On Allocations with Negative Externalities

Sayan Bhattacharya, Janardhan Kulkarni,
Kamesh Munagala, and Xiaoming Xu
Department of Computer Science
Duke University
Durham NC 27708-0129.
{bsayan,kulkarni,kamesh,xiaoming}@cs.duke.edu

Abstract. We consider the problem of a monopolist seller who wants to sell some items to a set of buyers. The buyers are strategic, unit-demand, and connected by a social network. Furthermore, the utility of a buyer is a decreasing function of the number of neighbors who do not own the item. In other words, they exhibit negative externalities, deriving utility from being *unique* in their purchases. In this model, any fixed setting of the price induces a sub-game on the buyers. We show that it is an exact potential game which admits multiple pure Nash Equilibria. A natural problem is to compute those pure Nash equilibria that raise the most and least revenue for the seller. These correspond respectively to the most optimistic and most pessimistic revenues that can be raised.

We show that the revenues of *both* the best and worst equilibria are hard to approximate within sub-polynomial factors. Given this hardness, we consider a relaxed notion of pricing, where the price for the same item can vary within a constant factor for different buyers. We show a 4-approximation to the pessimistic revenue when the prices are relaxed by a factor of 4. The interesting aspect of this algorithm is that it uses a linear programming relaxation that only encodes part of the strategic behavior of the buyers in its constraints, and rounds this relaxation to obtain a starting configuration for performing relaxed Nash dynamics. Finally, for the maximum revenue Nash equilibrium, we show a 2-approximation for bipartite graphs (without price relaxation), and complement this result by showing that the problem is NP-Hard even on trees.

1 Introduction

This paper considers pricing and allocations over a social network, when buyers derive utility from being *unique* in their purchase. Such *negative externalities* arise in several consumer goods where buyers derive value from "showing off" the product to friends lacking it. Consider the following example. For many years, the DVD publication industry has utilized the so called "double-dipping" policy, a term for releasing multiple versions of the same movie on discs. A quick search for the movie "The Matrix" shows besides the original, there are "The Matrix Revisited", "The Matrix: Platinum Limited Edition Collector's Set", and "The Ultimate Matrix Collection". It is often the case that these editions have the same core content (in this case, the movie), while they differ in the some "unique" different material that is packed with the disc, for instance, sound tracks of the music, toy character of the figures in the movie, etc. And it is often the case that the price discrepancy between these versions outweighs the "real value" or intrinsic value the extra material provides. An incentive as observed in [22] is that "..the extra 20 or 30 bucks ... is discreet enough to display without geek alarms flashing and whirling anytime I have friends and/or new people over". The same marketing policy exists in many other different industries, book publishing, expensive electronic gadgets to name a few.

The model for negative externalities we study in the paper is simple: The buyers are unit-demand and connected by a social network that we model as a graph; the edges in the graph represent friendships. There are two types of items, each with unlimited supply - a "cheap" item and an "expensive" item. Each user has an intrinsic value for each type of item; however, buyers of the expensive item also derive additional *extrinsic* utility from friends who only possess the cheaper item. We assume the edges in the graph are weighted, so that the extrinsic utility is the sum of the weights of edges leading to friends possessing the cheaper item.

In this model, a monopolist wishes to price the items to maximize revenue. Any fixed setting of prices induces a sub-game on the buyers where they decide which item to purchase in order to maximize their individual utility. If a buyer buys the cheaper item, then she gets its intrinsic valuation; else if she buys the more expensive item, she gets the sum of its intrinsic and extrinsic valuations. The whole process can be viewed as a strategic game occurring in two rounds: The seller commits to the two prices in advance, and then each buyer simultaneously decides which item to purchase. We term this the Pricing Game, and investigate two natural questions for a fixed setting of prices:

- What is the *pessimistic* Nash equilibrium, *i.e.*, one that raises minimum revenue? This gives a guarantee on the revenue the seller raises regardless of the behavior of the buyers; a risk-averse seller will choose prices at which this revenue is as large as possible.
- What is the *best* Nash equilibrium, *i.e.*, one that raises largest revenue? This will give the most optimistic view of the buyers' behavior, and is appropriate when the seller can recommend which item to buy via targeted advertising.

Given efficient algorithms for the above problems, the seller can treat them as subroutines while iterating over all possible prices. This will help her set the prices in such a way that maximizes the pessimistic (resp. optimistic) revenue.

Sequential Pricing: Though our model is motivated by negative externalities, the same model arises in an entirely different context that is well-studied in economics. In *sequential pricing*, buyers derive positive utility from neighbors who bought the item earlier in time, and strategically decide when to buy the item to maximize their utility. The goal of the seller is to decide the prices to set for each stage, with later stages having higher prices, so that the resulting subgame among the buyers raises large revenue. Such a model of externality is motivated by buyers deciding to wait on a purchase if a lot of her friends buy now, since waiting will lead to more product feedback and hence raise utility. This is a well-studied problem: Exact and approximation algorithms are known for models with no network effects [11, 7, 9, 13, 21], non-atomic buyers [5, 2], or with non-strategizing buyers [14, 1, 3]. Our formulation exactly models two-stage sequential pricing for a single item with strategic, atomic buyers on a network: Buyers derive both an intrinsic utility from the item as well as linear utility from neighbors who have bought the item in the previous stage. To map this to our problem, simply note that the buyers in the second stage of sequential pricing correspond to buyers of the more expensive item, and those in the first stage, to the cheaper item. As we show next, the presence of atomic, strategizing buyers on a network causes the problem to become structurally different from previously considered versions.

1.1 Our Results

We first show that the PRICING GAME is an exact potential game, and hence admits to a pure Nash Equilibrium (Theorem 1). As a corollary, the Nash dynamics converges in poly-time when the edge-weights are polynomially bounded integers. It is straightforward to check that not only are there multiple pure Nash equilibria in this game, but also they have vastly different revenue properties for the seller. This motivates us to focus on computing equilibria with certain optimality properties. In particular, we show the following results:

- For the Pessimistic (minimum) Nash equilibrium problem, we show that its revenue is NP-HARD to approximate within a factor of $O(n^{1/3-\epsilon})$, where n is the number of nodes in the social network (Theorem 2). Interestingly, our hardness results hold even when the buyers derive significantly large intrinsic utility compared to the externalities.
- In view of the above lower bound, we focus on a δ-relaxed notion of Nash equilibria,[1] where the seller posts different prices to different buyers of the expensive item, but these prices are within a factor of δ of each other. We give an algorithm (Theorem 6) for computing a 4-relaxed NE that is a 4-approximation to the pessimistic revenue. The algorithm uses a linear program relaxation to bound the revenue of the pessimistic NE. It is interesting to note that this linear program only encodes the constraint that the buyers do not deviate from buying the cheaper item to the more expensive item, allowing the buyers of the more expensive item to deviate. Based on the ideas from vertex cover, we present a rounding scheme for this LP, and use the resulting solution as a starting point for performing relaxed Nash dynamics.
- We show (Theorem 7) that the Best Nash equilibrium problem is also hard to approximate within a factor of $O(n^{1/3-\epsilon})$ by reduction from the independent set problem. However, in contrast with independent set, our problem is NP-Hard even on trees (Theorem 8). We finally present a 2-approximation to maximum revenue when the underlying network is bipartite (Theorem 9).

