

Differential Pricing with inequity aversion in social networks

Noga Alon, Tel Aviv University and Microsoft Research

Yishay Mansour, Tel Aviv University and Microsoft Research

Moshe Tennenholtz, Microsoft Research and the Technion - Israel Institute of Technology

We introduce and study the algorithmic problem of maximizing revenue in a network using differential pricing, where the prices offered to neighboring vertices cannot be substantially different. Our most surprising result is that the optimal pricing can be computed efficiently, even for arbitrary revenue functions. In contrast, we show that if one is allowed to introduce discontinuities (by deleting vertices) the optimization problem becomes computationally hard, and we exhibit algorithms for special classes of graphs. We also study a stochastic model, and show that a similar contrast exists there: For pricing without discontinuities the benefit of differential pricing over a single price is negligible, while for differential pricing with discontinuities the difference is substantial.

Categories and Subject Descriptors: F.2.2 [Nonnumerical Algorithms and Problems]

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Pricing, Algorithmic Game theory, Social Networks

1. INTRODUCTION

Differential pricing [19; 3] is the practice in which identical goods or services are sold in different prices to different customers. Some types of differential pricing are direct, and include providing special offers to new buyers in promotional sales, and special deals to loyal buyers or to another selected part of the population. Other types of differential pricing are less direct, and include techniques of quality differentiation and versioning, e.g., selling a book in hard-cover and soft-cover formats with corresponding prices.

Differential pricing may lead to externalities. While differential pricing is a general technique in economics and marketing, externality is a general term that refers to costs or benefits for a customer that was not involved in an actual economic transaction. For example, a customer might be convinced about the quality of a product due to the fact another customer bought it. While this is an example of a positive externality, one may find also a negative externality in some economic trades, e.g., a purchase of a fancy apartment by a criminal might badly effect the value of his neighbors' apartments.

In this paper we deal with externalities of differential pricing in a social network. A social network is a general setting consisting of a set of individuals and (typically binary) relations among them, e.g., a social network such as Twitter is modeled as a directed graph and a social network such as Facebook is modeled as an undirected graph. Notions of externality are central to social networks. For example, models of diffusion in social networks (see, e.g., [15]) are based on the assumption that there are positive externalities in product purchase (or in opinion formation) between neighbors (or friends) in a social network.

Noga Alon: Tel Aviv University and Microsoft Research, Herzlia, nogaa@tau.ac.il.

Yishay Mansour: Tel Aviv University and Microsoft Research, Herzlia, mansour@tau.ac.il.

Moshe Tennenholtz: Technion-Israel Institute of Technology and Microsoft Research, Herzlia, moshet@microsoft.com.

Permission to make digital or hardcopies 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 credits 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'13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

The above concepts of differential pricing, externalities, and social networks do not live in isolation from one another. For example, one may consider the idea of finding influential members in a social network, and by providing them appropriate incentive for product purchase influence purchases of other members. This approach tries to exploit positive externality in a social network using differential pricing. The idea would be to optimize a marketer revenue (see, e.g., [13]).¹ Another example is mechanism design for buyers with positive externalities [12]. However, negative externalities may also appear when differential pricing in social networks context is used. For example, differential pricing preferring new customers upon their friends who are loyal customers have resulted in surprisingly severe negative ramifications when employed by Amazon in the past (see [10] for this and related observations).

In this paper we introduce the first algorithmic study of differential pricing in a social networks setting, where the aim is to avoid a well known type of negative externality, titled inequity aversion (see, e.g., [9; 4]). Inequity aversion captures the need to avoid situations that a member in a social network will have significantly better deal than his peers in a social network. Inequity aversion is a most basic form of negative externality, to be distinguished from spiteful behavior, as in [16; 6].

Our motivation applies equally to “physical networks”. Consider a network of different stores of a large retail chain. The graph may represent the possibility of consumer to move from one store to the other. In this abstraction the nodes are stores and edge is proximity. The differential pricing allows for different prices at different stores. The constraint that the prices of neighboring stores should be similar is motivated by the desire of reducing consumer migration between the stores.

Our abstract model is the following. We have an underlying graph, which represents the relationships between the nodes (e.g., users). Each node has a revenue function that maps a price to a revenue. The constraint on the pricing is that neighboring nodes should receive similar offer prices. Our goal is to find a feasible set of prices that maximizes the total revenues.

For this basic model we derive an efficient algorithm to compute the revenue maximizing prices. We also show that in an expander graph, the fact that the difference between neighboring nodes is small implies that any feasible pricing is concentrated around a single price, and therefore the difference between the best single price and the differential pricing is small.

We extend our model by allowing discontinuities, by essentially ‘deleting’ nodes from the graph (and not receiving any revenue from them). This captures the situation in which offers are not made to some of the customers (and therefore no one can envy them). We show that the problem of revenue maximization in this model is computationally hard. For arbitrary revenue functions it is hard to approximate within $n^{1-\epsilon}$ for an n vertex graph, and even for simple revenue functions the exact optimization remains NP-complete. We give polynomial time algorithms for trees and bounded tree-width graphs. For graphs with small degrees or small separators we give approximation algorithms.

