Given a vector of players' strategies $s = (P_1, \ldots, P_n)$, let x_e be the number of agents whose strategy contains edge e. In the fair cost sharing game the cost to agent i is $\texttt{cost}_i(s) = \sum_{e \in P_i} \frac{c_e}{x_e}$ and the goal of each agent is to connect its terminals with minimum total cost.

It is well known that fair cost sharing games are potential games [15, 4] and the price of anarchy in these games is $\Theta(n)$ while the price of stability is $H(n)$ [4], where $H(n) = \sum_{i=1}^{n} 1/i = \Theta(\log n)$. We show in following that in the fair cost sharing game there exists a strategy for the advertising model producing an equilibria whose expected cost is at most an $O((1/\alpha) \log n)$ factor away from the cost of the optimal solution. Thus we get significant benefit from advertising in these games.

Before presenting the proof of our main results we first give two useful lemmas. The first one is well known characterization of the potential function $\Phi(S) = \sum_{e \in E} \sum_{x=1}^{x_e} f_e(x)$ of these games [15, 4], where $f_e(x)$ is cost felt by each user when there are x users on edge e.

LEMMA 3.1. *In the fair cost sharing game the for any joint action $s \in \mathcal{S}$ we have:* $\texttt{cost}(s) \leq \Phi(s) \leq H(n) \cdot \texttt{cost}(s)$.

The second lemma states the following useful property of a binomial random variable.

LEMMA 3.2. *Let X be a binomial random variable distributed $Bi(n, p)$. Then* $\mathbf{E}_X \left[\frac{c}{X+1}\right] = O\left(\frac{c}{p \cdot n}\right)$.

Proof: See Appendix. ∎

We start with our main result concerning fair cost-sharing games.

THEOREM 3.1. *For the fair cost sharing game there exists a strategy for the advertising model which has an expected cost of the final equilibrium at most $O((1/\alpha)(\log n) \cdot \texttt{cost}(\mathbf{OPT}))$.*

Proof: Fix some optimal solution $\mathbf{OPT}$. Let $s^{ad} = \mathbf{OPT}$, so the advertising strategy will be to tell each player i to use his path P_i^{OPT} in $\mathbf{OPT}$. Let R be the set of receptive players and $\bar{R} = N \setminus R$. For every edge e, let n_e^{opt} denote the number of people who use edge e in $\mathbf{OPT}$, and let us decompose this quantity into the number $n_{e,R}^{\text{opt}}$ of those in set R and the number $n_{e,\bar{R}}^{\text{opt}}$ of those not in R (both of which are random variables); so $n_e^{\text{opt}} = n_{e,R}^{\text{opt}} + n_{e,\bar{R}}^{\text{opt}}$.

We start by bounding the expected worst-case cost of the behavior s^{med} produced at the end of step three : that is, $\mathbf{E}_R[\max_{s^{med}=(s^{nr},s_R^{ad})} \texttt{cost}(s^{med})]$, where the max is taken over all behaviors s^{nr} that are Nash equilibria for players in $N \setminus R$ given that the behavior of players in R is s_R^{ad}. We will call this $\mathbf{E}[\texttt{cost}(s^{med})]$ for short. For every edge e let B_e denote the cost of edge e for a non-receptive player $i \in \bar{R}$ under the joint action s^{med}, and let A_e denote the cost of edge e for a receptive player $i \in R$ under the joint action s^{med}. Let $X_e = c_e/(1 + n_{e,R}^{\text{opt}})$ and let $X'_e = c_e/n_{e,R}^{\text{opt}}$. We clearly have $B_e \leq X_e$ and $A_e \leq X'_e$. So, for any player $i \notin R$ we have $\texttt{cost}_i(P_i^{OPT}, s_{-i}^{med}) \leq \sum_{e \in P_i^{OPT}} X_e$ and for any player $i \in R$ we have $\texttt{cost}_i(P_i^{OPT}, s_{-i}^{med}) \leq \sum_{e \in P_i^{OPT}} X'_e$. Since s^{med} is an equilibrium for the non-receptive players we have $s_i^{med} \in BR_i(s_{-i}^{med})$ for $i \in \bar{R}$. This implies the following upper bound on the expected total cost at the end of step three :

$$\mathbf{E}[\texttt{cost}(s^{med})] \leq \mathbf{E}[\sum_{i \notin R} \sum_{e \in P_i^{OPT}} X_e] + \mathbf{E}[\sum_{i \in R} \sum_{e \in P_i^{OPT}} X'_e].$$

Rearranging the sum over players to be a sum over edges we get:

$$\mathbf{E}[\texttt{cost}(s^{med})] \leq \mathbf{E}[\sum_e X_e \cdot n_{e,\bar{R}}^{\text{opt}}] + \mathbf{E}[\sum_e X'_e \cdot n_{e,R}^{\text{opt}}].$$