[1] Refer to the discussion in the beginning of Section 3.1 for more details.

We emphasize that all the three hardness results (Theorems 2, 7, 8) hold even if the edge-weights are small integers so that Nash dynamics converges quickly to *some* Nash equilibrium. In contrast, both the algorithms presented in this paper - a 4 approximation with relaxed prices to the revenue of the Pessimistic Nash Equilibrium, and a 2 approximation to the revenue of the Best Nash Equilibrium on bipartite graphs - run in polynomial time for arbitrary edge-weights.

Remark. All the missing proofs appear in the Appendix.

1.2 Related Work

Viral Marketing: Marketing strategies in social networks have been studied extensively, starting with the seminal paper by Kempe *et. al.* [17]. Several recent papers [2, 14, 1, 3, 12] consider pricing and auction design in social networks when buyers exhibit *positive* network externalities, meaning they derive positive utility from neighbors possessing the product. Our main contribution is to show that the structure induced by uniqueness and negative externalities is very different from positive externalities. In particular, we show that computing Nash Equilibria with desirable revenue properties ends up being hard to approximate, and we need to consider relaxed notions of equilibria for positive results.

Potential Games: Our problem is a special case of potential (or congestion) games [19], which always admit to pure Nash equilibria. Typically, these games have been studied from two perspectives - how bad can the welfare of the resulting equilibria be compared to optimal social welfare (*price of anarchy*) [18, 20]; and how quickly does the Nash dynamics converge [10, 8, 4]. The literature on price of anarchy aims to find worst case loss in efficiency due to strategizing across all game instances. In contrast, our focus is on *computing* the best or worst equilibria for a *specific* input network. It is relatively easy to check that the *worst-case* revenue of these equilibria across all networks can be $\Omega(n)$ factor off from the optimal revenue that can be raised with non-strategizing buyers; and furthermore, the revenues of best and worst equilibria can be separated by a factor of $\Omega(n)$. For general potential games, computing a pure Nash equilibrium is PLS-complete [10], and most literature in this domain have tried to circumvent this hardness by considering the notion of *approximate* Nash equilibrium [6]. Our focus is not on PLS-completeness; in fact, our problem is interesting (and hard) even in the regime where the Nash dynamics converges in polynomial time.

2 Notations and Preliminaries

There is a *cheap* item and an *expensive* item, each of them being available in unlimited supply. The buyers are unit demand. Consider an undirected graph $G = (V, E)$ with a weight function over the edges $w : E \to \mathbb{R}^+$. Each node $i \in V$ denotes a buyer, there is an edge $(i, j) \in E$ if buyers i, j are friends, and the weight of their friendship is given by the quantity $w(i, j)$. The price of the cheap (resp. expensive) item is set at p_1 (resp. p_2). Naturally, the price of the cheap item should be less than that of the expensive one, that is, $p_1 \leq p_2$. The pricing is *uniform* in the sense that the same item is offered at the same price to all the nodes. Say that a node is *black* (resp. *white*) if she buys the cheap (resp. expensive) item, while paying an amount p_1 (resp. p_2).

Definition 1. *The* externality $Ext(i)$ *of node i is the total weight of the edges it shares with black nodes. Thus, we have* $Ext(i) = \sum_{(i,j)\in E:\, j \text{ is black}} w(i,j)$.

Definition 2. *The* weighted degree $\mathcal{D}(i)$ *of node i is the total weights of the edges incident to it. Thus, we have* $\mathcal{D}(i) = \sum_{(i,j)\in E} w(i,j)$.

Any buyer purchasing the cheap (resp. expensive) item gets an *intrinsic* valuation of a (resp. b), where $0 \leq a \leq b$, and a, b are publicly known and fixed in advance. In addition, any buyer purchasing the expensive item gets an *extrinsic* valuation that increases as more of her friends purchase the cheap item, and her total valuation equals the sum of intrinsic and extrinsic valuations. To be more specific, the valuation of every black node is equal to a; whereas the valuation of a white node i is given by the expression $b + Ext(i)$.

The buyers have *quasi-linear* utilities. If a node i is black, then her utility is $U(i) = a - p_1$. Else if the node i is white, then her utility is $U(i) = b + Ext(i) - p_2$. And if node i does not purchase any item, then she gets *zero* utility. We assume $p_1 \leq a$ and $p_2 \leq b$. These two inequalities guarantee that every buyer purchases the item, or equivalently, every node is colored either black or white.

In this paper, we consider the setting where we are given the values of p_1, p_2, a and b. This induces a normal form game between the buyers, and we term it as the Pricing Game. The strategy of each node consists of choosing whether to be colored black or white. A coloring of the nodes defines a strategy profile, and in a Nash equilibrium, the color chosen by each node is a best response to the colors chosen by other nodes. Thus, node i is colored black iff:

$$a - p_1 \geq b + Ext(i) - p_2 \Rightarrow Ext(i) \leq p_2 - p_1 + a - b = \Delta \text{ (say)}$$

Similarly, node i is colored white if and only if $Ext(i) \geq \Delta$.[2]

The game can therefore be summarized as follows.

Pricing Game. In any pure Nash equilibrium, every node is colored either black or white, and each black (resp. white) node has an externality of at most (resp. at least) Δ. Let $B_{\mathcal{C}}$ (resp. $W_{\mathcal{C}} = V \setminus B_{\mathcal{C}}$) denote the set of black (resp. white) nodes under coloring $\mathcal{C}$. Each black (resp. white) node mays a payment of p_1 (resp. p_2). Thus:

$$\text{Seller's Revenue } = p_1 \times |B_{\mathcal{C}}| + p_2 \times |W_{\mathcal{C}}| = (\Delta + b - a) \times |W_{\mathcal{C}}| + p_1 \times |V_{\mathcal{C}}|$$

[2] If $\Delta < 0$, then it is a dominant strategy for each buyer to purchase the more expensive item, and the pure NE degenerates to the case where all nodes are colored white. Throughout the rest of the paper, we will consider the generic setting where $\Delta \geq 0$.

In the above equation, the second equality holds since $p_2 = p_1 + \Delta + b - a$.

Our first theorem shows that the Pricing Game is an exact potential game. Associate each coloring $\mathcal{C}$ with a potential $\phi(\mathcal{C})$, given by the total weight of black-black edges plus Δ times the number of white nodes. In terms of notations,

$$\phi(\mathcal{C}) = \Delta \times |W_{\mathcal{C}}| + \sum_{(i,j)\in E: i,j\in B_{\mathcal{C}}} w(i,j)$$

Theorem 1. *The* Pricing Game *is an exact potential game with potential function $\phi()$, and hence it admits a pure Nash equilibrium. If a, b, p_1, p_2 and all the edge weights $w(i,j)$ take integral values in the range $\{1, \ldots, \mu\}$, then the Nash dynamics converges within $\Theta(\mu|V|^2)$ steps.*

In this paper, we are interested in estimating the maximum (resp. minimum) revenue the seller may obtain from any pure Nash equilibrium. Towards this end, we define the following problems.