We further consider a simple stochastic model, where each node has a value drawn i.i.d. from a distribution over $[0, 1]$. This captures the situation where we do not have accurate information about the willingness to pay of the customers when we come up with offers. The revenue from a node is the price, if the price is below its value, and zero otherwise. We allow the prices to change by at most $\epsilon \geq 0$ between neighboring nodes. We show that in the basic model, for any connected graph and any value distribution the difference between the best single price and the differential pricing is small (vanishing as a function of ϵ). In contrast, we show that for a line graph with a uniform value distribution, there is

¹The economics literature focused on the other hand on the effects of differential pricing on social welfare, although not in the context of social networks [18; 17].

a linear gap between the best single price and the expected revenue of the optimal pricing with discontinuities, and this holds even for $\epsilon = 0$. Our method of computing the expected revenue of the optimal pricing is interesting in its own right. We use a dynamic program that computes the optimal pricing given a realization of the valuations. We model it as a Markov chain with an unbounded number of states, and then reduce it to a constant size Markov chain with rewards, for which we can compute the expected reward (which is identical to the average expected optimal revenue of our original problem).

The paper is organized as follows. Section 2 introduces our models for differential pricing in a network. The basic model and its related results are presented in Section 3. The extension to handle discontinuities is studied in Section 4. The stochastic model is addressed in Section 5. The final short section contains some concluding remarks and open problems.

2. MODEL

We are given a graph $G(V, E)$ and a finite set $P = \{0, 1, \dots, \ell\}$ of prices. For every node $v \in V$ there is an associated revenue function R_v that maps an offered price $p_v \in P$ assigned to node v to the revenue gained. (The representation of R_v is an array with ℓ entries.) The *total revenue* of a price vector p is $R(p) = \sum_{v \in V} R_v(p_v)$.

Without any constraint on the vector of prices p , the problem of maximizing the total revenue, is simply maximizing each R_v separately. The problem becomes more challenging when there are constraints imposed on the price vector p .

Model I: Bounded differences pricing. There is a difference vector α which assigns a non-negative number to each edge, i.e., $\alpha(v, w) \geq 0$ for $(v, w) \in E$. A price vector p is feasible with respect to a graph G and a difference vector α if for every edge $(v, w) \in E$ we have that $|p_v - p_w| \leq \alpha(v, w)$. The goal is to compute a feasible price vector p which maximizes the total revenue, i.e., $R(p) = \sum_{v \in V} R_v(p_v)$. Since the prices are integers, we will assume that also $\alpha(v, w)$ are integers.

Model II: Bounded differences pricing with discontinuities. This is an extension of the bounded differences pricing model, where we add a special price $p_v = \perp$ to P . By definition, $R_v(\perp) = 0$. In addition, in case $p_v = \perp$, there is no price restriction on edges incident with node v , i.e., we require $|p_v - p_w| \leq \alpha(v, w)$ only for $p_v, p_w \neq \perp$.

Model III: Fixed prices with discontinuities. This is a special case of the bounded differences pricing with discontinuities model, where $\alpha(v, w) = 0$ for any $(v, w) \in E$.

While we allow arbitrary revenue functions, in general, we sometimes focus on the interesting restricted class of *single value revenue function*. A revenue function R is a *single value revenue function* if for each node v there exists a value $val(v)$ such that for $p > val(v)$ we have $R_v(p) = 0$ and for $p \leq val(v)$ we have $R_v(p) = p$. We will also sometimes focus on the natural special case of a *uniform difference vector* α , assigning to each edge the same value.

3. BOUNDED DIFFERENCES PRICING

The main result of this section is an efficient algorithm that maximizes the total revenue in the bounded differences pricing model.

To start, assume that the revenue functions $R_v(\cdot)$ are concave and continuous². Then we can write the following concave program (which can be solved efficiently [5]).

²For simplicity, for the convex program we will assume that the prices are also continuous, i.e., $P = [0, \ell]$.

ConvexProg

$$\max \sum_{v \in V} R_v(p_v)$$

such that

$$\forall (u, v) \in E : p_v - p_w \leq \alpha(v, w) \text{ and } p_w - p_v \leq \alpha(v, w).$$

THEOREM 3.1. *If the revenue functions $R_v(\cdot)$ are concave then there is a polynomial time algorithm to compute the optimal prices.*

Our main result is showing a polynomial time algorithm for computing the optimal revenue for arbitrary revenue functions and an arbitrary graph. Our methodology would be the following. First, we present a fractional solution for the problem. Second, we show that the optimal solution is integral, and thus the problem can be solved efficiently.

We present a model for the fractional setting. Recall that we assume that the prices are discrete. For each node v we have $R_v(i)$ as the revenue when the price is $i \in P$ and $p_v(i)$ is the fraction given to price i at node v . This implies that p_v is a vector of length ℓ , such that $p_v(i) \geq 0$ and $\sum_{i=1}^{\ell} p_v(i) = 1$.