Note that $X'_e \leq 2X_e$ when $n_{e,R}^{\text{opt}} > 0$. This implies that

$$\mathbf{E}[\texttt{cost}(s^{med})] \leq 3\mathbf{E}[\sum_e X_e \cdot n_e^{\text{opt}}] = 3 \sum_e \mathbf{E}[X_e] \cdot n_e^{\text{opt}}.$$

So, to complete the proof we have to analyze $\mathbf{E}[X_e]$. Lemma 3.2 implies that $\mathbf{E}[X_e]$ is $O\left(c_e/(\alpha \cdot n_e^{\text{opt}})\right)$ for $n_e^{\text{opt}} \geq 1$. This implies that the expected cost at the end of step three satisfies:

$$\begin{aligned}\mathbf{E}[\texttt{cost}(s^{med})] &= O\left((1/\alpha)\sum_{e\in\mathbf{OPT}} c_e\right) \\ &= O((1/\alpha)\ \texttt{cost}(\mathbf{OPT})).\end{aligned}$$

The above equation together with Lemma 3.1 implies that the expected value of the potential function Φ at the end of step three is only an $O((1/\alpha)\log n)$ factor larger than the cost of $\mathbf{OPT}$, i.e.,

$$\mathbf{E}[\Phi(s^{med})] = O((1/\alpha)(\log n)\cdot \texttt{cost}(\mathbf{OPT})).$$

This implies that the expected cost of the final equilibrium at the end of step four , i.e., $\mathbf{E}[\texttt{cost}(s^f)]$, is at most that large, as desired. ■

Note that in fact the proof of Theorem 3.1 can be adapted to prove something stronger.

THEOREM 3.2. *Consider fair cost sharing games and a joint action* F. *Using $s^{ad} = F$ for the advertising model has an expected cost of the final equilibrium at most* $O((1/\alpha)(\log n)\cdot \texttt{cost}(\mathrm{F}))$.

3.1 Extensions A well studied extension of the fair cost sharing game is one where instead of a constant cost c_e, each edge has a cost $c_e(x)$ that is a nondecreasing but concave function of the number of players x using that edge [4]. For example, this can model a buy-at-bulk economy of scale for buying edges that can be used by more players. Notice that the cost of an edge $c_e(x)$ might increase with the number of players using it, but the cost per player $f_e(x) = c_e(x)/x$ decreases if $c_e(x)$ is concave. We can extend our result to this case as well.

THEOREM 3.3. *For the cost sharing game with nondecreasing concave cost functions $c_e(x)$, there exists a strategy for the advertising model which has an expected cost of the final equilibrium at most* $O((1/\alpha)(\log n)\cdot \texttt{cost}(\mathbf{OPT}))$.

Another extension [4] of the fair cost sharing game is one where each edge has both a cost function $c_e(x)$ and a latency function $d_e(x)$, where $c_e(x)$ is the cost of building the edge e for x users which the users will share, while $d_e(x)$ is the delay suffered by each user on edge e if x users are sharing the edge. The goal of each user will be to minimize the sum of his cost and his latency. If we assume that both the cost and latency for each edge depend only on the number of players using that edge, then the total cost felt by each user on the edge is $f_e(x) = c_e(x)/x + d_e(x)$. These games remain potential games [15, 4] (they are particular cases of congestion games). One can prove a lemma similar to Lemma 3.1 relating the cost and the value of the potential function for any given joint action. In particular, for linear delays we have:

LEMMA 3.3. *Consider the cost sharing game with delays where the cost function on edge e is $f_e(x) = c_e/x + l_e \cdot x$. For any joint action $s \in \mathcal{S}$ we have:* $\frac{1}{2}\texttt{cost}(s) \leq \Phi(s) \leq H(n)\cdot \texttt{cost}(s)$.

We show here how we can extend our results to deal with linear delays. The extension is not immediate though since the part of our argument in Theorem 3.1 that says that after step two, every non-receptive player i has a reasonably cheap option to try (namely its path in $\mathbf{OPT}$) is not clear anymore: since the original behavior s^{ini} was *arbitrary* there could exist edges with a much higher number of players on them under the joint action $(s_{\bar{R}}^{ini}, s_R^{ad})$ than in $\mathbf{OPT}$. In order to prove the desired result we instead argue the existence of a related "shadow" game, whose price of anarchy is not too large, and then relate performance of the behaviors as well as the optimum values between the two games.

THEOREM 3.4. *For the cost sharing game with delays where the cost function on edge e is $f_e(x) = c_e/x + l_e \cdot x$ there exists a strategy for the advertising model which has an expected cost of the final equilibrium at most* $O((1/\alpha)(\log n)\cdot \texttt{cost}(\mathbf{OPT}))$.