Problem 1 (BNE: Best Nash Equilibrium). Given an instance of the Pricing Game, find a pure Nash equilibrium that generates the maximum revenue.

Problem 2 (PNE: Pessimistic Nash Equilibrium). Given an instance of the Pricing Game, find a pure Nash equilibrium that generates the minimum revenue.

Our focus is *not* on studying Nash dynamics. Instead we ask the question: *Even if* the edge-weights are small integers so that Nash dynamics converges in polynomial time (Theorem 1), how hard is it to compute a pure Nash equilibrium with specific revenue properties? All our hardness results (Theorems 2, 7, 8) hold under this very natural scenario of small integer edge-weights. Our hardness results also hold when intrinsic valuations are large compared to each individual externality (the reductions set $a = b = 1$), and the seller wants to obtain non-negligible revenue from each item (the reductions set $p_1 = 1$). However, both the positive results in this paper - a 4 approximation with relaxed prices to the revenue of the Pessimistic Nash Equilibrium (Theorem 6), and a 2 approximation to the revenue of the Best Nash Equilibrium on bipartite graphs (Theorem 9) - do not make *any* assumption on the input, and hold for arbitrary edge-weights.

3 The Pessimistic Nash Equilibrium

First, we derive a strong inapproximability result for the Pessimistic Nash Equilibrium (PNE) problem by reducing it from Minimum Maximal Independent Set (MMIS). In the MMIS problem, we are given a graph, and the objective to find an independent set of minimum size such that every node in the graph is adjacent to at least one node in the independent set. Kann [16] shows that it is NP-hard to get a poly-time $O(n^{1-\epsilon})$ approximation for MMIS [16].

Theorem 2. *It is* NP-hard *to compute in poly-time a pure Nash equilibrium whose revenue is a $O(n^{1/3-\epsilon})$ approximation to the minimum revenue, where n is the number of nodes in the social network.*

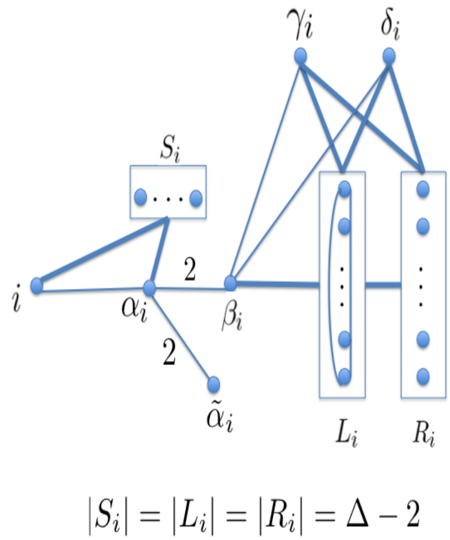

Fig. 1. The gadget associated with a node $i \in V$ in the inapproximability reduction for PNE problem. Each rectangle denotes a set of nodes.

Proof. Given an instance $G = (V, E)$ of the minimum maximal independent set problem, we reduce it to an instance of PNE as outlined below. The resulting weighted graph will be denoted by $\mathcal{G}$ (see Figure 1).

- Set $a = b = 1$, $p_1 = 1$, and $p_2 = |V|^2$, implying that $\Delta = (p_2 - p_1 + a - b) = \Theta(|V|^2)$. Consequently, in any pure Nash equilibrium of the PNE instance we construct, every black (resp. white) node should have an externality at most (resp. at least) Δ. Further, every black node makes a payment of 1, whereas every white node contributes $\Theta(|V|^2)$ towards total revenue.
- Start with $G = (V, E)$ and assign a weight of 3 to each edge $(i, j) \in E$.
- For all nodes $i \in V$:
  - Create a new set of nodes T_i of cardinality $(3\Delta - 1)$.
  - Partition T_i into four disjoint subsets as L_i, R_i, S_i and $\{\alpha_i, \tilde{\alpha}_i, \beta_i, \gamma_i, \delta_i\}$, where $|L_i| = |R_i| = |S_i| = \Delta - 2$.
  - Create a cycle connecting the nodes of L_i.
  - For all $u \in L_i, v \in R_i$, create an edge (u, v). For all $u \in L_i \cup R_i$, create the edges $(u, \gamma_i), (u, \delta_i)$. For all $u \in L_i$, create an edge (u, β_i). For all $u \in S_i$, create the edges $(u, \alpha_i), (u, i)$.
  - Create the edges $(\beta_i, \gamma_i), (\beta_i, \delta_i), (\alpha_i, \beta_i), (i, \alpha_i)$ and $(\alpha_i, \tilde{\alpha}_i)$.
  - Assign a weight of 2 to each of the edges (α_i, β_i) and $(\alpha_i, \tilde{\alpha}_i)$. Every other newly created edge gets a weight of 1.

Since all nodes in $S_i \cup \{\tilde{\alpha}_i\}$ have a weighted degree of 2, and $\Delta > 2$, we have:

Observation 1. *For all $i \in V$, the set of nodes $S_i \cup \{\tilde{\alpha}_i\}$ must be colored black in any pure Nash equilibrium.*

The next three lemmas will be crucial in deriving the lower bound.

Lemma 3. *Consider any pure Nash equilibrium in graph $\mathcal{G}$. The nodes of the set V that are colored black form a maximal independent set in $G = (V, E)$.*

Lemma 4. *Consider the pure Nash equilibrium in graph $\mathcal{G}$ with minimum revenue. If a node $i \in V$ is colored white, then the revenue from the set T_i is $\Theta(\Delta)$.*

Lemma 5. *Consider any pure Nash equilibrium in graph $\mathcal{G}$. If a node $i \in V$ is colored black, then the revenue from the set T_i is $\Theta(\Delta^2)$.*

Let B^* (resp. $W^* = V \setminus B^*$) denote the set of black (resp. white) nodes from V in *the pure Nash equilibrium of graph $\mathcal{G}$ that minimizes revenue.* By Lemmas 4, 5, the revenue of the coloring is given by the expression $|B^*| \times \Theta(\Delta^2) + |W^*| \times \Theta(\Delta)$. Since $\Delta = \Theta(|V|^2)$ and $|W^*|$ can be at most $|V|$, the first term dominates the second one, and revenue equals $|B^*| \times \Theta(|V|^4)$. Since the coloring *minimizes* the quantity $|B^*| \times \Theta(|V|^4)$, Lemma 3 implies that, upto a constant factor, the set B^* is a minimum maximal independent set in the graph $G = (V, E)$.

Now, let B (resp. $W = V \setminus B$) denote the set of black (resp. white) nodes from V in any pure Nash equilibrium of $\mathcal{G}$. By Lemma 5, the revenue obtained is at least $|B| \times \Theta(\Delta^2) = |B| \times \Theta(|V|^4)$. By Lemma 3, the set B is also a maximal independent set in the graph $G = (V, E)$. Hence an approximation for the PNE instance will imply an approximation for the MMIS of graph $G = (V, E)$.

