

Mechanisms and Allocations with Positive Network Externalities

ANAND BHALGAT, University of Pennsylvania
SREENIVAS GOLLAPUDI, Microsoft Research Search Labs
KAMESH MUNAGALA, Duke University

With the advent of social networks such as Facebook and LinkedIn, and online offers/deals web sites, network externalities raise the possibility of marketing and advertising to users based on influence they derive from their neighbors in such networks. Indeed, a user's knowledge of which of his neighbors "liked" the product, changes his valuation for the product. Much of the work on the mechanism design under network externalities has addressed the setting when there is only one product. We consider a more natural setting when there are multiple competing products, and each node in the network is a unit-demand agent.

We first consider the problem of welfare maximization under various different types of externality functions. Specifically we get a $O(\log n \log(nm))$ approximation for concave externality functions, a $2^{O(d)}$ -approximation for convex externality functions that are bounded above by a polynomial of degree d , and we give a $O(\log^3 n)$ -approximation when the externality function is submodular. Our techniques involve formulating non-trivial linear relaxations in each case, and developing novel rounding schemes that yield bounds vastly superior to those obtainable by directly applying results from combinatorial welfare maximization.

We then consider the problem of Nash equilibrium where each node in the network is a player whose strategy space corresponds to selecting an item. We develop tight characterization of the conditions under which a Nash equilibrium exists in this game. Lastly, we consider the question of pricing and revenue optimization when the users in the network are selfish agents, and their private information is the vector of valuations for different items. We show that for single parameter settings (when an agent's intrinsic valuation for every item can be described using one parameter), all our approximation results for welfare maximization extend to revenue maximization. For the multi-parameter setting, we design an $O(1)$ -approximate revenue optimal mechanism for IID agents, when the action of a single agent does not affect the externality enjoyed by remaining agents.

Categories and Subject Descriptors: F.2.2 [Theory of Computation]: Nonnumerical Algorithms and Problems; J.4 [Computer Applications]: Social and Behavioral Sciences

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Mechanism Design, Approximation Algorithms

1. INTRODUCTION

Externalities exist in many online commercial settings. Such settings include the hugely popular social networks such as Facebook and LinkedIn and more adhoc networks such as those induced by a particular Groupon offer, online blogs and product review pages. In each context, the externality can be respectively characterized by the

Anand Bhalgat is supported by NSF grants CCF-1116961 and IIS-0904314. Kamesh Munagala is supported by an Alfred P. Sloan Research Fellowship, a gift from Cisco, and by NSF grants CCF-0745761, CCF-1008065, and IIS-0964560. Part of this work was done while Munagala was visiting Twitter Inc.

Author's addresses: Anand Bhalgat, Department of Computer and Information Science, University of Pennsylvania; Sreenivas Gollapudi, Microsoft Research Search Labs and Kamesh Munagala, Department of Computer Science, Duke University.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

number of friend’s ”likes” in Facebook, the number of participants in an online deal offer and the number of positive and negative reviews/comments on a product web page. Such models of externalities raise the possibility of marketing and advertising to users based on influence they derive from such networks. In this paper, we consider allocation and pricing problems when agents in a network derive positive externalities from their ”neighbors” in the network. The broad question we study is not new: there has been a significant body of work that considers allocation and pricing (see [Hartline et al. 2008; Akhlaghpour et al. 2010; Candogan et al. 2010]), as well as equilibrium selection problems (see [Haghpanah et al. 2011; Bhattacharya et al. 2011]) in the context of allocating multiple copies of a *single item*. These works essentially assume users (or agents) in the network are unit-demand, and the item itself is available in unlimited supply; the literature mainly differs in the model of externality (positive, submodular, negative), as well as the class of allocation and pricing mechanisms (sequential pricing, auctions, differential pricing).

In this paper, we depart from the single item assumption, and consider the more realistic situation where there are multiple competing products available, and the users in the network are demand-constrained. We show that this leads to new algorithmic and technical challenges that are entirely absent in the single item setting. Though our theoretical model is very basic, it has the concrete practical aforementioned motivations. In these contexts, there are typically a large number of advertisers who represent diverse products. Each of these products is typically targeted at a specific set of users depending on reported and inferred interests. Furthermore, for each of the products, each user enjoys positive externalities that depend on which of his neighbors in the network the ad was shown to. In fact, in the case of Facebook, the user would know which of his neighbors ”liked” the ad. In another context, a web site showing multiple online coupon/deal offers may also report how many people have already purchased a coupon or shown interest in a deal. In order to charge the advertiser, the network will either use a simple model of externality a user will derive, or infer a more complex model from historical data. We do not consider the inference problem in this paper; instead, given a generic model of externality, we consider the computational and economic questions that arise in optimizing for this model.

1.1. Externality Model

Without further ado, we present the model that we study. The social network is modeled by a directed graph $G(V, E)$ on the users, where edges represent users who are connected in the network. This connection can either be symmetric (in networks such as Facebook) or asymmetric (in networks such as Twitter). In either case, we assume vertex j sees influence from vertices k such that there is a directed (resp. undirected) edge from k to j . Let N_j denote the set of neighbors k such that there is a directed edge from k to j . For simplicity in exposition, we assume that $j \in N_j$.

The seller has m distinct types of items (or ads), each with an unlimited supply. We denote a generic item by i . We define *intrinsic valuation* of agent j for item i as his valuation in absence of any externality effect; we denote it by v_{ij} . We denote the set of items for which agent j has non-zero intrinsic valuation by \mathcal{I}_j . Buyer j ’s valuation of item i is affected for the set of other agents in the network possessing the item; we refer to this effect as the *network externality*. If $S(j \in S)$ is the set of agents possessing the item i , then agent j ’s valuation for item i is given by $v_{ij}F_{ij}(S)$, where $F_{ij} : 2^{[n]} \rightarrow \mathbf{R}$ is the *externality function* for agent j and item i .^a In this paper, we restrict attention to positive, monotone and non-decreasing externality functions.

^aIn this paper, we consider a multiplicative model for simplicity of exposition; our results extend to an additive model of externality, such as $v_{ij} + F_{ij}(S)$ quite easily.

Our results differ depending on the precise functional form of the externality function. Suppose S_i is the set of vertices that are allocated item i .

Compact Externality. In this class of externality, for any given vertex $j(j \in S_i)$ and item i , we have $\forall S_i \subseteq V, F_{ij}(S_i) = \mathcal{F}_{ij}(|S_i \cap N_j|)$, where $\mathcal{F}_{ij} : \mathbb{N} \rightarrow \mathbb{R}^+$ is an arbitrary non-decreasing function. Within this class of functions, we distinguish *linear*, *convex*, and *concave* externality, corresponding respectively to the cases where \mathcal{F}_{ij} is linear, convex, and concave. If \mathcal{F}_{ij} is such that $\mathcal{F}_{ij}(1) = 0$ and $\mathcal{F}_{ij}(k) = 1, \forall k \geq 2$, then we term it as *step function* externality. In other words, a vertex requires at least one of his neighbors to be allocated the same item for his intrinsic valuation to realize.

Submodular Externality. In this case, for all $i, j, S \subseteq V, F_{ij}(S) = \mathcal{F}_{ij}(S)$, where the function $\mathcal{F}_{ij}(S) : 2^{[n]} \rightarrow \mathbb{R}^+$ is an arbitrary sub-modular function.

The key assumption that makes our model interesting is that vertices (or users) in the network are *unit demand*,^b and that the user cannot be allocated more than 1 item. This is typically imposed due to constraints on the number of ads that can be shown to a user without causing fatigue or cluttering his screen real estate.^c In this model, we consider several problems, studying both computational as well as economic aspects.

1.2. Problem Statements and Results

Allocation and Welfare Maximization. We first consider the problem of allocating items to users (or vertices) in order to maximize the overall welfare. This is a special case of combinatorial welfare maximization problems [Blumrosen and Nisan 2007]. In these problems, there are m agents and n items. Each item can be allocated to only one agent; if agent i is allocated subset S_i of items, then his welfare is $h_i(S_i)$. The goal is to allocate the items to agents in order to maximize the total welfare. For welfare maximization, the results depend on the nature of the functions h_i . If these are sub-modular, there is a constant factor approximation [Lehmann et al. 2001; Feige and Vondrák 2010; Calinescu et al. 2007]; these results extend in the so-called *demand oracle* model to sub-additive functions; on the other hand, if these functions are general, the problem encodes independent set and becomes hard to approximate [Blumrosen and Nisan 2007] to within a factor of \sqrt{n} where n is the number of items.