Proof: Fix some optimal solution $\mathbf{OPT}$ and let $s^{ad} = \texttt{advertise}(\mathcal{G}) = \mathbf{OPT}$. Namely, the advertising strategy will be to tell each player i to use his path P_i^{OPT} in $\mathbf{OPT}$. Let R be the set of receptive players. Let $n_{e,R}^{\text{opt}}$ denote the number of people in R who use edge e and let n_e^{opt} denote the number of people in $\mathbf{OPT}$ who use e.

By assumption, in step three all the users not in R settle on some equilibrium (given s_R^{ad}). Let n_e denote the number of users who are now on edge e. So, $n_e \geq n_{e,R}^{\text{opt}}$. We now define a new game $\mathcal{G}'$ with respect to the users in $\bar{R} = N \setminus R$, which is a congestion game with a linear latency function $h_e(n) = a_e + l_e \cdot n$, where $a_e = c_e/(1+n_e)$. Let $\mathbf{OPT}'$ denote the optimal cost for this game $\mathcal{G}'$.

We first claim that the behavior at the end of step three is also an equilibrium for users in $\bar{R}$ if we use the cost h_e instead of f_e on all edges e. In particular, suppose this was not the case. So, some user i currently using a set S of edges would prefer switching to the set $S - A + B$:

$$\sum_{e\in A}\left(\frac{c_e}{n_e+1} + l_e n_e\right) > \sum_{e\in B}\left(\frac{c_e}{n_e+1} + l_e(n_e+1)\right).$$

However, if we replace $\frac{c_e}{n_e+1}$ with $\frac{c_e}{n_e}$ on the LHS of the above equation, then the gap only gets larger. This means that i is not at equilibrium under f since it can benefit from switching.

Now we use the fact that the new game $\mathcal{G}'$ has a price of anarchy of $5/2$ [6]. So, the total cost in $\mathcal{G}'$ of the behavior of the non-receptive players at the end of Step three is $O(\mathbf{OPT}')$. Note now that the following hold:

(a) The cost of the non-receptive players at the end of Step three using cost functions f_e is at most twice their cost using functions h_e (since adding 1 to the denominator in going from f to h at best reduces the cost by a factor of 2).

(b) $\mathbf{E}(\mathbf{OPT}') = O((1/\alpha)\ \mathtt{cost}(\mathbf{OPT}))$. This is because one option for $\mathbf{OPT}'$ is to use the same paths as in $\mathbf{OPT}$, in which case:

(i) the $l_e \cdot n$ terms are the same as in $\mathbf{OPT}$, and

(ii) since $n_e \geq n^{\text{opt}}_{e,\bar{R}}$ and as we argued in Theorem 3.1 we have $\mathbf{E}[c_e/(1+n^{\text{opt}}_{e,\bar{R}})] = O\left(c_e/(\alpha \cdot n^{\text{opt}}_e)\right)$, this means the expected sum of the a_e terms in $\mathbf{OPT}'$ is $O\left(\sum_{e:n^{\text{opt}}_e>0} n^{\text{opt}}_{e,\bar{R}} \cdot c_e/(\alpha \cdot n^{\text{opt}}_e)\right)$.

These imply that the expected cost under f of the non-receptive players $\bar{R}$ at the end of Step three satisfies:

$$\mathbf{E}[\mathtt{cost}(s^{med}_{\bar{R}})] = O((1/\alpha)\ \mathtt{cost}(\mathbf{OPT})).$$

We now argue that the expected cost in the original game for the receptive players at the end of step three is also $O((1/\alpha)\,\mathtt{cost}(\mathbf{OPT}))$. In particular, the key issue is the latency term, since there could potentially be more non-receptive players on any given edge than in $\mathbf{OPT}$. However, if on a given edge there are more receptive players than non-receptive players in s^{med}, then we lose at most a factor of two compared to the latency cost in $\mathbf{OPT}$; on the other hand if there are more non-receptive players than receptive ones, then we are fine again because we have bounded the cost of the non-receptive players, so we can charge the cost of the receptive players to the cost of the non-receptive players, which we already bounded. So, that the expected cost at the end of step three satisfies:

$$\mathbf{E}[\mathtt{cost}(s^{med})] = O((1/\alpha)\ \mathtt{cost}(\mathbf{OPT})).$$

This together with Lemma 3.3 implies that the expected value of the potential function Φ at the end of step three is at most an $O((1/\alpha)\log n)$ factor larger than the cost of $\mathbf{OPT}$, which implies the cost of any final equilibrium at the end of step four is at most that large, as desired. ∎

Remarks: Note that in all the variants of the cost sharing game studied in this section, the cost of the final equilibrium reached is $O((1/\alpha)\log n)$ from the optimal cost while the price of stability is $\Theta(\log n)$ as shown in [4]. This implies that the difference in guarantee is only a factor $O(1/\alpha)$. Second, our proofs do not really require us to use $\mathbf{OPT}$ for $\mathtt{advertise}(\mathcal{G})$, but rather we can start with any solution F. We would then converge to a pure Nash equilibrium whose cost is at most $O((1/\alpha)\log n)$ from $\mathtt{cost}(\mathrm{F})$. For example, if we let F be the best Nash equilibrium, then since the price of stability is $O(\log n)$, the expected cost of the final equilibrium is within $O((1/\alpha)\log^2 n)$ of the optimal cost.