Finally, the graph $\mathcal{G}$ contains $\Theta(|V| \times \Delta) = \Theta(|V|^3) = n$ nodes. As a consequence, a $O(n^{1/3-\epsilon})$ approximation for the PNE instance will translate into a $O(|V|^{1-3\epsilon})$ approximation for the minimum maximal independent set of the graph $G = (V, E)$. This completes the reduction. □

3.1 Nash Equilibrium with Relaxed Prices

At first glance, it seems that a possible way to circumvent the hardness result (Theorem 2) is to focus on the relaxed notion of *approximate* Nash equilibrium, which is defined as a coloring where any node on flipping color can improve her utility by at most a constant factor. Unfortunately, such an approach is *not* going to work, for the following reason. In the proof of Theorem 2, we have $p_1 = a$, that is, the cheap item is offered at a price equal to its intrinsic valuation. Hence every black node gets *zero* utility. In this situation, any approximate Nash equilibrium must be an *exact* Nash equilibrium. As a result, the hardness lower bound carries over to the seemingly relaxed notion of approximate Nash equilibrium.

On a positive note, we provide a 4-approximation algorithm for the PNE problem when the constraint of uniform pricing is relaxed slightly, allowing the seller to offer different prices to different buyers purchasing the same item, under the condition that these prices are within a constant factor of each other. Our algorithm returns a coloring $\mathcal{C}$ satisfying the two properties described below.

– Every white node has an externality of at least $\Delta = (p_2 - p_1 + a - b)$, and every black node has an externality of at most 4Δ. It is easy to verify that such a coloring is a pure Nash equilibrium under the following *relaxed* pricing scheme: Each black (resp. white) node is offered the price $4p_2$ (resp. p_2) for the expensive item, whereas the cheap item is offered at the same price p_1 to every node.
– The revenue of the coloring $\mathcal{C}$ is *at most* 4 times the revenue of any pure Nash equilibrium where the cheap (resp. expensive) item is offered to all buyers at the same price p_1 (resp. p_2).

The algorithm is described in Figure 2. We give a LP relaxation for the PNE instance, encoding the constraint that every black node has an externality of at most Δ. Surprisingly, the LP does *not* enforce any lower bound on the externalities of the white nodes. Using a simple rounding scheme, we get an integral solution whose revenue is at most 4 times the minimum revenue. In the integral solution, every black node has an externality of at most 2Δ. Next, we employ an iterative improvement process that flips the colors of *bad* (see Definition 3) nodes. We show that the revenue does not increase during the iterative improvement process and that the process converges in polynomial time to a coloring that is a pure Nash equilibrium under the relaxed pricing scheme mentioned above.

We first write down the LP relaxation for the given PNE instance. Recall that the notation $\mathcal{D}(i)$ denotes the weighted degree of node i (Definition 2). Consider any pure Nash equilibrium. For all $i \in V$, if node i is colored white (resp. black), set $x_i = 1$ (resp. $x_i = 0$). For all $(i,j) \in E$, set $y_{ij} = 1$ if edge (i,j) has both its end vertices colored black, otherwise set $y_{ij} = 0$. The revenue is given by the objective value of LP-MIN. Constraint (3) states that if a node has weighted degree less than Δ, then it must be colored black. Constraint (2) ensures $y_{ij} = 1$ if and only if both endpoints of the edge (i,j) are black. Finally, Constraint (1) requires that if a node is black, then she has an externality of at most Δ. Thus, any pure Nash equilibrium is a feasible solution to LP-MIN, and optimal objective value of the LP lower bounds the minimum revenue obtainable from any pure Nash equilibrium.

$$\text{Minimize} \quad (\Delta + b - a)\sum_i x_i + p_1|V| \qquad \text{(LP-MIN)}$$

$$\begin{aligned}
\textstyle\sum_{j:(i,j)\in E} w(i,j)y_{ij} &\leq \Delta \quad && \forall i : \mathcal{D}(i) \geq \Delta && (1)\\
x_i + x_j + y_{ij} &\geq 1 && \forall (i,j) \in E && (2)\\
x_i &= 0 && \forall i : \mathcal{D}(i) < \Delta && (3)\\
x_i &\in [0,1]\ \forall i && && (4)
\end{aligned}$$

Definition 3. *A node is called* bad *if either it is white and has an externality less than Δ, or if it is black and has an externality more than 4Δ.*

Theorem 6. *Given an instance of the PNE problem, Algorithm RELAXED (Figure 2) returns a coloring of the nodes with the following properties.*

RELAXED

INPUT: A graph $G = (V, E)$, and some $\Delta > 0$.
OUTPUT: A coloring of the nodes.

Let $\{x_i^*, y_{ij}^*\}$ denote the optimal solution to LP-MIN.
Rounding: For all nodes $i \in V$:
If $x_i^* \geq 1/4$, then color node i as white; Else color node i as black.
Iterative Improvement:
WHILE there exists a *bad* [a] node $i \in V$: *Flip* the color of node i.

[a] See Definition 3

Fig. 2. Approximation Algorithm for the Pessimistic Nash Equilibrium (PNE) problem with relaxed prices.

- *The algorithm terminates in polynomial time.*
- *The revenue of the returned coloring is a 4-approximation to the optimal objective value.*
- *It returns a coloring where each white node has an externality of at least Δ and each black node has an externality of at most 4Δ.*

Proof. Let $\mathcal{C}$ denote the coloring obtained from the rounding step. Since every white node has $x_i^* \geq 1/4$, the revenue of coloring $\mathcal{C}$ is at most 4 times the optimal objective value of LP-MIN. We now show that the revenue can only decrease during the iterative improvement process, which converges in polynomial time.

Consider a node i which is black under coloring $\mathcal{C}$. Consider any other black node j that is adjacent to node i. It follows that $x_i^*, x_j^* < 1/4$, and by constraint (2) of LP-MIN, we have $y_{ij}^* \geq 1/2$. Now constraint (1) implies that $Ext(i) \leq 2\Delta$. Hence, under coloring $\mathcal{C}$, every black node i has $Ext(i) \leq 2\Delta$.

At any instant of the iterative improvement process, let W (resp. $B = V \setminus W$) denote the set of white (resp. black) nodes. Associate with every node i a variable called $bank(i)$, and define the potential of the system at any instant as $\Phi = \Delta \times |W| + \sum_{i \in V} bank(i)$. Initially, all the bank variables are set to 0. Hence, the potential of the coloring $\mathcal{C}$ is given by Δ times the number of white nodes. As the process unfolds, we adjust the bank variables according to the following rules.

- If a node i is flipped from white to black, then for all black nodes j adjacent to i, increment $bank(j)$ by an amount $w(i, j)$. Furthermore, set $bank(i)$ to 0.
- If a node i is flipped from black to white, then set $bank(i)$ to 0.

First, we observe that when a node i is flipped from white to black, we have:

$$\text{Decrease in Potential} = \Delta - \sum_{j \in B:\, (i,j) \in E} w(i, j) = \Delta - Ext(i) > 0$$

The last inequality holds since any node flipped from white to black has an externality strictly less than Δ. Next, we note that whenever a black node i is flipped to white, then the decrease in potential is $-\Delta + bank(i)$.