To see that our problem is a special case of combinatorial welfare maximization, we map the items i in our problem to agents, and the users j in the social network to items. If a set S_i of users is assigned item i , the function $h_i(S_i) = \sum_{j \in S_i} v_{ij} F_{ij}(S_i \cap N_j)$. We note that even if F_{ij} is submodular (or concave), the function h_i will not be submodular, so that we cannot directly use results from submodular welfare maximization.

We present approximation algorithms for the welfare maximization version of our problem. Our main technique involves formulating non-trivial linear relaxations in each case, and developing rounding schemes that yield bounds that vastly improve bounds obtainable by directly applying results from combinatorial welfare maximization. The key hurdle to using a simple relaxation is the non-linear nature of the externalities - this causes such relaxations to have arbitrarily bad integrality gaps. To get around this problem, we write the linear program over possible values of the externality. This in itself is a key technical contribution of this paper. Our techniques are

^bAgain, the unit demand assumption is for simplicity of exposition; our results easily extend to a user i having a limit K_i on demand.

^cWe assume that the agent can only be allocated one item (as opposed to being allocated many items but consuming only the one that yields maximum utility). This is motivated by applications like Facebook where there are inherent constraints on the number of ads that can be shown to a user, as well as by the nature of the externality - if a user does not consume an item, he cannot influence his neighbors using that item.

inspired by algorithms for the metric labeling problem [Kleinberg and Tardos 1999; Calinescu et al. 2004]; however, as already mentioned, since we are dealing with convex or concave functions, linear relaxations for the metric labeling problem cannot be applied directly, and we need significant new ideas.

We show the following approximation bounds, which depend both on the nature of the externality, as well as the nature of the underlying social graph. In these bounds, n is the number of vertices in the graph. In order to motivate the need for approximation algorithms, we first show (Section 3) that the welfare maximization problem is MAX-SNP hard even for a complete graph and linear externality.

- For concave externalities (Section 3.2), we present a $O(\log n(\log n + \log m))$ approximation for general graphs; the bound improves to a constant factor for complete graphs.
- For submodular externalities, we present a $O(\log^3 n)$ approximation in the demand oracle model. Its proof is deferred to the full version of the paper.
- For convex externalities (Section 3.3), we present a $2^{O(d)}$ approximation if the function \mathcal{F}_{ij} is bounded above by a degree d polynomial. As a special case, for linear externalities, we present a constant factor approximation.
- We get a $O(1)$ -approximation for step-function externalities; our LP relaxation in Section 3.4 directly implies a $O(1)$ approximate *truthful-in-expectation mechanism* (where the agents are vertices in the social network, with private valuation for each item). Further, we show that this result implies a $O(\log^2 n)$ approximation to the optimal revenue of a truthful-in-expectation mechanism in a prior-free setting.

Nash Equilibria. We next consider selfish behavior on the parts of users in the network. Suppose users could decide which items they are allocated, and choose that item which maximizes their individual welfare given the allocation to the other agents. Then does this process converge to a Nash Equilibrium, and if so how good is the welfare of this equilibrium compared to the optimal social welfare? In Section 4, we show that Nash equilibria with good price of anarchy exist when for any given item i and the set $S_i \subseteq V$ of users that are allocated item i , the externality effect observed by every user for item i is same. This corresponds to the setting where the underlying social network is a complete graph, and for any given item i , the externality function (F_{ij}) for all users is same. We show that these assumptions are tight, and a Nash equilibrium may not exist without these assumptions. Our techniques for showing existence involve careful construction of potential functions, hence showing that the Nash dynamics leads to a pure Nash equilibrium.

Bayesian Mechanism Design: Pricing and Revenue Maximization. We finally turn to the question of pricing and revenue optimization in Section 5. In our model, the users in the network are selfish agents, and their private information is the vector of valuations for different items. We assume that the externality function is public knowledge, since it is typically inferred from the network topology and shared interests.

Single Parameter Setting: In Section 5.1, we first consider the classic single-parameter setting, where for any given agent j , his intrinsic valuations for all items in \mathcal{I}_j are same; this *roughly* corresponds to scenarios where the competing products are of same quality. For each agent, this valuation is drawn independently from an agent-specific prior distribution. The goal is to design a truthful mechanism to maximize expected revenue, where the expectation is over the prior distributions. We combine techniques developed in this paper with techniques developed in [Hartline and Lucier 2010; Haghpanah et al. 2011] to show that, when the distributions are regular, all our approximation results for welfare maximization extend to revenue maximization as well, with the loss of a small constant factor in the approximation ratio. This provides

additional justification for focusing on welfare maximization as the core optimization problem.

We note that [Hartline and Lucier 2010] give a black-box reduction from an α -approximate welfare maximizing algorithm to a BIC mechanism that is an α -approximation to the optimal welfare; their approach depends upon making the allocation curve monotone in (real) valuations using the ironing technique. To design a BIC mechanism with an α -approximation to revenue, we use the α -approximate welfare maximizing algorithm on virtual valuations. If the allocation curve is not monotone in (real) valuations, then we make it monotone using ironing. With regular distributions, the virtual valuation is monotone in (real) valuation, and this step only increases the revenue of the mechanism.

Our result also extends to a setting where an agent’s valuations for different items are scaled by some publicly known nonnegative quantity relative to each other. To be precise, agent j is associated with a set of publicly known constants $\{c_{ij} | 1 \leq i \leq m\}$, such that when the parameter drawn from his distribution is v_j , then his intrinsic valuation for item i is $c_{ij} \times v_j$. This corresponds to a setting where an agent’s valuations of different items are completely correlated.

Multi-Parameter Setting: The key hurdle to extending the above results to the multi-parameter setting (where each agent’s intrinsic valuations for different items are different) is that there is no good characterization for even approximately optimal mechanisms in this setting, and furthermore, our results in Section 4 show that even without pricing, the agents’ behavior may often not admit a Nash equilibria. To overcome these challenges, we consider a simpler setting, where there are infinitely many non-atomic agents located on a complete graph, there is commonly known distribution G_i for item i , and agent j ’s intrinsic valuation for item i is drawn from G_i independent of other agents and items. The latter prior model is similar to that first studied in [Chawla et al. 2007], and subsequently in several papers [Chawla et al. 2010; Bhattacharya et al. 2010; Alaei 2011; Cai and Daskalakis 2011]. The reason for non-atomic agents is that the action of a single agent does not affect the externality enjoyed by the remaining agents - this in some sense, decouples the actions of the agents, and ensures that the Nash equilibrium is not affected by the action of a single agent. In this setting, under mild assumptions on agents’ behavior, we show a constant factor approximation to optimal revenue via a sequential posted pricing mechanism (Section 5.2). Our key technique involves formulating a quadratic program for the revenue of the optimal mechanism, showing that this quadratic program admits to a poly-time solution via dynamic programming, and using this relaxation to develop a posted pricing scheme.

1.3. Organization

The rest of the paper is organized as follows. In Section 2, we provide some definitions and relevant background material. In Section 3, we present our results for the welfare maximization problem. In Section 4, we present our results on the Nash equilibrium of the allocation game, and our results on Bayesian mechanism design are given in Section 5. Missing proofs can be found in the full version of the paper.

2. PRELIMINARIES

We have already presented the graph-based externality model in Section 1.1. This directly defines the allocation and welfare maximization problem as follows: Compute an allocation of items to the vertices so that the total welfare enjoyed by all vertices is maximized. We now proceed to define Nash equilibria and revenue maximization in this setting.

Allocation Game and the Nash Equilibrium. We study a game \mathcal{G} , where we have an infinite supply of each item, each buyer is a player, the strategy space for a buyer is to choose an item and his valuation of item i is $F_{ij}(S_i)v_{ij}$, where S_i is the set of buyers currently using item i . In a Nash equilibrium, player j uses the item $\operatorname{argmax}_i F_{ij}(S_i)v_{ij}$.

The *price of anarchy* (PoA) of a game is defined as the ratio of the optimal welfare to the minimum welfare in any Nash equilibrium, and the *price of Stability* (PoS) is defined as the ratio of the optimal welfare to the maximum welfare in a Nash equilibrium.