We would like to write a linear program for the bounded differences pricing that computes an optimal fractional solution. The first step is to show that we can check if two price distributions p_v and p_w are *compatible*. For deterministic prices p_v and p_w , they are compatible if the difference is at most $\alpha(v, w)$, i.e., $|p_v - p_w| \leq \alpha(v, w)$. In the fractional model we would like to be able to “move” prices to distance at most $\alpha(v, w)$ and transform p_v to p_w . Formally, two price vectors p_v and p_w are *matrix compatible* if there exists a stochastic $\ell \times \ell$ matrix $M_{v,w}$ such that: (1) if $|i - j| > \alpha(v, w)$ then $M_{v,w}[i, j] = 0$ and (2) $M_{v,w}p_v = p_w$.

An equivalent requirement is the following. Price vectors p_v and p_w are *threshold compatible* if for each k we have $\sum_{i=1}^{k+\alpha(v,w)} p_v(i) \geq \sum_{i=1}^k p_w(i)$ and $\sum_{i=1}^{k+\alpha(v,w)} p_w(i) \geq \sum_{i=1}^k p_v(i)$.

CLAIM 3.2. *Price vectors p_v and p_w are threshold compatible if and only if they are matrix compatible.*

PROOF. Clearly if price vectors p_v and p_w are matrix compatible they are threshold compatible. For the other direction, we use a max-flow problem on a bipartite graph. Define a bipartite graph $B(L_v, R_w)$, where $L_v = P = R_w$. There is an edge from $i \in L_v$ to $j \in R_w$ if $|i - j| \leq \alpha(v, w)$ and each such edge has infinite capacity (capacity 1 would be enough). Add a source node s and a sink node t , and connect the source node s to all nodes in L_v , where the capacity of (s, i) is $p_v(i)$. Similarly, connect all nodes in $j \in R_w$ to the sink node t where the capacity of (j, t) is $p_w(j)$.

Clearly, price vectors p_v and p_w are matrix compatible iff there is a flow of magnitude 1 in B from s to t . Consider a minimum cut in B . Assume that the cut has value strictly less than 1. Clearly the cut cannot have any edges between L_v and R_w . Thus the cut define subsets of nodes $S_v \subset L_v$ and $S_w = N(S_v)$, where $N(S) = \{j : (i, j) \in E \text{ for some } i \in S\}$. Since the cut has value less than 1 then $\sum_{i \in S_v} p_v(i) > \sum_{j \in S_w} p_w(j)$. It is not difficult to see that we may assume that both S_v and S_w are intervals, since if S_v is a union of intervals the inequality must hold for one of them and its set of neighbors. Let $S_v = [\beta_1, \beta_2]$ and $S_w = [\beta_1 - \alpha(v, w), \beta_2 + \alpha(v, w)]$. From the threshold constraints for $k = \beta_1 - \alpha(v, w) - 1$ we have that $\sum_{i=1}^{\beta_1-1} p_v(i) \geq \sum_{i=1}^{\beta_1-\alpha(v,w)-1} p_w(i)$. Since $\sum_{i=\beta_1}^{\beta_2} p_v(i) > \sum_{i=\beta_1-\alpha(v,w)}^{\beta_2+\alpha(v,w)} p_w(i)$ then $\sum_{i=1}^{\beta_2} p_v(i) > \sum_{i=1}^{\beta_2+\alpha(v,w)} p_w(i)$. Therefore p_v and p_w are also not threshold compatible. \square

Given the threshold compatibility, define the following linear program:

LinearProg1

$$\begin{aligned} & \max \sum_{v \in V} R_v(i) p_v(i) \\ & \text{such that} \\ & \forall (v, w) \in E \quad \forall k \in P: \sum_{i=1}^{k+\alpha(v,w)} p_v(i) \geq \sum_{i=1}^k p_w(i) \text{ and } \sum_{i=1}^{k+\alpha(v,w)} p_w(i) \geq \sum_{i=1}^k p_v(i) \\ & \forall v \in V \quad \forall k: p_v(k) \geq 0. \\ & \forall v \in V \quad \sum_k p_v(k) = 1. \end{aligned}$$

The first immediate claim is that the linear program provides an upper bound of the optimal value.

CLAIM 3.3. *The linear program LinearProg1 finds a fractional solution which is an upper bound on the optimal bounded differences pricing.*

Our surprising result is that there is *no integrality gap*. We show a simple way to round a fractional solution, and prove that the resulting integral solution is optimal.

Given an optimal fractional solution $p_v(i)$, let the price $price(v)$ at node v be the smallest non-zero price in $p_v(i)$. Formally, let $price(v) = \min\{i : p_v(i) > 0\}$. First we show that the resulting set of prices is feasible.