The last time instant when node i was colored black, we had $Ext(i) \leq 2\Delta$ and $bank(i)$ was set to 0. At the present instant, when we are flipping node i to white, $Ext(i)$ is at least 4Δ. During this time period, node i has therefore gained an externality of 2Δ. During the same time period, whenever a friend j of node i has been flipped to black, both the externality of node i and the variable $bank(i)$ have increased by the same amount $w(i,j)$. Consequently, at the present time instant, the variable $bank(i)$ has a value at least 2Δ, and when we are flipping node i to white, we are decreasing the potential by at least Δ.

The preceding discussion shows that the iterative improvement process never increases the potential. The highest potential is obtained at the beginning of the process when all bank variables are zero, and consequently, this value is bounded from above by the expression $\Delta \times |V|$. And whenever a node is flipped from black to white, the potential decreases by at least Δ. Since a flip from black to white must occur at least once amongst any sequence of $|V|+1$ consecutive flips, it follows that the process *decreases the potential* by an amount Δ within every $\Theta(|V|)$ steps. As a result, the process converges in $\Theta(|V|^2)$ steps.

Since the bank variables are initially set to zero and the potential decreases as the process unfolds, it implies that the number of white nodes (and hence revenue) in the final coloring must be lower than that of the coloring $\mathcal{C}$ obtained from the rounding step. Recall that the revenue of the coloring $\mathcal{C}$ is at most 4 times the minimum revenue, and the approximation guarantee follows.

The iterative improvement step is repeated till there are no bad nodes (Definition 3), and hence the final part of the theorem can be easily verified. □

4 The Best Nash Equilibrium

This section begins with a strong inapproximability result for the BNE problem.

Theorem 7. *Unless $NP = ZPP$, it is not possible to compute in polynomial time a pure Nash equilibrium whose revenue is a $O(n^{1/3-\epsilon})$ approximation to maximum revenue, where n is the number of nodes in the social network.*

In view of the above hardness result on general graphs, we consider the BNE problem on bipartite graphs. Unfortunately, the next theorem shows that the problem is NP-HARD even on trees (and hence on general bipartite graphs).

Theorem 8. *It is NP-hard to optimally solve the BNE problem on trees.*

On the positive side, we can show a simple 2 approximation for the BNE problem on bipartite graphs. Consider a BNE instance on a bipartite graph with partite-sets L, R and set of edges $E \subseteq L \times R$. Suppose that $|L| \geq |R|$, and the weighted degree (Definition 2) of every node in the graph is at least Δ. In this case, it is easy to check that coloring all the nodes in L (resp. R) as white (resp. black) gives a 2 approximation to the BNE problem. This idea can be extended to graphs where nodes can have arbitrary weighted degrees.

Theorem 9. *There is an efficient algorithm that gives* 2 *approximation to the BNE problem on bipartite graphs.*

5 Conclusion

We considered the setting of monopoly pricing over a social network in presence of negative externalities. All of our algorithmic results, both for BNE and PNE, can easily be adapted to the scenario where different buyers have different intrinsic valuations for the same item. In addition, the algorithms work even if the extrinsic valuation for the expensive item increases linearly with the externality *only upto a certain threshold*, and then remains fixed at the threshold value.

It will be interesting to extend our results to pricing with more than two items, and investigate other solution concepts such as mixed Nash equilibrium. We also note that our work assumed a perfect information model where the seller knows the valuations of the buyers, and we leave open the question of studying the effects of imperfect information under a Bayesian setting.

Acknowledgments: We thank Vincent Conitzer and Aneesh Sharma for several helpful discussions. This research was supported by Munagala's Alfred P. Sloan Research Fellowship, a gift from Cisco, and by NSF via grants CCF-0745761 and CCF-1008065.

References

[1] H. Akhlaghpour, M. Ghodsi, N. Haghpanah, V. Mirrokni, H. Mahini, and A. Nikzad. Optimal iterative pricing over social networks (extended abstract). In *WINE*, pages 415–423, 2010.

[2] N. Anari, S. Ehsani, M. Ghodsi, N. Haghpanah, N. Immorlica, H. Mahini, and V. Mirrokni. Equilibrium pricing with positive externalities (extended abstract). In *WINE*, pages 424–431, 2010.

[3] D. Arthur, R. Motwani, A. Sharma, and Y. Xu. Pricing strategies for viral marketing on social networks. In *WINE*, pages 101–112, 2009.

[4] B. Awerbuch, Y. Azar, A. Epstein, V. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *EC*, pages 264–273, 2008.

[5] B. Bensaid and J. Lesne. Dynamic Monopoly Pricing with Network Externalities. *International Journal of Industrial Organization*, 14(6):837–855, 1996.

[6] A. Bhalgat, T. Chakraborty, and S. Khanna. Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In *EC*, pages 73–82, 2010.

[7] J. Bulow. Durable-goods Monopolists. *The Journal of Political Economy*, 90(2):314–332, 1982.

[8] S. Chien and A. Sinclair. Convergence to approximate nash equilibria in congestion games. In *SODA*, pages 169–178, 2007.

[9] R. Coase. Durability and Monopoly. *Journal of Law and Economics*, pages 143–149, 1972.

[10] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *STOC*, pages 604–612, 2004.

[11] F. Gul, H. Sonnenschein, and R. Wilson. Foundations of dynamic monopoly and the coase conjecture. *Journal of Economic Theory*, 39(1):155–190, June 1986.
[12] N. Haghpanah, N. Immorlica, V. S. Mirrokni, and K. Munagala. Optimal auctions with positive network externalities. In *EC*, pages 11–20, 2011.
[13] O. Hart and J. Tirole. Contract Renegotiation and Coasian Dynamics. *The Review of Economic Studies*, 55(4):509–540, 1988.
[14] J. Hartline, V. Mirrokni, and M. Sundararajan. Optimal Marketing Strategies over Social Networks. In *WWW*, pages 189–198, 2008.
[15] J. Håstad. Clique is hard to approximate within n$^{\text{1-epsilon}}$. In *FOCS*, pages 627–636, 1996.
[16] V. Kann. Polynomially bounded minimization problems that are hard to approximate. *Nordic Journal of Computing*, 1(3):317–331, 1994.
[17] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the Spread of Influence Through a Social Network. In *ACM SIGKDD*, pages 137–146, 2003.
[18] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
[19] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
[20] T. Roughgarden and É. Tardos. How bad is selfish routing? In *FOCS*, pages 93–102, 2000.
[21] J. Thépot. A Direct Proof of the Coase Conjecture. *Journal of Mathematical Economics*, 29(1):57–66, 1998.
[22] http://www.washingtonpost.com/wp-dyn/articles/A55414-2004Dec10.html, Dec. 2004.

A Appendix

A.1 Proof of Theorem 1

Proof. Given any coloring $\mathcal{C} : V \rightarrow \{\text{black, white}\}$, let $U_{\mathcal{C}}(i)$ and $Ext_{\mathcal{C}}(i)$ respectively denote the utility of node i and the externality of node i. Further, let $B_{\mathcal{C}}$ (resp. $W_{\mathcal{C}} = V \setminus B_{\mathcal{C}}$) denote the set of black (resp. white) nodes under coloring $\mathcal{C}$. The potential function associated with coloring $\mathcal{C}$ is given by

$$\phi(\mathcal{C}) = \Delta \times |W_{\mathcal{C}}| + \sum_{(i,j)\in E: i,j \in B_{\mathcal{C}}} w(i,j)$$

Suppose we take a coloring $\mathcal{C}$ and flip the color of any one node i to get a new coloring $\mathcal{C}'$, while the colors of all other nodes remain the same. We show $U_{\mathcal{C}'}(i) - U_{\mathcal{C}}(i) = -(\phi(\mathcal{C}') - \phi(\mathcal{C}))$. This will prove that the Pricing Game is an exact potential game, guaranteeing the existence of a pure Nash equilibrium. We note that $Ext_{\mathcal{C}}(i) = Ext_{\mathcal{C}'}(i)$, and consider two possible scenarios.