Mechanism Design and Incentive Compatibility. In a *mechanism*, buyers report their *bids* or valuations for each item to the seller, and the seller chooses to allocate service/items to the buyers based on the bids. Each buyer wishes to maximize his value of service/item received minus the price paid. Let $\mathbf{v} = [v_1, v_2, \dots, v_n]$ be the reported bid vector, and $\mathbf{v}_j = [v_{1j}, v_{2j}, \dots, v_{mj}]$ is buyer j 's bid, where v_{ij} is buyer j 's intrinsic valuation for item i . Given mechanism \mathcal{M} and bid vector \mathbf{v} , we use $w_j(\mathcal{M}(\mathbf{v}))$ and $p_j(\mathcal{M}(\mathbf{v}))$ to indicate buyer j 's welfare and payment in the mechanism. Mechanism \mathcal{M} is said to be *incentive compatible* if for each $j, \mathbf{v}, \mathbf{v}'_j$, we have

$$w_j(\mathcal{M}(\mathbf{v})) - p_j(\mathcal{M}(\mathbf{v})) \geq w_j(\mathcal{M}(\mathbf{v}_{-j}, \mathbf{v}'_j)) - p_j(\mathcal{M}(\mathbf{v}_{-j}, \mathbf{v}'_j))$$

it is said to be *truthful in expectation* if for each $j, \mathbf{v}, \mathbf{v}'_j$, we have

$$\mathbf{E}_{\mathcal{M}} [w_j(\mathcal{M}(\mathbf{v})) - p_j(\mathcal{M}(\mathbf{v}))] \geq \mathbf{E}_{\mathcal{M}} [w_j(\mathcal{M}(\mathbf{v}_{-j}, \mathbf{v}'_j)) - p_j(\mathcal{M}(\mathbf{v}_{-j}, \mathbf{v}'_j))]$$

and it is said to be *Bayesian incentive compatible* if for each $\mathbf{v}_j, \mathbf{v}'_j$, we have

$$\mathbf{E}_{\mathbf{v}_{-j}, \mathcal{M}} [w_j(\mathcal{M}(\mathbf{v})) - p_j(\mathcal{M}(\mathbf{v}))] \geq \mathbf{E}_{\mathbf{v}_{-j}, \mathcal{M}} [w_j(\mathcal{M}(\mathbf{v}_{-j}, \mathbf{v}'_j)) - p_j(\mathcal{M}(\mathbf{v}_{-j}, \mathbf{v}'_j))]$$

In a *single dimensional mechanism*, each buyer's total value for a set of goods is specified by a single parameter. Suppose buyer j 's parameter is drawn independently from a distribution with CDF G_i and PDF g_i . Then the virtual valuation $\varphi_j(v)$ for value v is given by $v - \frac{1-G_j(v)}{g_j(v)}$. We say that distribution G_j is *regular* if $\varphi(v)$ is a non-decreasing function of v .

3. ALLOCATION PROBLEM

In this section, we address the problem of welfare maximization under positive network externalities. Recall the externality model from Section 1.1; our goal is to allocate items to the users in order to maximize total welfare enjoyed by all the users. We present approximation algorithms for this problem under various models of externality: In Section 3.2, we address the case where the externality function is concave for each vertex. Our result for convex externality is addressed in Section 3.3, and finally, we consider the case of step-function externality in Section 3.4. Our result on submodular externality is deferred to the full version of the paper.

In each case, our techniques build on carefully formulating a linear programming relaxation. The key hurdle to using a simple relaxation is the non-linear nature of the externalities - this causes such relaxations to have arbitrarily bad integrality gaps. To get around this problem, we write the linear program over possible values of the externality. This in itself is a key technical contribution of this paper.

Our welfare maximization LP for step-function externalities in Section 3.4 also directly implies a $O(1)$ approximate truthful-in-expectation mechanism (where the agents are vertices in the social network, with private valuation for each item). Further, we show that this result implies a $O(\log^2 n)$ approximation to the optimal revenue of a truthful-in-expectation mechanism.

3.1. Hardness of Approximation

Before presenting our approximation algorithms, we show that the welfare maximization problem is MAX-SNP hard even for linear externality and the complete graph model. We note that the minimization version of the problem (metric labeling for uniform metric) is known to be APX-hard [Kleinberg and Tardos 1999; Chekuri et al. 2001].

THEOREM 3.1. *The welfare maximization problem for linear externality and complete graphs is MAX-SNP hard.*

PROOF. The welfare maximization problem can be formulated as a variant of set cover. There are m sets (corresponding to the items), and n elements corresponding to the agents. The valuations v_{ij} are either 0 or 1. Denote the sets by S_i ; we say $j \in S_i$ if $v_{ij} = 1$. Allocating the items to the agents corresponds to assigning each element to one of the sets it belongs to. Since the externality is linear and the graph is complete, the welfare is simply computed as follows: Let n_i denote the number of elements assigned to S_i (or the number of agents assigned item i). Then, the welfare is $\sum_{i=1}^m n_i^2$. The goal is to find an assignment of elements to sets to maximize this quantity.

We will reduce the hardness from the Exact 3 Cover (X3C) problem. In this problem, we are given k elements X and a collection of sets, T with three elements each. The goal is to find a cover of the elements so that each element is covered by exactly one set. In the maximization version of this problem, the goal is to find the *largest* collection of sets so that each element is in exactly one set. This problem is known to be MAX-SNP-hard [Kann 1991]. Suppose it is NP-Hard to distinguish between instances of X3C where the maximum value is k and where it is at most $q = k(1 - \epsilon)$. In the former case, the welfare maximization problem has optimal value $\sum_{i=1}^k 3^2 = 9k$. In the latter case, there are q sets that cover 3 elements each, and the remaining $3(k - q)$ elements can only be covered by sets in a fashion so that each new set covers at most 2 elements. It is now easy to check that the optimal welfare is at most $9q + 6(k - q) = k(9 - 3\epsilon)$. For constant $\epsilon > 0$, the gap between the optimal welfare in the two cases is a constant, showing MAX-SNP hardness. \square

3.2. Concave Externality

Our result for concave externality is given in Theorem 3.2, and we consider the special case of a complete graph in Theorem 3.7.

THEOREM 3.2. *Given any undirected graph G , if $F_{ij}(S) = \mathcal{F}_{ij}(|S \cap N_j|)$ where \mathcal{F}_{ij} is a non-decreasing concave function, then there exists a polynomial time algorithm to compute an allocation with welfare $\Omega\left(\frac{OPT}{\log n(\log n + \log m)}\right)$.*

PROOF. In order to formulate a linear program, we perform a sequence of transformations to the input graph. We randomly partition the graph into two parts V_1 and V_2 , so that each vertex is independently allocated to either part with probability $1/2$. For a given partition of vertices into (V_1, V_2) , we will compute *restricted welfare*, which is derived only from the externality imposed by vertices in V_2 on vertices in V_1 . It is easy to check that since \mathcal{F}_{ij} is concave, the expected value of the restricted welfare is at least $OPT/4$, where the expectation is over the random partitioning^d.

Assume $n = |V|$ is a power of 2 for simplicity. For each vertex $j \in V_1$ and each item i , we create $\log n$ copies (i, j, a) , for $a = 0, 1, 2, 4, \dots, n$. Copy (i, j, a) is only used if j is

^dInfluence-and-Exploit approach has been used in the past for handling network externalities, see [Hartline et al. 2008; Haghpanah et al. 2011]. However it does not make our problem “easy” by any means and our contributions begin only after this step.

assigned item i , and there are at least a neighbors of j in V_2 which are also assigned this item. The value of (i, j, a) is $u(i, j, a) = v_{ij} \mathcal{F}_{ij}(a)$. Similarly, for each $k \in V_2$ and each item i , we create a copy (i, k) , which is used if k is assigned item i . We can map the optimal assignment to a selection of vertices of this form, so that for each $j \in V_1$ or $k \in V_2$, at most one copy is chosen. It is easy to check that the value of this allocation is within a factor of 2 of the optimal restricted welfare for this partition, where the loss of a factor 2 is because we scaled the a 's in powers of 2.