4 Load Balancing Games

In this section we concentrate on load balancing games (see [16]) defined as follows. There are n jobs and m unrelated machines. Each player is associated with a job, so have n players. Every job can impose a load on one of the machines, so for every player j its set of feasible actions is to assign job j to some machine i, $i \in \{1,\ldots,m\}$. Each job j has associated a cost $c_{i,j}$ for running on machine i. 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., $\mathtt{cost}_j(s) = L_{s_j}(s)$. For the social cost we use the *makespan*, which is the load on the most loaded machine, i.e., $\mathtt{cost}(s) = \max_i L_i(s)$. The price of stability in this games is 1, since there is always a pure Nash equilibrium which is also socially optimal [11].

First, we show a strong negative result: as long as there are two players (jobs) that are not re-assigned, the ratio of the cost of the equilibrium produced to OPT can be unbounded. In addition, we show that this negative result holds even in the viral marketing model, and even if we restrict the adversary to choosing equilibria s^{nr} that are reachable via better-response dynamics from the initial state s^{ini} (we call such an adversary a *reasonable adversary*).

THEOREM 4.1. *There is a load balancing game with $m = n$ unrelated machines such that in the viral marketing model, for any set R of at most $n-2$ players, the cost of the final Nash equilibrium might be unbounded with respect to the social optimum. (I.e., for any advertisement joint action s^{ad} and set of players R, $|R| \leq n-2$, there is an equilibrium $s^f \in U(s^{ad}, R)$ s. t. $\mathtt{cost}(s^f) \geq 1$ while $\mathbf{OPT} = \epsilon$.) Moreover, there exist initial joint actions s^{ini} such that this can occur via a reasonable adversary.*

Proof: Consider the following load balancing game $\mathcal{G}$. Job j has cost ϵ on machine j and 1 on any other machine. The social optimum assigns job j to machine j and has cost ϵ. In $\mathcal{G}$, a joint action is a Nash equilibrium if and only if it allocates to each machine a single job. Let an *empty* machine be a machine with no job assigned to it.

Consider an initial joint action s^{ini} which assigns an even job $2k$ to machine $2k-1$ and an odd job $2k-1$ to machine $2k$. Note that this initial assignment s^{ini} is a Nash

equilibrium and has cost 1. Therefore, the ratio of the initial configuration s^{ini} cost and the social optimum is $1/\epsilon$, which is unbounded since ϵ is arbitrary.

Let R be an arbitrary set of $n-2$ jobs and let s^{ad} be the recommendation for players in R. Let j_1 and j_2 be the two players not in R and $i_1 = s^{ini}_{j_1}$ and $i_2 = s^{ini}_{j_2}$ the machines on which they run in s^{ini}, respectively. Given our initial joint action s^{ini} we know that $j_1 \neq i_1$ and $j_2 \neq i_2$. Also, if $j_1 = i_2$ then $j_2 = i_1$ and vice versa. Hence, either $j_1 = i_2$ and $j_2 = i_1$ or $j_1 \neq i_2$ and $j_2 \neq i_1$. We will show that for our s^{ini} for any s^{ad} and any R, even for a reasonable adversary the process can terminate in a final Nash equilibrium s^f which has a cost of at least 1.

If there is a player k in R which s^{ad} assigns to a machine $s^{ad}_k \neq k$, then consider the following dynamics. Let all the jobs except k reach any equilibrium s' for them (say, using a best response dynamics). In s' there is no empty machine, since otherwise some job $j \neq k$ can improve its cost by moving to the empty machine. This implies that s' is a pure Nash equilibrium of the game $\mathcal{G}$ and therefore $s^f = s'$. Since the cost of job k in s^f is 1, we have that $\texttt{cost}(s^f) = 1$, and we are done. Therefore, assume that s^{ad} assigns each job $j \in R$ to its least load machine, i.e., $s^{ad}_j = j$.

We have two remaining cases to analyze. The first case is when $j_1 \neq i_2$ and $j_2 \neq i_1$. In this case s^{ad} assigns job i_1 to machine i_1 and job i_2 to machine i_2, and each machine has a load of $1+\epsilon$, while machines j_1 and j_2 are empty. Jobs j_1 and j_2 can then undergo a better-response process and select the following equilibrium: job j_1 selects machine j_2 and job j_2 selects machine j_1, having a cost of 1 for each. Since this is an equilibrium we also reached s^f which has a cost of 1. In the second case $j_1 = i_2$ and $j_2 = i_1$. In this case after s^{ad} each machine has a single job, and hence we are at an equilibrium which has a cost of 1. ∎

For the case of two machines and an arbitrary number of jobs we can derive the following result.