Case 1. Node i is black under coloring $\mathcal{C}$, and white under coloring $\mathcal{C}'$. Here,

$$U_{\mathcal{C}'}(i) - U_{\mathcal{C}}(i) = b + Ext_{\mathcal{C}'}(i) - p_2 - (a - p_1) = Ext_{\mathcal{C}}(i) - \Delta$$

As node i changes her color from black to white, the total weight of black-black edges reduces by $Ext_{\mathcal{C}}(i)$ and the number of white nodes increases by 1. Thus,

$\phi(\mathcal{C}') - \phi(\mathcal{C}) = -Ext_{\mathcal{C}}(i) + \Delta = -(U_{\mathcal{C}'}(i) - U_{\mathcal{C}}(i))$.

Case 2. Node i is white under coloring $\mathcal{C}$, and black under coloring $\mathcal{C}'$. Here, $U_{\mathcal{C}'}(i) - U_{\mathcal{C}}(i) = -(Ext_{\mathcal{C}}(i) - \Delta)$. As node i changes her color from white to black, the total weight of black-black edges increases by $Ext_{\mathcal{C}}(i)$ and the number of white nodes decrease by 1. Thus, $\phi(\mathcal{C}') - \phi(\mathcal{C}) = Ext_{\mathcal{C}}(i) - \Delta$. Hence the PRICING GAME is an exact potential game.

Finally, it is easy to verify that the potential value is at most $\Theta(\mu n^2)$. Since every step in the Nash dynamics decreases this potential by an integral amount, the process must converge within $\Theta(\mu n^2)$ steps. □

A.2 Proof of Lemma 3

Proof. Consider an edge $(i, j) \in E$. Node i (resp. node j) is friend with the set S_i (resp. S_j). By Observation 1, both the nodes i, j gets an externality of at least $(\Delta - 2)$. The edge (i, j) has weight 3. Thus, both the nodes i, j cannot be colored black simultaneously, or else each of them will get an externality of $(\Delta + 1)$. The black nodes in V, therefore, form an independent set.

Now consider a node $i \in V$. If none of her friends in the set V are colored black, then she gets an externality of at most $\Delta - 1$ (when node α_i is colored black). Hence node i must be colored black. In other words, every white node in V must be adjacent to at least one black node in V, and as a result, the black nodes in V form a maximal independent set in $G = (V, E)$. □

A.3 Proof of Lemma 4

Proof. Color all the nodes in the set $\{\alpha_i, \tilde{\alpha}_i\} \cup L_i \cup R_i \cup S_i$ as black, and all three nodes in the set $\{\beta_i, \gamma_i, \delta_i\}$ as white. Verify that each black (resp. white) node in the set T_i has an externality of at most (resp. at least) Δ, and total revenue from T_i is $O(\Delta)$. Since the set T_i contains $\Omega(\Delta)$ nodes, minimum revenue from T_i is at least $\Omega(\Delta)$. The lemma follows. □

A.4 Proof of Lemma 5

Proof. Since the set T_i contains $O(\Delta)$ nodes, even if we color all of them white, revenue from the set is at most $O(\Delta^2)$. In the following paragraphs, we show that when node i is colored black, revenue from the T_i is at least $\Omega(\Delta^2)$.

Both the nodes γ_i, δ_i are connected to $2\Delta - 4$ nodes in the set $L_i \cup R_i$ via unit-weight edges. If either node γ_i or node δ_i is colored black, at least $\Delta - 4$ nodes from the set $L_i \cup R_i$ will have to be colored white, and the revenue will be $\Omega(\Delta^2)$. For the rest of the proof, assume that both the nodes γ_i, δ_i are colored white. Hence the maximum possible externality of any node in the set R_i can be at most $\Delta - 2$ (apart from the set L_i, all her friends are white). It follows that all nodes in the set R_i must be colored black.

By Observation 1, all nodes in the set $S_i \cup \{\tilde{\alpha}_i\}$ must be colored black, and node i is black by assumption. Consequently, node α_i has an externality of at least $\Delta + 1$ (she is friend with the set $\{i\} \cup S_i \cup \{\tilde{\alpha}_i\}$), and she must be colored white. Next, node β_i can potentially have an externality of at most $\Delta - 2$ (apart from the set L_i, all her friends are white), and she must be colored black.

Finally, consider the set L_i. Each node in L_i has an externality of at least $\Delta - 1$ (she is friend with the set $\beta_i \cup R_i$). Recall that nodes in L_i form a cycle. No three consecutive nodes on the cycle can be colored black; for if three nodes $x, y, z \in L_i$, connected by edges (x, y) and (y, z), are colored black, then the middle node y will get an externality of $\Delta + 1$. Therefore, at least $(\Delta - 2)/3$ nodes from the set L_i must be colored white, resulting in a revenue of $\Omega(\Delta^2)$. □

A.5 Proof of Theorem 7

Proof. The reduction is from the Maximum Independent Set problem. It is known [15] that unless $NP = ZPP$, the Maximum Independent Set problem cannot be approximated in polynomial time within a factor of $O(n^{1-\epsilon})$.

Given an instance of the maximum independent set problem $G = (V, E)$, we construct an instance of BNE as outlined below. The resulting graph will be denoted by $\mathcal{G}$.

- Set $a = b = 1$, $p_1 = 1$, and $p_2 = |V|^2$, implying that $\Delta = \Theta(|V|^2)$. Equivalently, in any pure Nash equilibrium of the BNE instance we construct, every black (resp. white) node should have an externality at most (resp. at least) Δ. Further, every black node makes a payment of 1, whereas every white node contributes $\Theta(\Delta) = \Theta(|V|^2)$ towards total revenue.

- Start with the graph $G = (V, E)$. Assign a weight of 2 to each edge $(i, j) \in E$.

- For each node $i \in V$:
  - Create a new set of nodes T_i of cardinality $(3\Delta - 3)$.
  - Partition T_i into three mutually disjoint subsets L_i, R_i and S_i of equal size, that is, $|L_i| = |R_i| = |S_i| = \Delta - 1$.
  - For all $u \in L_i, v \in R_i$, create an edge (u, v) of weight 1.
  - For all $u \in L_i \cup S_i$, create an edge (u, i) of weight 1.

In the graph $\mathcal{G}$, consider any node $i \in V$. Note that all nodes in S_i have a weighted degree of 1, and they must be colored black. Further, if node i is colored black, then we can color the set of nodes R_i as black and the set of nodes L_i as white, obtaining a revenue of $\Theta(\Delta^2)$ from the set T_i. On the other hand, if node i is colored white, then no node in the set $L_i \cup R_i$ can get an externality of Δ. As a consequence, no node in the set $L_i \cup R_i$ can be colored white, and we get a revenue of $\Theta(\Delta)$ from the set T_i.