Linear Program: Randomly choose one value for a in the powers of 2 between 1 and n , say a^* , and only consider the welfare derived from vertices in V_1 with a^* neighbors in V_2 . Let $y(i, j) = 1$ if copy (i, j, a^*) is chosen for $j \in V_1$. Similarly, let $z(i, k) = 1$ if copy (i, k) is chosen for $k \in V_2$. Then the expected welfare of the following integer program (IP1) is a $O(\log n)$ approximation to OPT, where the expectation is over the random choice of a . The rounding scheme for the linear relaxation of IP1 is given in Algorithm 1, and we bound its welfare in Lemma 3.3.

$$\begin{aligned} \text{Maximize } & \sum_{i,j} u(i, j, a^*) y(i, j) \\ \sum_{k \in N_j \cap V_2} \min(z(i, k), y(i, j)) & \geq a^* y(i, j) \quad \forall i, j \in V_1 \\ \sum_i z(i, k) & \leq 1 \quad \forall k \in V_2 \\ \sum_i y(i, j) & \leq 1 \quad \forall j \in V_1 \\ y(i, j), z(i, k) & \in \{0, 1\} \quad \forall i, j \in V_1, k \in V_2 \end{aligned}$$

ALGORITHM 1: Rounding Scheme for Linear Relaxation of IP1

Select an integer ℓ uniformly at random between 0 and $2(\log n + \log m)$, set $\gamma \leftarrow 2^{-\ell}$;

Round down values of $y(i, j)$ to the nearest power of 2;

For each $i, j \in V_1$, set all $y(i, j) \neq \gamma$ to 0;

For each $i, k \in V_2$, set $z(i, k) \leftarrow \min\{z(i, k), \gamma\}$;

for each item i do

 Choose δ_i u.a.r. in $[0, 1]$;

if $\delta_i \leq \gamma$ **then**

 For each $j \in V_1$ with $y(i, j) = \gamma$ that is unassigned so far, assign item i to it w.p. $1/3$;

 For each $k \in V_2$ that is unassigned so far, assign item i to it w.p. $\min\left\{1/3, \frac{z(i, k)}{3\gamma}\right\}$;

end

end

LEMMA 3.3. *The expected welfare of the solution computed by Algorithm 1 is a $O(\log(n) + \log(m))$ approximation to IP1.*

PROOF. Rounding down values of $y(i, j)$ for $j \in V_1$ to the nearest power of 2 leads to a factor 2 loss in the approximation, as \mathcal{F}_{ij} is concave for each i, j . Values of $y(i, j)$ s smaller than $\frac{1}{n^2 m^2}$ can be neglected with a loss of $\frac{1}{nm}$ factor: if their contribution is more $\frac{1}{mn}$ factor of the LP objective, then scale their values up by a factor of nm , and assign a value of 0 to other $y(i, j)$ s. If such transformation is possible, then it is a contradiction since the welfare is now more than the LP objective. Clearly packing constraint for each vertex $j \in V_1$ is not violated. Assigning the value of $1/m$ to each $z(i, k)$ makes the LP feasible.

Thus we get $2(\log n + \log m)$ groups of vertices in V_1 . If we randomly choose one of these groups, and neglect all other groups in V_1 , then the welfare is affected by a

factor of $2(\log n + \log m)$. Thus we get that the expected welfare of the following integer program (IP2) is $\Omega\left(\frac{\text{OPT}}{\log n(\log(n)+\log(m))}\right)$, where the expectation is over the choice of γ .

$$\text{Maximize } \sum_{i,j} u(i, j, a^*)y(i, j)$$

$$\begin{aligned} \sum_{k \in N_j \cap V_2} z(i, k) &\geq a^*y(i, j) \quad \forall i, j \in V_1 \\ \sum_i z(i, k) &\leq 1 \quad \forall k \in V_2 \\ \sum_i y(i, j) &\leq 1 \quad \forall j \in V_1 \\ y(i, j) &\in \{0, \gamma\} \quad \forall i, j \in V_1 \\ z(i, k) &\in [0, \gamma] \quad \forall i, k \in V_2 \end{aligned}$$

To analyze the rounding algorithm, consider any $(i, j \in V_1)$ pair with $y(i, j) = \gamma$ (after scaling). Its contribution to the objective of IP2 is $u(i, j, a^*)y(i, j)$. We will calculate its expected contribution to the integer solution. We first note an important observation.

LEMMA 3.4. *For any item i and vertex $j \in V_1 \cup V_2$, the probability that j is not assigned to any item before i is at least $2/3$.*

PROOF. For a vertex $j \in V_2$, the probability that he is assigned item i' (conditioned on no previous allocation) is

$$\min\left\{\frac{1}{3}, \frac{z(i', j)}{3\gamma}\right\} \times \Pr[\delta_i \leq \gamma] \leq \frac{z(i', j)}{3}$$

Thus the probability that he is assigned an item before i is at most $\sum_{i' < i} \frac{z(i', j)}{3} \leq \frac{1}{3}$.

For a vertex $j \in V_1$, for any i' with $y(i', j) = \gamma$, the probability that he is assigned item i' (conditioned on no previous allocation) is $\frac{\gamma}{3} = \frac{y(i', j)}{3}$, and it is 0 otherwise. Thus the probability that he is assigned an item before i is at most $\sum_{i' < i} \frac{y(i', j)}{3} \leq \frac{1}{3}$. \square

From now on, fix an $(i, j \in V_1)$ pair with $y(i, j) = \gamma$, and condition on the event that j has not been assigned to any item before i . In this event, for any $k \in V_2 \cap N_j$, define $\bar{z}(i, k) = z(i, k)$ if k has not yet been assigned an item yet, and $\bar{z}(i, k) = 0$ otherwise.

$$\text{LEMMA 3.5. } \mathbf{E}[\sum_{k \in V_2 \cap N_j} \bar{z}(i, k)] \geq \frac{1}{3} \sum_{k \in V_2 \cap N_j} z(i, k) \geq \frac{a^*\gamma}{3}$$

PROOF. The probability that j is not allocated before i is at least $2/3$, and the same holds for k as well. Using union bound, the probability that both j and k are available at round i is $1/3$. The lemma follows by summing over k , and by noting that $y(i, j) = \gamma$. \square

$$\text{LEMMA 3.6. } \Pr\left[\sum_{k \in N_j \cap V_2} \bar{z}(i, k) \geq \frac{a^*\gamma}{6}\right] \geq \frac{1}{5}$$

PROOF. Follows from the previous lemma, since $\frac{\sum_{k \in N_j \cap V_2} \bar{z}(i, k)}{\sum_{k \in N_j \cap V_2} z(i, k)} \leq 1$ always. \square

We note $\Pr[\delta_i \leq \gamma] = \gamma$. Conditioned on the event in which (a) $j \in V_1$ is unassigned before item i , (b) $\sum_{k \in N_j \cap V_2} \bar{z}(i, k) \geq \frac{a^*\gamma}{6}$, and (c) $\delta_i \leq \gamma$, we analyze the contribution of pair (i, j) to the welfare. The probability of this event is $\Omega(\gamma)$. In this event, the probability that item i is assigned to vertex j is $1/3$. As each $z(i, k) \leq \gamma$, using Chernoff's bound, the probability that item i is assigned to $\Omega(a^*)$ vertices in $N_j \cap V_2$ is at least $\Omega(1)$, and this event is independent of assignment of item i to vertex $j \in V_1$. Thus using concavity of \mathcal{F}_{ij} , the contribution of (i, j) to the integral solution is $\Omega(\gamma u(i, j, a^*))$, where as its contribution to IP2 is $u(i, j, a^*)y(i, j) = \gamma u(i, j, a^*)$. This completes the proof.

Complete Graph. In the following theorem, we consider the special case when G is a complete graph, and show a constant approximation.