CLAIM 3.4. *For every edge $(v, w) \in E$ we have that $|price(v) - price(w)| \leq \alpha(v, w)$.*

PROOF. For contradiction, assume that it does not hold for edge $(v, w) \in E$ and $price(v) < price(w) - \alpha(v, w)$. Let $i = price(v)$. It implies that we have $p_v(i) > 0$ and for any $j \leq i + \alpha(v, w)$ we have $p_w(j) = 0$. This implies that p_v and p_w do not obey the edge condition with $k = i$, which contradicts that p_v and p_w are part of a feasible solution. \square

THEOREM 3.5. *The prices $price(v)$ are an optimal solution of the bounded differences pricing.*

PROOF. Claim 3.4 shows that $price(v)$ is a feasible solution. For contradiction, assume that $price(v)$ has strictly lower value than the fractional price vectors $p_v(i)$. We will show that we can “delete” the $price(v)$ solution and remain with a feasible solution of a higher value, contradicting the fact that we started with an optimal solution.

Let $\epsilon > 0$ be smaller than any non-zero value in the fractional solution, i.e., $\epsilon < \min_{i, v: p_v(i) > 0} p_v(i)$. Then we can write the fractional solution as a convex combination of an ϵ fraction of the $price(v)$ solution and a $1 - \epsilon$ fraction of the remaining solution (after subtracting $price(v)$). Formally, define p'_v such that $p'_v(i) = (p_v(i) - \epsilon I\{i = price(v)\}) / (1 - \epsilon)$. The new solution $p'_v(i)$ is feasible, since we are subtracting ϵ from both sides of the edge constraints (here it is important that $price(v)$ and $price(w)$ are the minimal non-zero entry). Therefore, p'_v has a strictly higher value than p_v , which is a contradiction to the optimality. \square

We have established the following.

COROLLARY 3.6. *There is a polynomial time algorithm that computes the optimal bounded differences pricing.*

3.1. Totally Unimodular

Here is an alternative proof that the bounded differences pricing has no integrality gap. The main benefit of the alternate proof is that it captures a much wider family of feasibility constraints, and allows to extend the basic model, while maintaining the fact that the optimal solution is computable in polynomial time.

Recall that a matrix A is called *totally unimodular* if the determinant of any square submatrix of it is either 0, or 1 or -1 . It is an easy corollary of Cramer’s rule that any vertex of a polytope given by $Ax \leq b$ for an integral vector b and such a matrix A has

integral coordinates. In our case the LP we consider has variables we denoted by $p_v(k)$ for $v \in V$ where $G = (V, E)$ and where $k = 1, 2, \dots, \ell$, for a total of $n\ell$ variables. The defining inequalities are

- (1) $\forall (v, w) \in E \quad \forall k: \sum_{i=1}^{k+\alpha(v,w)} p_v(i) - \sum_{i=1}^k p_w(i) \geq 0.$
- (2) $\forall (v, w) \in E \quad \forall k: \sum_{i=1}^{k+\alpha(v,w)} p_w(i) - \sum_{i=1}^k p_v(i) \geq 0.$
- (3) $\forall v \in V \quad \forall k: p_v(k) \geq 0.$
- (4) $\forall v \in V \quad : \sum_k p_v(k) = 1.$

We need the following theorem of Ghouila Hourri (1962) [11].

THEOREM 3.7. *A matrix A is totally unimodular if and only if for any nonempty subset of its columns there is a linear combination of these columns with coefficients $\{-1, 1\}$ so that the resulting vector is a vector of $\{0, 1, -1\}$.*

We will show that this is the case for our matrix, denoted by A . The rows of the matrix are the inequalities and the columns are the variables. We can think of the variables (columns) as partitioned to n blocks, one per node $v \in V$ and in each block the ℓ variables $p_v(k)$ in an increasing order of k . It is worthwhile to consider the structure of the complete matrix. Any inequality of type (1) and (2) has a prefix of $+1$ in one block and a prefix of -1 in another block. Any inequality of type (3) has a single $+1$ and an inequality of type (4) has the entire block $+1$.

Given a nonempty subset S of the set of columns, assign the columns signs as follows: the columns of S in each block get alternating signs starting with $+1$, that is: $+1, -1, +1, -1, \dots$ (In the next block we start again with $+1$). Let us consider the different inequalities. For an inequality of type (3), we clearly have the result either $+1$ or -1 , if the relevant column is in S , since it had a single $+1$, and otherwise the result is 0. For an inequality of type (4) we have a block of $+1$, so when we multiply by an alternating $+1$ and -1 sequence the sum is either $+1$ or 0. For an inequality of type (1) or (2), the sum over the $+1$ prefix is either $+1$ or 0 (same as for any inequality of type (4)). The sum over the -1 prefix is either -1 or 0. This implies that the sum of the values is either $+1$, -1 or 0. Therefore, we have established the following theorem.

THEOREM 3.8. *The matrix A is totally unimodular.*

Remark: Using the proof based on the totally unimodular matrix we can show many interesting extensions to our bounded differences pricing model.