Let B (resp. W) denote the set of nodes in V that are colored black (resp. white) in any pure Nash equilibrium of $\mathcal{G}$. The preceding discussion indicates that the total revenue is given by the expression $|B| \times \Theta(\Delta^2) + |W| \times \Theta(\Delta)$. Since we set $\Delta = \Theta(|V|^2)$ and $|W|$ can be at most $|V|$, the first term dominates the second one, and revenue equals $|B| \times \Theta(|V|^4)$. Thus, approximating the maximum revenue reduces, upto a constant factor, to approximating the maximum number of nodes in V that can be colored black. To complete the proof, we show in the next paragraph that the black nodes in V form an independent set. Since the graph $\mathcal{G}$ has $|V| \times \Theta(\Delta) = \Theta(|V|^3) = n$ nodes, any $O(n^{1/3-\epsilon})$ approximation for the BNE instance will imply an $O(|V|^{1-3\epsilon})$ approximation for the maximum independent set of the graph $G = (V, E)$.

Consider any two nodes $i, j \in V$ that are friends of each other. Recall that all nodes in the set $S_i \cup S_j$ must be colored black. Since $|S_i| = |S_j| = \Delta - 1$, each of the nodes i, j already has an externality of $\Delta - 1$. Thus, nodes i, j can never be colored black simultaneously, or else each of them will get an externality of $(\Delta + 1)$. In other words, black nodes in V form an independent set in the input graph $G = (V, E)$. □

A.6 Proof of Theorem 8

Proof. Given an instance of the 3SAT problem with n variables $\{x\}$ and m clauses $\{c\}$, we reduce it to an instance of the BNE problem on tress as follows (See Figure 3).

- Set $a = b = 1$, $p_1 = 1$. Set p_2 to be sufficiently large so that Δ is also sufficiently large. In particular, any $p_2 \geq 5$ will work.

- For each variable x and its negation $\neg x$:
  - Create the nodes v_x and $v_{\neg x}$.
  - Create an edge $(v_x, v_{\neg x})$ of weight 2.
  - Create a new set of nodes Λ_x of cardinality $(\Delta - 1)$.
  - Create another new set of nodes $\tilde{\Lambda}_x$ of cardinality $(\Delta - 1)$.
  - For all $u \in \Lambda_x$: Create an edge (u, v_x) of weight 1.
  - For all $u \in \tilde{\Lambda}_x$: Create an edge $(u, v_{\neg x})$ of weight 1.

- For each clause c:
  - Create the nodes q_c, r_c.
  - Create an edge (q_c, r_c) of weight 2.
  - Create two new sets of nodes $\tilde{Q}_c, \tilde{R}_c$, where $|\tilde{Q}_c| = \Delta - 3$, and $|\tilde{R}_c| = \Delta - 1$.
  - For all $u \in \tilde{Q}_c$: Create an edge (u, q_c) of weight 1.
  - For all $u \in \tilde{R}_c$: Create an edge (u, r_c) of weight 1.
  - Create a new set of nodes S_c of size k. Set $k >> m + n$. Let $S_c = \{s_c^1, \ldots, s_c^k\}$.
  - For all $i \in \{1, \ldots, k\}$:
    - Create an edge (r_c, s_c^i) of weight 1.
    - Create a new set of nodes T_c^i of cardinality $\Delta - 1$.
    - For all $u \in T_c^i$: Create an edge (u, s_c^i) of weight 1.

– For all variables x and all clauses c:
  - If the literal x belongs to clause c, create the edge (v_x, q_c) of weight 1. Else if the literal $\neg x$ belongs to clause c, create the edge $(v_{\neg x}, q_c)$ of weight 1.

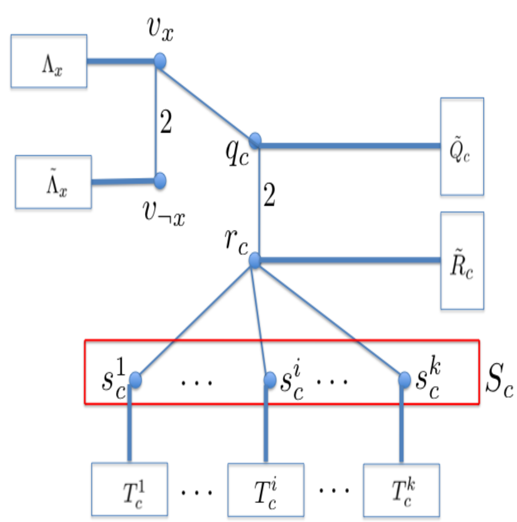

Fig. 3. NP-hardness reduction from 3SAT to BNE on trees. Clause c contains literal x. Each rectangle denotes a set of nodes.

Note that the graph constructed above is a tree. An optimal algorithm for the BNE instance finds a pure Nash equilibrium with maximum number of white nodes. To complete the proof, we show (Lemmas 10, 13) that the BNE instance has a pure Nash equilibrium with at least $k \times m$ white nodes if and only if the 3SAT instance has a satisfying assignment. Recall that m is the number of clauses, n is the number of variables, and $k >> m + n$.

Lemma 10. *If the 3SAT instance has a satisfying assignment, then the BNE instance has a pure Nash equilibrium with at least $k \times m$ white nodes.*

Proof. Note the satisfying assignment for the 3SAT instance, and color the nodes of the BNE instance as follows.

– For all variables x:
  - Color the set of nodes $\Lambda_x \cup \tilde{\Lambda}_x$ as black.

  - In the satisfying assignment, if $x = 1$, color node v_x as black and node $v_{\neg x}$ as white. Else if $x = 0$, color node v_x as white and nodes $v_{\neg x}$ as black.

- For all clauses c:
  - Color the set of nodes $\tilde{Q}_c \cup \tilde{R}_c \cup_{i\in\{1...k\}} T_c^i$ as black.
  - Color node r_c as black, and node q_c as white.
  - Color the set of nodes S_c as white.

Verify that the resulting coloring is a pure Nash equilibrium, that is, every black (resp. white) node has an externality of at most (resp. at least) Δ. Since there are m clauses, and for each clause c, the set of nodes S_c have been colored white, the total number of white nodes is at least $m \times k$. □

The next observation follows from the fact that a node with weighted degree (Definition 2) strictly less than Δ must be colored black in any pure Nash equilibrium.

Observation 2. *Consider any pure Nash equilibrium of the BNE instance. For all variables x, the set of nodes $\Lambda_x \cup \tilde{\Lambda}_x$ must be black. Similarly for all clauses c, the set of nodes $\tilde{Q}_c \cup \tilde{R}_c \cup_{i\in\{1...k\}} T_c^i$ must be black.*

Observation 2 states that certain nodes of the BNE instance must be black in any pure Nash equilibrium. Consequently, all the white nodes must come from the remaining set $\mathcal{V} = \cup_c\{r_c, q_c\} \cup_x \{v_x, v_{\neg x}\} \cup_c S_c$. The first two components, that is $\cup_c\{r_c, q_c\}$ and $\cup_x\{v_x, v_{\neg x}\}$, together contain $2(n+m)$ nodes. Whereas the third component, that is $\cup_c S_c$, contains km nodes. Overall, there are $2(n+m)+km$ nodes in the set $\mathcal{V}$. The next two claims will be crucial in proving Lemma 13.