THEOREM 4.2. *There is a load balancing game with two unrelated machines and n jobs such that in the viral model, for any set R of at most $n/2-1$ players, the cost of the final Nash equilibrium might be unbounded with respect to the social optimum.*

Proof: Assume that n is even, i.e., $n = 2k$. We have k jobs of type I, defined such as their cost on machine 1 is ϵ and on machine 2 is 1; we also have k jobs of type II, defined such as their cost on machine 1 is 1 and on machine 2 is ϵ. In s^{ini} all the jobs of type I are on machine 2 and all the jobs of type II are on machine 1, which is a Nash equilibrium that has a cost of k (compared to $\mathbf{OPT}$ which has a cost of $k\epsilon$).

Suppose that R includes k_1 jobs of type I and k_2 jobs of type 2, and advertises s^{ad} for them. Consider the following Nash equilibrium s^{med} for the players not in R. We select k_1 jobs from type II and k_2 jobs of type I and pair them with the jobs in R where in each pair one job is of type I and the other of type II. (Since R is strictly less than half the jobs, i.e., $|R| \leq k-1$, we can do it, and there will be at least one type I job and one type II job remaining.) For each pair of matched jobs $j_1 \in R$ and $j_2 \notin R$ we set the action of job j_2 to be the opposite machine of j_1, i.e., $s^{med}_{j_2} = 3 - s^{ad}_{j_1}$. This implies that the pair's contribution on each machine is identical. Therefore at s^{med} we have that the load on both machines is identical, and hence it is a Nash equilibrium. Since there is a pair of jobs that did not move from their action in s^{ini} the cost is at least 1, while the optimal cost is $k\epsilon$. ∎

We can show that the result above is almost tight in the sense that there is always a set of players R of size $(1-/m)n$, such that in the viral marketing model the final equilibrium is always optimal.

THEOREM 4.3. *For any load balancing game there exists a strategy for the viral model with a set R of at most $(1-1/m)n$ players such that any final equilibrium s^f has optimal cost.*

Proof: Let $s^{ad} = \mathbf{OPT}$ for some $\mathbf{OPT}$. Let k be the machine with most players in s^{ad}, i.e., $k = \arg\max_i |B_i(s^{ad})|$. Let $R = N \setminus B_k(s^{ad})$, and note that since at least n/m players are in $B_k(s^{ad})$, then $|R| \leq n - n/m$. Note that given that s^{ad} is performed by R, any job $j \in B_k(s^{ad})$ has a best response whose cost is at most $\texttt{cost}_j(\mathbf{OPT})$, since in s^{med} only jobs in $B_k(s^{ad})$ would select machine k. Therefore at the end of step three we have that $\texttt{cost}(s^{med}) = \texttt{cost}(\mathbf{OPT})$. Since the best response dynamics does not increase the cost and eventually converge to a pure Nash equilibrium [11], we have $\texttt{cost}(s^f) = \texttt{cost}(\mathbf{OPT})$, as desired. ∎

5 Consensus Games, Cut Games, and Party Affiliation Games

In this section we consider three related classes of games, played by users who are viewed as vertices in a connected, undirected simple graph $G = (N, E)$ with n vertices, where $N = \{1, ..., n\}$. We will first describe the most general game of *party affiliation* and then discuss the special cases of *consensus games* and *cut games*.

In *party affiliation games* [7] the set of edges E is partitioned into positive and negative edges, denoted by PE and NE respectively. Each player i has two actions r or b, i.e., $\mathcal{S}_i = \{r, b\}$. A player has cost 1 for each incident positive edge on which he disagrees with his neighbor, and cost 1 for each negative edge on which he agrees with his neighbor, i.e., $\texttt{cost}_i(s) = \sum_{(i,j)\in PE} \mathbf{I}_{(\mathbf{s_i} \neq \mathbf{s_j})} + \sum_{(i,j)\in NE} \mathbf{I}_{(\mathbf{s_i} = \mathbf{s_j})}$. The overall social cost is the sum of the costs of all the users, plus 1, i.e., $\texttt{cost}(s) = 1 +$

$\sum_{i \in N} \texttt{cost}_i(s)$.[2] It is straight forward to show an exact potential function [15] for the party affiliation game, simply let the potential be $\Phi(s) = (\texttt{cost}(s) - 1)/2$. Also, in any party affiliation game the social optimum is a Nash equilibrium, thus the Price of Stability is 1.