- (1) We can add equalities of the form $\sum_{i=\beta_v}^{\gamma_v} p_v(i) = 0$ (or $= 1$) that is, take some consecutive values in one block only and forbid prices there (or insist the prices are there) and the same proof shows that's totally unimodular too.
- (2) We can add constraints of the form $p_v(i) = 0$, for any v and any price i we wish, as this clearly keeps the matrix totally unimodular. This means we can actually force the available subset of allowed set prices for each vertex v to be any subset of P we wish.
- (3) So far we assumed that α is symmetric, i.e., $\alpha(v, w) = \alpha(w, v)$. We can extend the model to non-symmetric α and achieve the same result. We can have a directed graph $G(V, E)$ and have: $\forall (v, w) \in E \quad \forall k: \sum_{i=1}^{k+\alpha(v,w)} p_v(i) - \sum_{i=1}^k p_w(i) \geq 0$, and the same proof that the matrix is totally unimodular still holds.

3.2. Expanders

Expander graphs are a natural class of graphs, and a random graph (with some mild assumptions on the parameters used) is an expander graph with high probability. We will be interested in understanding the resulting bounded difference pricing on this natural class of

graphs, where in this section we consider a uniform α with $\alpha(u, v) = 1$ for every edge, and single revenue value functions.

The main observation here is that since the function assigning prices to vertices is a Lipschitz function, then almost all the vertices have almost the same price, i.e., the function is nearly a constant. This implies that for expander graphs, the feasibility constraint forces almost all nodes to have a very similar price.

Definition 3.9. A graph $G(V, E)$ has *expansion* λ if for any set $S \subset V$ such that $|S| \leq |V|/2$ we have $|N(S)| \geq (1 + \lambda)|S|$, where $N(S) = \{v | \exists w \in S, (w, v) \in E\}$.

THEOREM 3.10. *Given a graph $G(V, E)$ with n vertices and expansion λ , any assignment of integers to nodes, i.e., $p : V \rightarrow Z$, such that $|p(v) - p(w)| \leq 1$ if $(v, w) \in E$, satisfies the following: There is a value i so that for any $\epsilon > 0$ all nodes v but at most ϵn have $p(v) \in [i - k, i + k]$ where $k = O(\frac{1}{\lambda} \ln \frac{1}{\epsilon})$.*

PROOF. Let i be the median value of the function p , that is, at least $n/2$ vertices satisfy $p(v) \geq i$ and at least $n/2$ vertices satisfy $p(v) \leq i$.

Fix $\epsilon > 0$, let S_l be the set of nodes with prices less than $i - k$ and let S_h be the set of nodes with prices more than $i + k$, where k will be determined later. Let $N^j(S_l)$ denote the set of all vertices of distance at most j from S_l . Thus $N^0(S_l) = S_l$. By the expansion assumption, if $|N^j(S_l)| \leq n/2$, then $|N^{j+1}(S_l)| \geq (1 + \lambda)|N^j(S_l)|$. It thus follows that for $k = \lceil \frac{\ln 1/\epsilon}{\ln(1+\lambda)} \rceil$, if $|S_l| \geq \epsilon n/2$, then $|N^k(S_l)| \geq n/2$. But this is impossible, as all vertices in $N^k(S_l)$ must have prices below i , which is the median. It follows that $|S_l| \leq \epsilon n/2$. The same reasoning implies that $|S_h| \leq \epsilon n/2$, providing the required result. \square

From a pricing perspective, the simple result above implies that a constant price provides a good approximation for single value revenue functions on expanders.

COROLLARY 3.11. *Assume $G(V, E)$ is an expander with expansion λ , and for any $v \in V$ the revenue R_v is a single value revenue function over the prices $\{1, \dots, \ell\}$, with $\alpha(u, v) = 1$ for every edge (u, v) . Let OPT be the optimal value of the bounded differences prices. Then, there exists a single price p such that*

$$OPT - \sum_{v \in V} R_v(p) = O\left(\frac{n}{\lambda} \ln \ell\right)$$

PROOF. From Theorem 3.10 we have,

$$OPT - \sum_{v \in V} R_v(p) \leq \epsilon n \ell + O\left(\frac{n}{\lambda} \ln \frac{1}{\epsilon}\right).$$

Setting $\epsilon = 1/\ell$ gives the desired result. \square

4. BOUNDED DIFFERENCES PRICING WITH DISCONTINUITIES

In contrast to the case of bounded differences pricing, where we showed a polynomial time algorithm, when we add discontinuities the problem becomes hard (NP-complete). We show the hardness in Section 4.1. We give exact algorithms for special classes of graph (trees and graphs of bounded tree width) in Section 4.2 and approximation algorithms (for bounded degree and graphs with small separators) in Section 4.3.

4.1. Hardness

LEMMA 4.1. *For any graph $G(V, E)$, there is a set of revenue functions, such that for $\alpha(u, v) = 1$ for every edge, the total revenue of the optimal bounded differences pricing with discontinuities is equal to the maximum size of an independent set in $G(V, E)$.*

PROOF. We need to define the revenue functions to define the reduction. For node v_i let revenue $R_{v_i}(i \cdot (2n)) = 1$ and for any other price $R_{v_i}(p) = 0$.