THEOREM 3.7. *Given any clique G on n vertices, if $F_{ij}(S) = \mathcal{F}_{ij}(|S|)$ where \mathcal{F}_{ij} is a non-decreasing concave function, then there exists a polynomial time algorithm to compute an allocation with welfare $\Omega(\text{OPT})$.*

PROOF. We create n copies of item i , namely $(i, 1), (i, 2), \dots, (i, n)$, where the copy (i, a) corresponds to item i being given to a vertices. We treat each copy as a different item. Let $u(i, j, a) = v_{ij} \mathcal{F}_{ij}(a)$ denote the valuation if agent j is assigned to item i . Let $x(i, j, a) = 1$ if vertex j is assigned to item (i, a) and let $y(i, a) = 1$ if item (i, a) is opened. We have the following integer program that upper bounds the optimal welfare, and the rounding scheme is given in Algorithm 2.

$$\begin{aligned} & \text{Maximize } \sum_{i,j,a} u(i, j, a)x(i, j, a) \\ & \sum_j x(i, j, a) \geq ay(i, a) \quad \forall i, a \\ & \sum_{i,a} x(i, j, a) \leq 1 \quad \forall j \\ & x(i, j, a) \in [0, y(i, a)] \quad \forall i, j, a \\ & y(i, a) \in \{0, 1\} \quad \forall i, a \end{aligned}$$

To analyze the above algorithm, we note that the contribution of an (i, j, a) pair to

ALGORITHM 2: Rounding Scheme for the Complete Graph

Solve the linear relaxation of the integer program;

Scale down all variables by a factor of 3;

for each item (i, a) do

 Choose δ_{ia} uniformly at random in $[0, 1]$;

if $y(i, a) \geq \delta_{ia}$ **then**

 Item (i, a) is marked active;

 For each vertex j that is unassigned so far, independently assign it to (i, a) w.p. $\frac{x(i, j, a)}{3y(i, a)}$;

end

end

the LP objective is $u(i, j, a)x(i, j, a)$. The following lemma bounds its contribution to the integral solution; its proof is deferred to the full version of the paper. The theorem then follows by summing over all (i, j, a) pairs.

LEMMA 3.8. *The expected contribution of vertex j and item (i, a) to the integral solution is $\Omega(x(i, j, a)u(i, j, a))$.*

□

3.3. Convex Externality

In this section, we consider the case when the externality function is a non-decreasing convex function, the following theorem states our result.

THEOREM 3.9. *When the externality function is convex, and bounded above by a polynomial of degree d , there exists a polynomial time algorithm to compute a $2^{O(d)}$ approximation.*

This gives a $O(1)$ approximation when d is a constant. In the rest of the section, we prove Theorem 3.9. As in the case of concave externality, we partition the graph into two halves, V_1 and V_2 , by assigning each vertex independently to one of the halves. It is easy to check that the expected value of the restricted welfare is at least $\frac{\text{OPT}}{4^d}$, where the expectation is over random partitioning.

The following LP upper bounds the optimal restricted welfare (Lemma 3.10). In the LP, $x(i, j, a)$ is the probability that vertex $j \in V_1$ is assigned item i and a of its neighbors in V_2 are also assigned item i , and $z(i, k)$ is the probability that vertex $k \in V_2$ is assigned item i . Furthermore, $y(i, j, a)$ is *roughly* the probability that a vertices in $N_j \cap V_2$ are assigned item i . We define $u(i, j, a)$ to be $v_{ij} \mathcal{F}_{ij}(a)$. The rounding scheme is given in Algorithm 3.

$$\begin{aligned} & \text{Maximize} && \sum_{i,j \in V_1, a} x(i, j, a) u(i, j, a) \\ & \sum_{\ell \leq a \leq |N_j \cap V_2|} a \times y(i, j, a) \leq \sum_{k \in N_j \cap V_2} \min\{z(i, k), \sum_{\ell \leq a \leq |N_j \cap V_2|} y(i, j, a)\} && \forall i, j \in V_1, a \\ & \sum_{i,a} x(i, j, a) \leq 1 && \forall j \in V_1 \\ & \sum_a y(i, j, a) = 1 && \forall i, j \in V_1 \\ & \sum_i z(i, k) \leq 1 && \forall k \in V_2 \\ & x(i, j, a), y(i, j, a), z(i, k) \in [0, 1] && \forall i, j \in V_1, k \in V_2, 0 \leq a \leq |N_j \cap V_2| \end{aligned}$$

LEMMA 3.10. *The above LP is a relaxation to the optimal integral solution for the restricted welfare problem.*

PROOF. Given an instance of the problem, and an optimal solution to the instance, we construct a feasible solution for the LP as follows. We set $x(i, j, a) = 1$ if vertex $j \in V_1$ is assigned item i and a of his neighbors in V_2 also possess item i . Similarly, we set $z(i, k) = 1$ if vertex $k \in V_2$ is assigned item i . Furthermore, we assign $y(i, j, a) = 1$ if exactly a neighbors of vertex $j \in V_1$ have item i . Other variables are zero. Clearly the welfare remains unaffected. Now we check the feasibility of the solution. Constraints 1, 3, 4, 5 are trivially satisfied, and it remains to check the feasibility of the 2nd constraint. Fix vertex $j \in V_1$, let i be the item assigned to him and a be the number of his neighbors possessing the item. Then for $\ell > a$, both LHS and RHS are zero. For $\ell \leq a$, both LHS and RHS are exactly a . This proves the lemma. \square

ALGORITHM 3: Rounding Scheme for Convex Externality

```

for each item  $i$  do
  Choose  $\delta_i$  u.a.r. in  $[0, 1]$ ;
  for each vertex  $k \in V_2$  do
    if  $z(i, k) \geq \delta_i$  and no item has been assigned to  $k$  so far then
      Assign item  $i$  to  $k$  w.p.  $1/2$ ;
    end
  end

```

end

For each vertex $j \in V_1$, pick the best item based on its externality to V_2 ;

To prove the theorem, it suffices to establish that, for any given vertex $j \in V_1$, his expected welfare is $1/2^{O(d)}$ fraction of his contribution to the LP objective. Consider the round in which item i is assigned to vertices in V_2 , and let S

be the (random) set of vertices in $N_j \cap V_2$ that are not assigned any item previously. For each value of $a = 0, 1, 2, \dots, |N_j \cap V_2|$, we define event $\mathcal{E}(i, j, a)$ that is true when the threshold used for assigning item i to vertices in V_2 lies in the range $(\sum_{a+1 \leq \ell \leq |N_j \cap V_2|} y(i, j, \ell), \sum_{a \leq \ell \leq |N_j \cap V_2|} y(i, j, \ell)]$. We define $\mathcal{C}(i, j, a)$ as follows:

$$\mathcal{C}(i, j, a) = \frac{1}{y(i, j, a)} \left(\sum_{k \in N_j \cap V_2} \min \left\{ \max \left\{ z(i, k) - \sum_{a+1 \leq \ell \leq |N_j \cap V_2|} y(i, j, \ell), 0 \right\}, y(i, j, a) \right\} \right)$$

In other words, $\mathcal{C}(i, j, a)$ measures the LP mass for item i in $N_j \cap V_2$ in event $\mathcal{E}(i, j, a)$, it roughly corresponds to the expected number of neighbors of j in event $\mathcal{E}(i, j, a)$. We would want this quantity to be a .

We note important properties of these events: (a) $\mathcal{E}(i, j, a)$ is disjoint from $\mathcal{E}(i, j, a' \neq a)$, (b) $\Pr[\mathcal{E}(i, j, a)] = y(i, j, a)$ (by definition of rounding for vertices in V_2), and (c) for each $0 \leq \ell \leq |N_j \cap V_2|$,

$$\sum_{\ell \leq a \leq |N_j \cap V_2|} [\mathcal{C}(i, j, a) \times y(i, j, a) | y(i, j, a) \neq 0] \geq \sum_{\ell \leq a \leq |N_j \cap V_2|} a \times y(i, j, a) \quad (1)$$

This property follows by the second constraint in the LP, and the definition of $\mathcal{C}(i, j, a)$. It implies the following: given vertex $j \in V_1$ and item i , even if the expected number of his neighbors that are assigned item i in event $\mathcal{E}(i, j, a)$ is less than a , the expected number of neighbors that are assigned item i in events $\mathcal{E}(i, j, \geq a)$ is at least the required value.

For simplicity in analysis, we consider a modified algorithm for assigning an item to vertex $j \in V_1$, selecting the best item can only improve the welfare of the allocation. Consider items in order $1, 2, \dots, m$. Let i be the item in consideration. If j has not been assigned any item until now, let $\mathcal{E}(i, j, a)$ be the event that was true during assignment of item i to vertices in V_2 ; then assign item i to j w.p. $\frac{x(i, j, a)}{(2 \times y(i, j, a))}$, conditioned on $y(i, j, a) \neq 0$. We note that

$$\sum_{i, a} \frac{x(i, j, a) \Pr[\mathcal{E}(i, j, a)]}{2 \times y(i, j, a)} = \sum_{i, a} \frac{x(i, j, a) \times y(i, j, a)}{2 \times y(i, j, a)} = 1/2$$

Thus the probability that he is not assigned any item is at least $1/2$.