Claim 11. *Consider a pure Nash equilibrium with at least km white nodes. Then for all clauses c, the node q_c must be colored white.*

Proof. Suppose for some clause c, the node q_c is black. As a result, node r_c gets an externality that is at least $\Delta + 1$ (she is friend with $\{q_c\} \cup \tilde{R}_c$), and she must be white. It follows that each node s_c^i gets an externality strictly less than Δ (her only friends are the set of nodes $\{r_c\} \cup T_c^i$), and all the nodes in S_c are black. The number of white nodes, therefore, is upper bounded by $|\mathcal{V} \setminus S_c| = 2(n+m)+km-k < km$. The last inequality holds since $k >> m+n$, and we get a contradiction. □

Claim 12. *Suppose for all clauses c, the node q_c is colored white. Then for all variables x, one of the nodes $v_x, v_{\neg x}$ is colored white, while the other one is colored black.*

Proof. By Observation 2, the set of nodes $\Lambda_x \cup \tilde{\Lambda}_x$ must be black. Now, if both the nodes v_x and $v_{\neg x}$ are black, then each of them gets an externality of $\Delta + 1$, a contradiction. On the other hand, if both the nodes v_x and $v_{\neg x}$ are white, then each of them gets an externality of $\Delta - 1$ (since for all the clauses c, node q_c is white). Again we reach a contradiction. Thus, the nodes $v_x, v_{\neg x}$ must be assigned different colors. □

Lemma 13. *If the BNE instance has a pure Nash equilibrium with at least $k \times m$ white nodes, then the 3SAT instance has a satisfying assignment.*

Proof. Fix a pure Nash equilibrium with at least km white nodes. A literal denotes a variable or its negation. Consider any clause c consisting of the disjunction of three literals l_1, l_2, l_3. Node q_c is friend with the set of nodes

$$\{v_{l_1}, v_{l_2}, v_{l_3}, r_c\} \cup \tilde{Q}_c$$

The set $\{r_c\} \cup \tilde{Q}_c$ contributes exactly $\Delta - 1$ to the weighted degree (Definition 2) of node q_c. On the other hand, Claim 11 states that the node q_c is white. Thus, at least one of the nodes $v_{l_1}, v_{l_2}, v_{l_3}$ must be black. In other words, every clause c contains a literal l such that the node v_l is black. Further, for all literals l, the nodes v_l and $v_{\neg l}$ get different colors (Claims 11, 12).

Now, consider an assignment for the 3SAT instance where for all literals l, if node v_l is black (resp. white), the value of l is set to 1 (resp. 0). By the arguments outlined in the previous paragraph, the resulting assignment is consistent (every literal and its negation gets complementary values) and satisfying (every clause has at least one literal set to 1). This completes the proof. □

□

A.7 Proof of Theorem 9

Proof. Given a BNE instance on a bipartite graph with partite-sets L, R and set of edges $E \subseteq L \times R$, the algorithm described in Figure 4 finds a coloring of the nodes into black and white. We show that the output coloring is a pure Nash equilibrium and approximates maximum revenue by a factor of 2.

Recall that the weighted degree (Definition 2) of a node i is denoted by $\mathcal{D}(i)$. For each node i, Algorithm 4 maintains two variables $Ext(i)$ and $M(i)$. The variable $Ext(i)$ denotes the current externality (Definition 1) of node i. The variable $M(i)$ denotes the total weight of the edges node i shares with her friends who are still uncolored. Initially no node is colored, and we have $Ext(i) = 0, M(i) = \mathcal{D}(i)$ for all nodes $i \in L \cup R$.

Algorithm 4 starts by iteratively coloring the rigid nodes (Definition 4) and updating the $Ext(i), M(i)$ variables. When there are no more rigid nodes, let L' (resp. R') denote the still uncolored subset of L (resp. R). If $|L'| \geq |R'|$, then all nodes in the set L' are colored white and all nodes in the set R' are colored black. Else if $|L'| < |R'|$, then all nodes in the set L' are colored black and all nodes in the set R' are colored white.

Definition 4. *Suppose a subset of the nodes $L \cup R$ have been colored. An uncolored node i is called* rigid *if either $Ext(i) > \Delta$ or $Ext(i) + M(i) < \Delta$. In the former (resp. latter) case, node i must be colored white (resp. black) if we cannot change the colors of the nodes that have already been colored, and the final coloring is to be a pure Nash equilibrium.*

GREEDY

INPUT: An instance of BNE problem on bipartite graph $(L, R, E \subseteq L \times R)$.
OUTPUT: A coloring of the nodes.

FOR ALL nodes j: $Ext(j) \leftarrow 0$; $M(j) \leftarrow \mathcal{D}(j)$;
WHILE there exists some rigid [a] uncolored node i
  IF $Ext(i) + M(i) < \Delta$ THEN Node i is colored black;
    FOR ALL nodes j adjacent to i:
      $Ext(j) \leftarrow Ext(j) + w(i, j)$; $M(j) \leftarrow M(j) - w(i, j)$;
  ELSE IF $Ext(i) > \Delta$ THEN Node i is colored white;
    FOR ALL nodes j adjacent to i:
      $M(j) \leftarrow M(j) - w(i, j)$;
Let L' (resp. R') be the still uncolored subset of L (resp. R);
IF $|L'| \geq |R'|$ THEN all nodes in L' are colored white and in R' colored black;
ELSE IF $|L'| < |R'|$ THEN all nodes in L' are colored black and in R' colored white;

[a] See Definition 4

Fig. 4. Approximation algorithm for the BNE problem on bipartite graphs.

Let $S = (L \cup R) \setminus (L' \cup R')$ be the set of nodes considered rigid during some iteration of the WHILE loop. Without any loss of generality, assume $|L'| \geq |R'|$ (see Figure 4).

In the final coloring, each white (resp. black) node from the set S, by definition, has an externality of at least (resp. at most) Δ. Now consider a node $i \in L'$. At the end of the execution of the WHILE loop: node i has $Ext(i) + M(i) > \Delta$. In the bipartite graph, all the neighbors of i that are still uncolored belong to the set R'. In the final coloring, all nodes in the set R' are black, and node i is white. As a consequence, node i gets an externality of at least Δ, and plays the best response in the final coloring. It can be shown similarly that each node $j \in R'$ also plays best response in the final coloring. Thus, the coloring returned by the algorithm in Figure 4 is a pure Nash equilibrium.

Let $Rev(S)$ denote the revenue obtained from the node set S. Every pure Nash equilibrium will color nodes belonging to the set S exactly the same way as done by our algorithm, and get the same revenue $Rev(S)$. The revenue obtained from the remaining set of nodes $L' \cup R'$, pretending all of them can be colored white, is upper bounded by $p_2 \times (|L'| + |R'|)$. But our algorithm gets a revenue at least $p_2 \times \max(|L'|, |R'|)$ from the set $L' \cup R'$. This gives the desired 2 approximation. □