Consensus games are a special case of party affiliation games where all the edges in G are positive edges, i.e., $NE = \emptyset$. The two social optimal solutions in a consensus game are "all blue" and "all red", both of which are also a Nash equilibrium. On the other hand, for an even number of players, let G be the clique K_n with a perfect matching removed, e.g., $E = \{(i,j) : j \neq i\} - \{(2i, 2i-1) : n/2 \geq i \geq 1\}$. Consider the joint action s in which even players play r and odd players play b, i.e., $s_{2k} = r$ and $s_{2k+1} = b$. This is a Nash equilibrium, since each player has exactly half its neighbors the same color and exactly half of the opposite color. This results in a social cost of $\Omega(n^2)$ and thus the Price of Anarchy for consensus games is $\Omega(n^2)$.

Cut games (see [7]) are a special case of party affiliation games where all the edges in G are negative edges, i.e., $PE = \emptyset$. Thus cut games have the opposite objective from consensus games. As mentioned above, in a cut game the optimal solution is a Nash equilibrium and so the Price of Stability is 1. However, the problem of *finding* an (approximately) optimal solution is the Min-UnCut problem [1] for which the best efficient approximation algorithm known has approximation ratio $O(\sqrt{\log n})$. As with consensus games, the Price of Anarchy for cut-games can be as bad as $\Omega(n^2)$. For instance, if G is the complete bipartite graph $K_{n/2,n/2}$, then coloring half the nodes on the left and half the nodes on the right blue, and coloring half the nodes on the left and half the nodes on the right red, is a Nash equilibrium with cost $\Omega(n^2)$ (and yet the optimal solution has cost 1 since the graph is bipartite).

We first show that if all nodes have degree $\omega(\log n)$, then in the advertising model all these games have a sharp threshold at $\alpha = 1/2$: any constant $\alpha > 1/2$ is sufficient to produce an optimal or near-optimal solution, and yet there exist families of graphs for which any constant $\alpha < 1/2$ yields a solution of cost as nearly bad as possible (a factor $\Omega(n^2)$ worse than optimal). We begin with the simpler case of consensus games.

THEOREM 5.1. *For consensus games in which each node has degree $\omega(\log n)$ there is a sharp threshold at $\alpha = 1/2$ in the advertising model: for any constant $\alpha > 1/2$ there exists a strategy such that with high probability the final equilibrium is the optimal solution (ratio of 1), and yet for any constant $\alpha < 1/2$ there exist graphs such that for any advertising strategy with high probability the final equilibrium will be a factor $\Omega(n^2)$ worse than optimal.*

[2]The "+1" is just to ensure the cost is nonzero so that all ratios are well-defined.

Proof: For the upper bound, the advertising strategy is simply to tell all nodes to become color red, i.e., $s^{ad} = (r, \ldots, r)$. By Hoeffding bounds, each node with degree at least $\log n/(\alpha - 1/2)^2$ has more than half of its neighbors in set R with probability at least $1 - 1/n^2$. Therefore, by the union bound all nodes have this property with probability at least $1 - 1/n$, and so with high probability at the end of step three all nodes are red, i.e., $s^{med} = (r, \ldots, r)$, which is optimal.

For the lower bound, let $\gamma = 1/2 - \alpha$ and consider a graph consisting of two cliques of size $n/2$, where each vertex has $\gamma n/8$ neighbors in the other clique. Suppose initially we have one clique red and the other clique blue. Since γ is a positive constant, for sufficiently large n we have that with high probability each node has at most a $1/2 - \gamma/2$ fraction of its neighbors in set S. However, since each node initially has only a $\gamma/4$ fraction of its neighbors of the other color, this will not be sufficient to cause any of the nodes not in S to change color in step three. Therefore, in step three, all nodes in S will simply revert to their original color and we again have $\Omega(n^2)$ badly-colored edges. ∎

The key to the upper bound above is that with high probability the set R satisfies the property that every vertex not in R has more than half its neighbors in R. For cut games and more generally party-affiliation games, we will need a bit more (in particular because **OPT** no longer necessarily has zero cost for every player). Specifically let us say that a set S is a β-dominating set if every vertex not in S has more than a $1/2 + \beta$ fraction of its neighbors in S. Hoeffding bounds imply that in the advertising model with $\alpha > 1/2 + 2\beta$, with high probability the set of receptive players is a β-dominating set, so long as all nodes have degree $\omega(\log n)$:

LEMMA 5.1. *For any party affiliation game in which each node has degree $\omega(\log n)$, for any constant $\alpha > 1/2 + 2\beta$ with probability $1 - o(1)$ the set of receptive players is a β-dominating set in the advertising model.*

We now show the following property of β-dominating sets in party affiliation games.