This implies that in any connected component (of the graph obtained after the deletion of all edges (v, w) where either $p_v = \perp$ or $p_w = \perp$) we can get at most a revenue of 1. (If nodes v_i and v_j are connected by a path, and node v_i has revenue 1, then $p_{v_i} = i \cdot (2n)$. The price in node v_j has to be in the range $[i \cdot (2n) - n, i \cdot (2n) + n]$ which does not include the price $j \cdot (2n)$ and implies that $R_{v_j} = 0$.)

This implies that the maximum total revenue is exactly the size of the maximum independence set. \square

Using the well known results on approximability of independent set [14], we get the following.

THEOREM 4.2. *Computing the optimal bounded differences pricing with discontinuities in an n vertex graph is NP-complete, and is hard to approximate within a factor of $n^{1-\epsilon}$ assuming that NP does not have polynomial time randomized algorithms.*

The above reduction uses fairly unnatural revenue functions. We would like to establish the hardness for single value revenue function. We are able to show the hardness for a slightly larger class which includes in addition the constant revenue function, i.e., $R_v(p) = c$ for any $p \neq \perp$.

THEOREM 4.3. *The problem of computing the optimal bounded differences pricing with discontinuities is NP-complete even for revenue functions which are either single value or constant.*

PROOF. The reduction is from the Multiway (or Multiterminal) cut problem, which is the following: Given a graph $G(V, E)$ and a specified set S of k vertices in it, find the minimum number of edges needed to separate all vertices of S , that is, the minimum number of edges that can be deleted so that no two of the vertices of S stay in the same connected component of the resulting graph. This is known to be NP-hard even for $k = 3$, see [8].

We proceed with a reduction to our problem. Let $G(V, E)$ and $S = \{v_1, \dots, v_k\}$ be an instance of the multiway cut problem. Put $|V| = n$.

Construct a new input graph H for our problem: this is the graph obtained from G by replacing each edge e of G by a path of length 2 in which the middle vertex is denoted by e (this is sometimes called the 1-subdivision of G). Thus H has $|V| + |E|$ vertices and $2|E|$ edges. We demand prices assigned to neighboring vertices to differ by at most 1, that is, our difference vector α is 1 for every edge.

We next define the revenue functions for each vertex of H . For each vertex e that came from an edge of G the revenue is 1 for any prescribed price (besides \perp which gives revenue 0, as usual). For every vertex of H that came from a vertex $v \in V \setminus S$ the revenue is n^2 for every prescribed price (and 0 for \perp). Note that n^2 is larger than $|E|$, hence it never pays to assign \perp to such a vertex, as one can always delete all vertices of the form e , lose only $|E|$ from the total potential revenue, and make every vertex of H in what's left isolated, getting the maximum possible revenue for it. For each vertex $v_i \in S$ define $r_i = 10i \cdot n^2$. For price $p_i \leq r_i$ the revenue is p_i and otherwise it is zero.

Note that one possible choice of prices is to remove (that is, assign \perp to the members of) a set of vertices of H that corresponds to a minimum set of edges of G that separates the members of S . We can then assign the component of v_i in what's left price $r_i = 10i \cdot n^2$. It is not difficult to check that this is in fact the optimum revenue one can get. Indeed, if two vertices of S stay in the same component than we lose a lot for those, as the distance between them in H is less than $2n$ (since it is less than n in G). Similarly, if we assign \perp to a vertex of H corresponding to a vertex (and not an edge) of G then we lose too much and hence this cannot be done in an optimum solution. It follows that the maximum

possible revenue in our problem is exactly $|E| + n^2(|V| - k) + \sum_{i=1}^k 10in^2 - f$, where f is the minimum number of edges needed to be removed to separate all vertices of S .

This completes the proof. \square

4.2. Exact Algorithms

We show exact algorithms for special graphs.

THEOREM 4.4. *There is an $O(n\ell)$ time algorithm to compute the optimal bounded differences pricing with discontinuities when $G(V, E)$ is a tree.*

PROOF. We use a dynamic programming algorithm to find the optimal prices. Given the tree, we assign one node to be the root, and direct all the edges accordingly.

The dynamic programming algorithm, for each sub-tree T_v , rooted at v , for each possible price for v (including \perp), computes the optimal revenue of the subtree. Let $opt(T_v, k)$ be the revenue of the optimal assignment for T_v conditioned on $p_v = k$.

Given a node w , assume we computed the optimal assignment for each child subtree of T_w . Given a price $p_w = k$, a child v of w can have prices in the range $[k - \alpha(v, w), k + \alpha(v, w)]$ or \perp (i.e., at most $2\alpha(v, w) + 2$ prices). For each edge (w, v) and subtree T_v , the optimal revenue at v is $opt_{v,w}(k) = \max_{i \in [k - \alpha(v, w), k + \alpha(v, w)] \cup \{\perp\}} opt(T_v, i)$. Then $opt(T_w, k) = R_w(k) + \sum_{v: child(v, w)} opt_{v,w}(k)$, where $child(v, w)$ is derived from the tree structure. This computation is done from the leaves to the root. In the root r we simply maximize $\max_k R_r(k) + \sum_{v: (r, v) \in E} opt_{v,r}(k)$. The running time of the algorithm is $O(n\ell)$ provided $\alpha(u, v) \leq O(1)$ for every edge (u, v) . \square

We can extend the dynamic programming algorithm for trees to apply to bounded tree width graphs, and derive the following result.