Recall that S is the (random) set of vertices in $N_j \cap V_2$ that are not assigned any item before i . We define a random variable $X(i, j, a)$ as follows:

$$X(i, j, a) = \frac{1}{y(i, j, a)} \left(\sum_{k \in S} \min \left\{ \max \left\{ z(i, k) - \sum_{a+1 \leq \ell \leq |N_j \cap V_2|} y(i, j, \ell), 0 \right\}, y(i, j, a) \right\} \right)$$

$X(i, j, a)$ is defined in a similar way as $\mathcal{C}(i, j, a)$, however $X(i, j, a)$ only considers vertices in $N_j \cap V_2$ which are unallocated so far, whereas $\mathcal{C}(i, j, a)$ considers all vertices in $N_j \cap V_2$. As the probability that $k \in N_j \cap V_2$ is not assigned an item prior to i is at least $1/2$, then for any a (with $y(i, j, a) \neq 0$), we have

$$\mathbf{E}_S[X(i, j, a)] \geq \frac{\mathcal{C}(i, j, a)}{2}$$

Thus the expected number of vertices in $N_j \cap V_2$ that are assigned item i in case of event $\mathcal{E}(i, j, a)$ is at least $\frac{\mathcal{C}(i, j, a)}{4}$. The second $1/2$ comes by the fact that vertex $k \in V_2$ is assigned an item i w.p. $1/2$ when $z(i, k)$ is greater than the threshold for item i .

Using convexity of \mathcal{F}_{ij} , the expected value of externality when $j \in V_1$ is assigned item i is

$$\sum_a \mathcal{F}_{ij}(\mathcal{C}(i, j, a)/4) \times \Pr[\mathcal{E}(i, j, a)] \times \frac{x(i, j, a)}{2 \times y(i, j, a)} \times \frac{1}{2} = \frac{1}{2^{O(d)}} \sum_a \mathcal{F}_{ij}(\mathcal{C}(i, j, a)) \times x(i, j, a)$$

We note an important property of the optimal LP solution: as \mathcal{F}_{ij} is non-decreasing, there exists an ℓ such that for $a > \ell$ we have $x(i, j, a) = y(i, j, a)$ and $x(i, j, a) = 0$ for $a < \ell$; furthermore $0 \leq x(i, j, \ell) \leq y(i, j, \ell)$. Thus the externality can be rewritten as (neglecting the constant term)

$$x(i, j, \ell) \times \mathcal{F}_{ij}(\mathcal{C}(i, j, \ell)) + \sum_{\ell+1 \leq a \leq N_j \cap V_2} y(i, j, a) \times \mathcal{F}_{ij}(\mathcal{C}(i, j, a))$$

Using monotonicity of \mathcal{F}_{ij} and Equation 1, we get

$$\geq x(i, j, \ell) \times \mathcal{F}_{ij}(\ell) + \sum_{\ell+1 \leq a \leq N_j \cap V_2} y(i, j, a) \times \mathcal{F}_{ij}(a) = \sum_{0 \leq a \leq N_j \cap V_2} \mathcal{F}_{ij}(a) \times x(i, j, a)$$

This completes the proof.

3.4. Step Function Externality

In this section, we consider the setting when the externality is a step function. The following theorem states our result.

THEOREM 3.11. *There exists a truthful in expectation mechanism that achieves a $O(1)$ -approximation to the optimal welfare.*

PROOF. We partition the graph into two halves V_1 and V_2 , by randomly assigning each vertex to either side w.p. $1/2$. It can be easily seen that the expected value of the optimal restricted welfare is at least $\frac{OPT}{4}$. The following LP upper bounds the optimal restricted welfare. In the LP, $y(i, j, k)$ is the probability that vertex $j \in V_1$ derives externality for item i from vertex $k \in V_2$, and $x(i, j)$ is the probability that item i is assigned to vertex $j \in V_1 \cup V_2$.

$$\text{Maximize } \sum_{i, j \in V_1} x(i, j) v_{ij}$$

$$\begin{aligned} x(i, j) &= \sum_{k \in N_j \cap V_2} y(i, j, k) & \forall i, j \in V_1 \\ y(i, j, k) &\leq x(i, k) & \forall i, j \in V_1, k \in N_j \cap V_2 \\ \sum_i x(i, \ell) &\leq 1 & \forall \ell \in V_1 \cup V_2 \\ x(i, \ell), y(i, j, k) &\in [0, 1] & \forall i, j \in V_1, k \in V_2, \ell \in V_1 \cup V_2 \end{aligned}$$

We give a rounding scheme such that the expected welfare of each vertex $j \in V_1$ is exactly $\frac{(e-1)}{4e}$ fraction of his welfare in LP. We then compute the VCG price of the fractional allocation for each vertex in V_1 , and the price asked to vertex $j \in V_1$ is $\frac{(e-1)}{4e}$ fraction of his VCG price. Clearly this mechanism is truthful-in-expectation. Now we illustrate the rounding scheme: assign item i to vertex $k \in V_2$ w.p. $x(i, k)$ independent of other vertices in V_2 , and if k is assigned item i , consider every vertex $j \in V_1$ such that $k \in N_j \cap V_2$, and assign item i to j w.p. $\frac{y(i, j, k)}{2x(i, k)}$, irrespective of past allocation to j . In this process, a vertex $j \in V_1$ can be assigned multiple items, as well as multiple copies of the same item. However, the probability that he is allocated at least one copy of item i is $\frac{(1-1/e)x(i, j)}{2}$. Fix item i , conditioned on $j \in V_1$ receiving item i , the expected number of additional items he receives is $1/2$. Thus item i is the only item allocated

to $j \in V_1$ w.p. at least $\frac{(e-1)}{4e} \cdot x(i, j)$, let this value be $x'(i, j)$. In this event, we drop the allocation to j w.p. $x'(i, j) - \frac{(e-1)}{4e} \cdot x(i, j)$. This completes the proof of the theorem. \square

We get the following corollary for the revenue maximization; its proof is deferred to the full version of the paper.

COROLLARY 3.12. *There exists a truthful-in-expectation mechanism with expected revenue $\Omega\left(\frac{\text{OPT}}{\log^2 n}\right)$.*

The mechanism is based on declaring a random reserve price of $\frac{\text{OPT}}{2^r}$ for $1 \leq i \leq 2 \log n$, and the price asked to a vertex is the max of his VCG price and the reserve price. The assumption about the knowledge of OPT can be removed by another loss of $O(\log n)$ in the approximation factor, using techniques in [Chakraborty et al. 2009].

4. NASH EQUILIBRIUM

In this section, we consider a game \mathcal{G} , where we have an infinite supply of each item, each vertex is a player, the strategy space for a player is to choose an item and his valuation of item i is $F_{ij}(S_i)v_{ij}$, where S_i is the set of players currently using item i . In a Nash equilibrium, player j uses the item $\text{argmax}_i \{F_{ij}(S_i)v_{ij}\}$.

Complete Graphs. We consider the case when the underlying graph is a complete graph, and the externality for an item solely depends upon the number of other players possessing the item, i.e. for every item i and player j , there exists a function $\mathcal{F}_{ij} : \mathbb{N} \rightarrow \mathbb{R}^+$ such that $F_{ij}(S) = \mathcal{F}_{ij}(|S|)$ for all S . We let \mathcal{F}_{ij} to be an arbitrary non-decreasing concave function.

In Lemma 4.1, we show existence of a Nash equilibrium, when the externality function of all players is same for any given item. Lemma 4.2 establishes that this assumption is tight, and a Nash equilibrium may not exist if this assumption is removed. Lemma 4.3 analyzes the PoA and PoS of this game when a Nash equilibrium exists^e.