LEMMA 5.2. *For party-affiliation games, if the set R following the advertising strategy is a β-dominating set, then we can produce a solution within an $O(1/\beta)$ factor of optimal.*

Proof: We argue by considering two kinds of nodes: those with less than a β fraction of their incident edges incurring a cost (of one) in **OPT** (call those "low-cost" nodes), and those with more than a β fraction of their incident edges incurring a cost in **OPT** (call those "high-cost" nodes). The advertising strategy is to tell nodes to behave according to

$\mathbf{OPT}$, i.e., $s^{ad} = \mathbf{OPT}$.[3] Since R is a β-dominating set, each low-cost node will change in step three to its color in $\mathbf{OPT}$ (because that color minimizes its cost with a majority of its neighbors). For the high-cost nodes, we may not produce the desired behavior; however, no matter how the high-cost nodes behave, they cannot incur a cost that is more than a $1/\beta$ factor worse than their cost in $\mathbf{OPT}$ (by definition of "high cost"). Therefore, the total cost by the end of step three is at most $(1+1/\beta)\mathbf{OPT}$. Finally, the cost can only improve via the best-response process in Step four . ∎

For $\alpha > 1/2$, by setting $\beta = (\alpha - 1/2)/2$, Lemma 5.1 and Lemma 5.2 imply that the cost of the final equilibrium is within an $O(1)$ factor of optimal. For $\alpha < 1/2$, the $\Omega(n^2)$ bound from Theorem 5.1 still applies; thus exhibiting is a sharp threshold at $\alpha = 1/2$ in the advertising model.

THEOREM 5.2. *For party-affiliation games in which each node has degree $\omega(\log n)$, for any constant $\alpha > 1/2$ there exists a strategy for the advertising model such that with high probability the final equilibrium has cost $O(\mathbf{OPT})$. Moreover, for any constant $\alpha < 1/2$ there exist party-affiliation games such that for any advertising strategy with high probability the final equilibrium will be a factor $\Omega(n^2)$ worse than optimal.*

Unfortunately, in the low-degree case, no value $\alpha < 1$ is sufficient to achieve a cost that is even within a $o(n)$ factor of $\mathbf{OPT}$, because the graph G could consist of a collection of constant-sized components, and with high probability a constant fraction of these components would have no member in R. On the other hand, note that if the set of receptive players R constitutes a β-dominating set, then by Lemma 5.2 we produce a solution within an $O(1/\beta)$ factor of optimal. This immediately translates to a result in the viral marketing model, where we can select the set of receptive players. The following lemma constructs a β-dominating set in an arbitrary graph.

LEMMA 5.3. *For any graph G and $\beta < 1/6$, there is a β-dominating set R of size at most $(31/32)n$.*

Proof: Let S' be a random set of vertices, where each vertex in S' with probability $1/2 + \beta$, and let $S'' = \{v : v \text{ does not have at least } \lceil d_v(1/2 + \beta)\rceil \text{ neighbors in } S'\}$, where d_v is the degree of v. Consider $R = S' \cup S''$. The set R is by construction a β-dominating set, so we simply have to argue about size. In particular, we show that any vertex v has some constant probability of not being chosen in $S' \cup S''$.

With probability $1/2 - \beta$ the vertex v is not in S'. In order for v not to be in S'' it needs at least $\lceil d(1/2 + \beta)\rceil$ neighbors in S'. Consider the following event: the first 3 neighbors of v are in S' (probability $(1/2+\beta)^3$) and from the remaining $d-3$ neighbors of v at least $\lfloor (d-3)(1/2+\beta)\rfloor$ are in S' (this has probability at least $1/2$ since this is less than or equal to the median of the distribution). If this event occurs, since $\beta < 1/6$, then v has at least $\lfloor (d_v-3)(1/2+\beta)\rfloor + 3 = \lfloor d_v(1/2+\beta) + 3(1/2-\beta)\rfloor \geq \lfloor d_v(1/2+\beta) + 1\rfloor \geq \lceil d_v(1/2+\beta)\rceil$ neighbors in S' and therefore it is not in S''. Since v is not in S' nor S'' it is not in R. The probability of the event is at least $1/2(1/2-\beta)(1/2+\beta)^3 > 1/32$ for $\beta < 1/6$. So, the expected size of R is at most $(31/32)n$ and therefore a set of at most that size must exist. ∎

For the case of consensus games we can show improved bounds (since we only need a strict majority for each node). Call a set S a *strict dominating set* if for every node v not in S the strict majority of its neighbors are in S, i.e. $\lceil \frac{d_v+1}{2} \rceil$. For consensus games it is sufficient to have a strict dominating set R to guarantee an optimal solution, since then we are guaranteed that all the players will switch to the color of the majority of their neighbors. The following lemma derives bounds for strict dominating sets.

LEMMA 5.4. *For any graph G in which each node has odd degree, there is a strict dominating set R of size at most $n/2$, and for general degrees, there is a strict dominating set R of size at most $(23/27)n$.*

Proof: First, consider the case that each node has odd degree. Consider an arbitrary Nash equilibrium of the *cut game* in the given graph G, i.e., all the edges are negative edges, and let R be the minority color class in this equilibrium (hence, the size of R is at most $n/2$). Since it is a Nash equilibrium of the cut game, each vertex not in R has a majority of its neighbors in the set R (a strict majority since all degrees are odd).