THEOREM 4.5. *There is an $O(n^k \ell^k)$ time algorithm to compute the optimal bounded differences pricing with discontinuities when $G(V, E)$ is a graph of tree width k .*

4.3. Approximation

We show approximation algorithms for certain classes of graph. The first is using the maximum degree of the graph.

THEOREM 4.6. *For a graph $G(V, E)$ with maximum node degree D , the bounded differences pricing with discontinuities can be approximated with a multiplicative factor of $D + 1$.*

PROOF. For each node v , let $R_v^{max} = \max_p \{R_v(p)\}$ be the maximum revenue we can get from v (at some price) and $p_v^{max} = \arg \max R_v(p)$. Now we use the greedy algorithm. Take the node with the maximum R_v^{max} . Assign v the price p_v^{max} . Set all the neighbors w of v to $p_w = \perp$.

We claim that the benefit of the greedy algorithms is at least $\sum_{v \in V} R_v^{max} / (D + 1)$. The proof follows from the observation that when we assign a value p_v^{max} , then $R_v^{max} \geq R_w^{max}$ for any neighbor w of v which is not yet assigned a price. This implies that $D \cdot R_v^{max} \geq \sum_{w \in N(v)} R_w^{max}$. Therefore, the revenue we get is at least $\sum_v R_v^{max} / (D + 1)$. \square

THEOREM 4.7. *For any n -vertex graph $G(V, E)$ with a fixed excluded minor (and in particular for any planar graph), the bounded differences pricing with discontinuities can be approximated with an additive factor of ϵMn , where M is the maximum revenue that can be obtained from a single vertex and $\epsilon > 0$ is an arbitrary absolute positive constant.*

PROOF. Fix $\epsilon > 0$. As shown in [1], in any n -vertex graph with a fixed excluded minor one can eliminate (efficiently) ϵn nodes, and remain with connected components, each of size at most $O(1/\epsilon^2)$.

We can eliminate the ϵn nodes by assigning them prices of \perp . In each constant size connected component we can now find the precise optimum by exhaustively enumerating all possible price vectors (allowing the use of \perp here too). Clearly we outperform the optimal pricing in each component, and lose only in the pricing of the nodes eliminated. \square

5. STOCHASTIC MODEL

In this section we assume that each node v has a single value revenue function. Recall that this implies that node v has a value $val(v)$. The revenue from node v is $R_v(p)$ and is equal to p if $p \leq val(v)$ and 0 if $p > val(v)$. In the stochastic model we assume that $val(v)$ is distributed uniformly in $[0, 1]$. (This would be more convenient than having discrete prices and values.)

Given the distribution of $val(v)$, the expected revenue at price p is $p(1-p)$. Therefore, the single fixed price that maximizes, a priori, the revenue is $p = 1/2$ and its expected revenue is $1/4$. This will be the basic benchmark against which we will compare all our results.³ Our main goal is to study how much can be gained by changing the prices in either the bounded differences pricing model or the bounded differences pricing with discontinuities model. The end result is that in the bounded differences pricing model, very little is gained compared to a single fixed price. This holds for any connected graph and any distribution over $val(v)$, assuming that it is sampled independently and with the same distribution for all nodes. In contrast, allowing discontinuities creates a gap between the revenue of the optimal fixed price and the optimal fixed price with discontinuities.

5.1. Bounded differences pricing: line graph

The line graph has nodes $\{1, \dots, n\}$, where the edges are $(i, i + 1)$ for $1 \leq i \leq n - 1$. We assume that $\alpha(i, i + 1) = \epsilon$, hence the price at nodes i and $i + 1$ can differ by at most ϵ , i.e., $|p_{i+1} - p_i| \leq \epsilon$. We derive both a possibility result (lower bounding the average revenue) and an impossibility result (upper bounding the average revenue). Both bounds deviate from the revenue of the single best price by some function of ϵ that tends to 0 as ϵ tends to 0. Later we extend our results from a line graph to an arbitrary connected graph and from a uniform distribution to an arbitrary distribution. We will start with the *possibility result*.

CLAIM 5.1. For the line graph there is a pricing strategy for the bounded differences pricing that has expected total revenue of $0.25 + \Theta(\sqrt{\epsilon})$.

PROOF. Our bounded differences pricing would be based on the realized values of the nodes. We will partition the line to segments, and show that in each segment we can set bounded differences pricing, given the prices, such that the expected revenue is above 0.25.