LEMMA 4.1. *There exists a pure Nash equilibrium in game \mathcal{G} if for every $j, j' \in V, S \subseteq V, \mathcal{F}_{ij}(|S|) = \mathcal{F}_{ij'}(|S|)$. Furthermore, every sequence of improving moves leads to a Nash equilibrium.*

PROOF. For each item i , define $\Phi_i = \sum_{1 \leq k \leq |S_i|} \log(\mathcal{F}_{ij}(k)) + \sum_{j \in S_i} \log(v_{ij})$, and consider a potential function $\Phi = \sum_{1 \leq i \leq m} \Phi_i$. We analyze a move by player j changing his strategy from item i to item i' . His welfare for item i is $\mathcal{F}_{ij}(|S_i|)v_{ij}$, and his welfare after his move to item i' is $\mathcal{F}_{i'j}(|S_{i'} + 1|)v_{i'j}$. If the move is an improving move, then

$$\begin{aligned} \mathcal{F}_{i'j}(|S_{i'} + 1|)v_{i'j} &\geq \mathcal{F}_{ij}(|S_i|)v_{ij} && \text{i.e.,} \\ \log(\mathcal{F}_{i'j}(|S_{i'} + 1|)) + \log(v_{i'j}) &\geq \log(\mathcal{F}_{ij}(|S_i|)) + \log(v_{ij}) \end{aligned}$$

The change in the potential is

$$\Delta(\Phi_{i'}) + \Delta(\Phi_i) = \log(\mathcal{F}_{i'j}(|S_{i'} + 1|)) + \log(v_{i'j}) - \log(\mathcal{F}_{ij}(|S_i|)) - \log(v_{ij})$$

^e[Gairing et al. 2006] study convergence properties of congestion games with player specific latency functions. The game considered in Lemma 4.1 (for any given item, players have an identical externality function, and the externality graph is complete) can be seen as an un-weighted congestion game with player specific delay functions with a resource in the latter game corresponding to an item in the former, and the potential function used in Lemma 4.1 is similar to the potential function used in Section 3.1 by [Gairing et al. 2006]. However, once we relax our game to allow player specific externality functions (the cases considered in Lemmas 4.2 and 4.3) or when the graph is not complete, the two classes of games are fundamentally different.

which is greater than 0 for an improving move by a player. As there are only finitely many distinct states in the game, it implies the existence of a pure Nash equilibrium. \square

We note that the result in Lemma 4.1 holds for arbitrary \mathcal{F}_{ij} 's as well.

LEMMA 4.2. *There exists an instance of game \mathcal{G} with no Nash equilibrium.*

PROOF. Consider an instance with 3 players and 3 types of items, $v_{ij} = 1$ for all i, j . Players' externality functions satisfy following inequalities.

$$\begin{aligned} \mathcal{F}_{11}(2) &> \mathcal{F}_{21}(2) > \mathcal{F}_{21}(1) > \mathcal{F}_{31}(2) > \mathcal{F}_{11}(1) > \mathcal{F}_{31}(1) \\ \mathcal{F}_{22}(2) &> \mathcal{F}_{32}(2) > \mathcal{F}_{32}(1) > \mathcal{F}_{12}(2) > \mathcal{F}_{22}(1) > \mathcal{F}_{12}(1) \\ \mathcal{F}_{33}(2) &> \mathcal{F}_{13}(2) > \mathcal{F}_{13}(1) > \mathcal{F}_{23}(2) > \mathcal{F}_{33}(1) > \mathcal{F}_{23}(1) \end{aligned}$$

Furthermore, we have $\mathcal{F}_{ij}(3) = \mathcal{F}_{ij}(2)$ for each i, j . It can be noted that such externality functions can be implemented using concave functions. Now we prove that this game does not have any Nash equilibrium. (a) Player 1 can not be allocated item 3 in a Nash equilibrium, as he can switch to item 2 even if no other player is using item 2. Similarly players 2 and 3 cannot be allocated items 1 and 2 respectively in a Nash equilibrium. (b) If all players are allocated distinct items, then players 1, 2 and 3 have to be allocated items 2, 3 and 1 respectively, and in this case player 3 will switch to item 3. (c) If players 1 and 2 are allocated the same item, then it has to be item 2, and in this case player 3 needs to be allocated item 1. In this case, player 1 will switch to item 1. Similar contradiction can be shown when players 2 and 3 or players 1 and 3 are allocated the same item. This completes the proof. \square

LEMMA 4.3. *In the above mentioned game, if \mathcal{F}_{ij} is concave for each i, j , and*

- each v_{ij} is either 1 or 0, then the price of stability is 1 and the price of anarchy is between $2n/3$ and n .
- v_{ij} s are arbitrary, then price of stability is at least 2 and the price of anarchy is between $2n/3$ and n .

The proof of Lemma 4.3 is deferred to the full version of the paper.

5. REVENUE OPTIMAL MECHANISMS

In this section, we consider the problem of designing Bayesian incentive compatible mechanisms for optimizing revenue, where the buyers' intrinsic valuations are private knowledge, and we are specified a prior over these valuations. We consider the single dimensional setting in Section 5.1, and our result for the multidimensional setting is given in Section 5.2.

5.1. Single Dimensional Valuation

In this section, we consider the setting where the intrinsic valuation of buyer j for each item in \mathcal{I}_j is same and it is drawn from a publicly known distribution with CDF G_j (and PDF g_j). We further assume that these distributions are regular. We note Myerson's characterization for single dimensional mechanism in Theorem 5.1, and Theorem 5.2 states our result for this setting.

THEOREM 5.1. *Given any truthful mechanism \mathcal{M} , let $vx_j(v)$ be the expected value of the goods allocated to buyer j when his reported parameter is v , then by Myerson's characterization, his expected payment for valuation v is $vx_j(v) - \int_0^v x_j(v')dv'$. Furthermore, the expected revenue of \mathcal{M} from buyer j is $E_v[\varphi_j(v)x_j(v)]$.*

THEOREM 5.2. *Given any positive network externality setting, let β be the expected ratio of the restricted welfare to the optimal welfare for random partitioning, and α be the approximation factor of the allocation problem for the restricted welfare. Then there exists a BIC mechanism with an $\alpha\beta$ -approximation to the revenue.*

PROOF. (Sketch) For a given set of buyers' bids, partition the vertices into V_1 and V_2 , by assigning each vertex to either side w.p. $1/2$. Let $V_3 \subseteq V_1$ be the set of vertices with non-negative virtual valuations. Neglect vertices in $V_1 - V_3$ and solve the allocation problem on (V_3, V_1) on virtual valuations. The expected virtual welfare for the instance (V_3, V_2) is at least β factor of the optimal virtual welfare. We use the α -approximation algorithm for the restricted virtual welfare maximization. In this process, the allocation to buyers in V_2 is independent of their valuations, and they are assigned some item among the set of items they are interested in, while their payment is 0. Allocation for buyers in V_1 depends upon their valuations.

Let $v x_j(v)$ be buyer j 's expected welfare when his intrinsic valuation is v and he is assigned to side V_1 . If $x_j(v)$ is not monotone in v , then we make it monotone using the ironing technique from [Hartline and Lucier 2010]. As the distributions are regular, the virtual valuation for buyer j is monotone in his real valuation, the ironed allocation curve is monotone in his virtual valuations as well, and this step only increases the revenue of the mechanism. Let $\bar{x}_j(v)$ be the allocation curve after applying the ironing procedure.

Payments are computed as follows: buyer j 's payment is 0 when he is assigned to V_2 ; when he is assigned to V_1 , his payment for valuation v is $v\bar{x}_j(v) - \int_0^v \bar{x}_j(v')dv'$. Using Myerson's characterization, the mechanism is incentive compatible in Bayesian sense, and it remains to bound the revenue. As the expected virtual welfare of the allocation is at least $\alpha\beta$ fraction of the expected optimal virtual welfare, we get an $\alpha\beta$ approximation for the revenue. \square

5.2. Multidimensional Valuation: Complete Graph, IID and Non-Atomic Buyers

In this section, we consider the setting, when there are $n \rightarrow \infty$ unit demand buyers. Each buyer's valuation for item i is drawn independently from a known distribution with CDF G_i (and PDF g_i), furthermore each distribution is regular. The underlying graph is complete, so that the externality seen by a buyer is simply a fixed function of the fraction of all buyers who own the same item. We assume the linear model of externality in the rest of the section: if n_i is the number of buyers who own item i , and v_{ij} is the intrinsic valuation of buyer j for item i , his overall valuation is simply $v_{ij}n_i$, where the valuation v_{ij} is drawn from G_i .

The eventual mechanism is a posted pricing. When j th buyer in the sequence makes a decision about the purchase, he is aware about the allocation of items to previous buyers, and not to the buyers later in the sequence. This necessitates a need to define buyers' behavior. We assume that buyers are *completely risk averse*, and each buyer buys an item that maximizes his utility only considering the past allocation. Our techniques can be extended to a setting when buyers are *risk neutral* and buyers make purchase decisions based on the expected allocation in the future. Our result is stated in the following theorem.