For the case of general degrees we use a randomized argument similar to Lemma 5.3. Let S' be a random set of vertices, where each vertex is in S' with probability $2/3$, and let $S'' = \{v : v \text{ does not have a strict majority of neighbors in } S'\}$. We will then let $R = S' \cup S''$. Set R by construction has the property that all vertices not in R have a strict majority of their neighbors in R so we simply have to argue about size. In particular, any vertex v has at least a probability $1/3$ of not being chosen in S' and probability at least $4/9$ having a strict majority of its neighbors in S' (the worst case is when v has degree 2). So, the expected size of R is at most $(23/27)n$ and therefore a set of at most that size must exist. ∎

We can now deduce the following theorem for the viral marketing model.

[3] In general, this advertising strategy can be computationally hard to compute. For cut-games, however, at the loss of an extra $O(\sqrt{\log n})$ factor we can instead use the Min-UnCut approximation algorithm of [1], and the rest of the argument proceeds in the same way.

THEOREM 5.3. *In the viral marketing model: (1) For any party affiliation game there is a set R of $(31/32)n$ players that guarantees a solution within an $O(1)$ factor of optimal. (2) For consensus games, there is a set R of at most $n/2$ players if all degrees are odd, or a set R of at most $(23/27)n$ players for general degrees, that guarantees an optimal solution.*

6 Conclusions

In this paper we consider the question of to what extent can a "public advertising campaign" cause behavior to switch from a bad equilibrium to a good one even if only a fraction of people actually follow the given advice, and do so only temporarily. Unlike the notion of price of stability we do not assume everyone takes the advice, and unlike the notion of value of altruism we assume everyone in the end will act in their own interest. We show that for some natural games (cost-sharing), it is enough to induce a small fraction to behave well in order to reach a good equilibrium, whereas for others (load balancing) one needs to reach nearly the entire population, and yet others (party affiliation) have a threshold property. Thus, we provide an interesting metric along which games can differ: how much "effort" (e.g., advertising dollars) a central authority might need to incur in order to induce good behavior in them.

While we have described our results using advice that consists of the global optimum behavior, in all cases this can be replaced with advising the best equilibrium at an extra cost of the price of stability. In a sense (viewed in reverse) this can be thought of as asking "how stable is the price of stability": starting from a good equilibrium, can a small shock to the system produce a bad state from which natural dynamics could not recover?

References

[1] A. Agarwal, M. Charikar, K. Makarychev, and Y. Makarychev. O($\sqrt{logn}$) approximation algorithms for Min UnCut, Min 2CNF deletion, and directed cut problems. In *Proc. 37th STOC*, 2005.

[2] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *Proc. 17th ACM-SIAM SODA*, 2006.

[3] N. Andelman, M. Feldman, and Y. Mansour. Strong Price of Anarchy. In *Proc. 18th ACM-SIAM SODA*, 2007.

[4] 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.

[5] M. Charikar, H. Karloff, C. Mathieu, J. Naor, and M. Saks. Online multicast with egalitarian cost sharing. In *Proc. 20th SPAA*, 2008.

[6] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proc. 37th STOC*, 2005.

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

[8] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):236 – 259, 2004.

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

[10] E. D. Demaine, M.T. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. The price of anarchy in network creation games. In *Proc. 17th SODA*, 2006.

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

[12] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. 22nd PODC*, 2003.

[13] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2):226 – 239, 2005.

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

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

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

[17] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236 – 259, 2002.

[18] Y. Sharma and D. P. Williamson. Stackelberg thresholds in network routing games or the value of altruism. In *ACM Conference on Electronic Commerce*, 2007.

A Additional Proofs

We prove here a useful property of a binomial random variable.

Lemma 3.2 Let X be a binomial random variable distributed $Bi(n,p)$. Then $\mathbf{E}_X\left[\frac{c}{X+1}\right] = O\left(\frac{c}{p\cdot n}\right)$.

Proof: We have

$$\begin{aligned}
\mathbf{E}_X\left[\frac{n+1}{X+1}\right] &= \sum_{i=0}^{n}\frac{n+1}{i+1}\binom{n}{i}p^i(1-p)^{n-i} \\
&= \frac{1}{p}\sum_{i=0}^{n}\binom{n+1}{i+1}p^{i+1}(1-p)^{n-i} \\
&= \frac{1}{p}-\frac{(1-p)^{n+1}}{p}
\end{aligned}$$

This implies that

$$\begin{aligned}
\mathbf{E}_X\left[\frac{c}{X+1}\right] &= \frac{c}{p\cdot(n+1)}-\frac{c(1-p)^{n+1}}{p\cdot(n+1)} \\
&= \frac{c}{n+1}\left[\frac{1}{p}-\frac{(1-p)^{n+1}}{p}\right] \\
&= O\left(\frac{c}{p\cdot n}\right). \quad \blacksquare
\end{aligned}$$