THEOREM 5.3. *There exists an 16-approximation algorithm to the revenue of the optimal mechanism when buyers are IID, and the externality network is a complete graph.*

In the rest of the section, we prove Theorem 5.3. We first note an important property in the following lemma that lets us reduce the problem to a single dimensional setting, similar to Chawla et al [Chawla et al. 2010].

LEMMA 5.4. *Given an instance I of the given problem, reduce it to a single dimensional problem I_{copies} as follows: split each buyer into m copies, where the i th copy is interested only in item i , with a restriction that only one copy of a buyer can be allocated an item. Then the revenue of the optimal mechanism on I_{copies} is no less than the revenue of the optimal mechanism on \mathcal{I} .*

We shall denote the copies of buyer j by $j1, j2, \dots, jm$. Let \mathcal{M}_{OPT} be the optimal mechanism for I_{copies} . Given a buyer j in I , let $y_i(v)$ be the probability that ji is allocated an item when his valuation is v and let $p_i(v)$ be his expected payment when his valuation is v . Let $n_i = \sum_v y_i(v)g_i(v)$, i.e. the expected number of buyers that buy item i . By Myerson's characterization and taking expectation over all possible realizations, we have $\sum_v p_i(v)g_i(v) = n_i (\sum_v \varphi_i(v)y_i(v)g_i(v))$. The following quadratic program upper bounds the revenue of the optimal mechanism.

$$\begin{aligned} & \text{Maximize } n \times \left(\sum_{i,v} p_i(v)g_i(v) \right) \\ & \sum_i n_i \leq n \\ & \sum_v p_i(v)g_i(v) = n_i \times (\sum_v \varphi_i(v)y_i(v)g_i(v)) \quad \forall i \\ & n_i = n \times (\sum_v y_i(v)g_i(v)) \quad \forall i \\ & y_i(v) \in [0, 1] \quad \forall i, v \end{aligned}$$

As the distributions are regular, $\varphi_i(v)$ is non-decreasing in v for each item i . Thus the optimal solution to the convex program has the following property: for every i , y_i is either a step function such that $y_i(v) = 0$ for $v \leq v_i$ and $y_i(v) = 1$ for $v > v_i$, or it is a convex combination of two step functions with two thresholds $v_{i1} < v_{i2}$ such that v_{i1} is the largest valuation less than v_{i2} . The value of v_i (or v_{i1}, v_{i2}) can be considered as the price for the item i .

Solving the quadratic program: Using the characterization of the optimal solution, given a value of n_i , the values of $y_i(v)$ and $p_i(v)$ are uniquely determined. Hence it suffices to guess the values of n_i s. We use a DP with the following state space: $A[i, n']$ is a solution to the convex program with n' buyers and first i items. We solve the DP considering the number of buyers in increments of $\frac{\epsilon^3 n}{m^2}$. We now establish that this DP gives a $(1 + O(\epsilon))$ -approximate solution: in the optimal solution, for each item i with $n_i \geq \frac{\epsilon^2 n}{m^2}$, we round down n_i to the nearest multiple of $\frac{\epsilon^3 n}{m^2}$. We then arbitrarily pick an item with more than $\frac{\epsilon n}{m}$ buyers allocated to it (if such item exists), and remove $\frac{\epsilon^2 n}{m}$ buyers from it. These steps reduce the revenue of the optimal allocation by a factor of $O(\epsilon)$. For each item i with $n_i < \frac{\epsilon^2 n}{m^2}$ buyers, we compute the best allocation up to $\frac{\epsilon^2 n}{m^2}$ buyers. In this step, the revenue can only increase, and the number of extra buyers required in this process is $\frac{\epsilon^2 n}{m}$, which is compensated by the previous step.

The pricing mechanism is given below, and the approximation bound follows from Lemma 5.5.

- We offer free items to first $\sum n_i/2$ buyers: first $n_1/2$ buyers are given item 1 for free and other items are withheld. Next $n_2/2$ buyers are offered item 2 for free and other items are not shown to them, and so on. We note that $\sum_i n_i/2 \leq n/2$. We skip all buyers between $\sum_i n_i/2 + 1$ to $n/2$.
- For buyer j between $n/2 + 1$ and n , we offer items as follows: let n'_i be the number of buyers that already have item i . Then for each item i , show item i to buyer j independently w.p. $1/2$ and ask a price of $(n'_i + 1)v_i$.

LEMMA 5.5. For any j between $n/2 + 1$ and n , the expected revenue generated from the j th buyer is at least $\frac{\text{OPT}}{8n}$.

PROOF. The expected number of items for which buyer j has valuation more than the price is exactly $\frac{\sum_i v_i}{2n} \leq \frac{1}{2}$. Hence w.p. $\frac{\sum_i v_i}{2n}$, he has valuation for at least one item more than its price. The probability that he is eligible for at least two items is at most $\frac{\sum_i v_i}{4n}$. Thus w.p. at least $\frac{\sum_i v_i}{4n}$, he is eligible for exactly one item. Furthermore, when he buys the item, his price is at least $1/2$ of the quadratic program price for the item. This completes the proof. \square

We note that, our technique can be modified to get a 16-approximation for concave externality and $O(2^d)$ - approximation for convex externality.

REFERENCES

- AKHLAGHPOUR, H., GHODSI, M., HAGHPANAH, N., MIRROKNI, V. S., MAHINI, H., AND NIKZAD, A. 2010. Optimal iterative pricing over social networks (extended abstract). In *WINE*. 415–423.
- ALAEI, S. 2011. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*. 512–521.
- BHATTACHARYA, S., GOEL, G., GOLLAPUDI, S., AND MUNAGALA, K. 2010. Budget constrained auctions with heterogeneous items. In *STOC*. 379–388.
- BHATTACHARYA, S., KULKARNI, J., MUNAGALA, K., AND XU, X. 2011. On allocations with negative externalities. In *WINE*. 25–36.
- BLUMROSEN, L. AND NISAN, N. 2007. *Algorithmic Game Theory*. Cambridge University Press, Chapter Combinatorial Auctions.
- CAI, Y. AND DASKALAKIS, C. 2011. Extreme-value theorems for optimal multidimensional pricing. In *FOCS*. 522–531.
- CALINESCU, G., CHEKURI, C., PÁL, M., AND VONDRÁK, J. 2007. Maximizing a submodular set function subject to a matroid constraint (extended abstract). In *IPCO*. 182–196.
- CALINESCU, G., KARLOFF, H. J., AND RABANI, Y. 2004. Approximation algorithms for the 0-extension problem. *SIAM J. Comput.* 34, 2.
- CANDOGAN, O., BIMPIKIS, K., AND OZDAGLAR, A. 2010. Optimal pricing in networks with externalities. *Submitted*.
- CHAKRABORTY, T., HUANG, Z., AND KHANNA, S. 2009. Dynamic and non-uniform pricing strategies for revenue maximization. In *FOCS*. 495–504.
- CHAWLA, S., HARTLINE, J. D., AND KLEINBERG, R. D. 2007. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*. 243–251.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *STOC*. 311–320.
- CHEKURI, C., KHANNA, S., NAOR, J. S., AND ZOSIN, L. 2001. Approximation algorithms for the metric labeling problem via a new linear programming formulation. In *SODA*.
- FEIGE, U. AND VONDRÁK, J. 2010. The submodular welfare problem with demand queries. *Theory of Computing* 6, 1, 247–290.
- GAIRING, M., MONIEN, B., AND TIEMANN, K. 2006. Routing (un-) splittable flow in games with player-specific linear latency functions. In *ICALP*. 501–512.
- HAGHPANAH, N., IMMORLICA, N., MIRROKNI, V. S., AND MUNAGALA, K. 2011. Optimal auctions with positive network externalities. In *ACM Conference on Electronic Commerce*. 11–20.
- HARTLINE, J. D. AND LUCIER, B. 2010. Bayesian algorithmic mechanism design. In *STOC*. 301–310.
- HARTLINE, J. D., MIRROKNI, V. S., AND SUNDARARAJAN, M. 2008. Optimal marketing strategies over social networks. In *WWW*. 189–198.
- KANN, V. 1991. Maximum bounded 3-dimensional matching is MAX SNP-complete. *IPL* 37, 27–35.
- KLEINBERG, J. M. AND TARDOS, É. 1999. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields. In *FOCS*. 14–23.
- LEHMANN, B., LEHMANN, D. J., AND NISAN, N. 2001. Combinatorial auctions with decreasing marginal utilities. In *ACM Conference on Electronic Commerce*. 18–28.