We partition the line into segments of length $\ell = 3/\sqrt{\epsilon}$. Consider a single such segment. The probability that none of the ℓ nodes have values in the range $[0.5 - \sqrt{\epsilon}, 0.5 + \sqrt{\epsilon}]$ is at least $e^{-6-o(1)}$, which is a constant bounded away from 0. If the event does not hold then the fixed price $1/2$ has still an expected average payoff of 0.25. If the event holds, consider the following bounded differences pricing. We start at price 0.5 and increase by jumps of ϵ until $0.5 + \sqrt{\epsilon}$, using $\ell/3$ nodes, keep the price at $0.5 + \sqrt{\epsilon}$ for the next $\ell/3$, and return to price 0.5 using the remaining $\ell/3$ nodes. In this case the pricing has an expected average payoff of more than $0.25 + \sqrt{\epsilon}/6$ per node in the segment. Indeed, conditioned on having no value in $[0.5 - \sqrt{\epsilon}, 0.5 + \sqrt{\epsilon}]$, the probability of each vertex in the segment to have a value exceeding $0.5 + \sqrt{\epsilon}$ is exactly $1/2$, and when this happens for a vertex in the middle third of the segment, the revenue obtained is $0.5 + \sqrt{\epsilon}$. \square

For the *impossibility result* we will show that no pricing strategy has an expected revenue significantly above $1/4$ per node.

³Technically, we consider the average expected revenue by normalizing it per node.

THEOREM 5.2. *For the line graph, the expected optimal bounded differences pricing total revenue is at most $0.25 + O(\epsilon^{1/3})$ per node.*

PROOF. For the proof, partition the line into n/ℓ segments, each of size ℓ (to be determined later to be $\Theta(\epsilon^{-2/3})$).

Consider a realization of the values $val(i)$ in a segment, and let the optimal pricing for the realization be $opt(i)$. Note that by definition of bounded differences pricing $\max\{opt(i)\} - \min\{opt(i)\} \leq \epsilon\ell$. Therefore, if we select a single fixed price $p = \min\{opt(i)\}$ the revenue on the realization would decrease by at most $\epsilon\ell$ per node. We can further round this price to an integral multiple of $\epsilon\ell$, losing again at most $\epsilon\ell$ per node. This way, the total number of possible prices is only $\frac{1}{\epsilon\ell}$. Moreover, the difference between the rounded prices obtained for two consecutive segments is less than $2\epsilon\ell$, implying that knowing the price of a segment, there are at most 5 possible prices for the next segment. It follows that altogether there are only $\frac{1}{\epsilon\ell}5^{n/\ell}$ ways to choose the rounded prices for all the segments.

For each fixed choice of the prices, the expected value of the revenue over the whole line is at most $n/4$. Moreover, by the Chernoff-Hoeffding Inequality (c.f., e.g., [2], Appendix A), the probability of getting revenue at least $(1/4 + \lambda)n$ is at most $e^{-\lambda^2 n/3}$. Therefore, the probability that there is a pricing among the $\frac{1}{\epsilon\ell}5^{n/\ell}$ rounded ones considered above that gives total revenue exceeding $(1/4 + \lambda)n$ is at most

$$\frac{1}{\epsilon\ell}5^{n/\ell}e^{-\lambda^2 n/3}.$$

Taking, say, $\lambda^2 = 6/\ell$ we conclude that with very high probability none of these gives revenue exceeding $(1/4 + \lambda)n$. This provides a bound of $(1/4 + \lambda + 2\epsilon\ell)n$ for the optimum. The desired result is obtained by taking $\ell = \epsilon^{-2/3}$ (and hence $\lambda = \Theta(\epsilon^{1/3})$). \square

5.2. Bounded differences pricing: Arbitrary graphs

It is not hard to see that the impossibility result above extends to any connected graph. This can be done by extending the last proof for an arbitrary tree. It is shorter, however, to use the known result that the cube of any connected graph is Hamiltonian (see [7]). This means that for any connected graph G , the graph obtained by connecting two vertices of G iff the distance between them in G is at most 3, contains a Hamilton cycle (and hence also a Hamilton path). It follows that one can number the vertices of the original graph so that the distance between any two consecutive vertices is at most 3, and hence their prices can differ by at most 3ϵ . The following result thus follows from the result for the line.

COROLLARY 5.3. *For any connected graph $G(V, E)$ with uniform $\alpha = \epsilon$, the expected optimal bounded differences pricing total revenue per node is at most $0.25 + O(\epsilon^{1/3})$.*

5.3. Bounded differences pricing: Arbitrary distribution

In this section we consider an arbitrary distribution for the values $val(v)$. When comparing the expected revenue of the optimal fixed price and the expected revenue of the optimal bounded differences pricing, there is a sharp contrast between the ratio, which can be high, and the absolute difference, which will be low. For the ratio, we show that it is essentially unbounded (Theorem 5.4). For the absolute difference it will be a function of ϵ (Theorem 5.5).

THEOREM 5.4. *For any $k \geq 1$ and $\epsilon > 0$, there exists a distribution D^k over $[0, 1]$, such that for any single fixed price the expected revenue is at most $\epsilon^k/(1 - \epsilon)$ and the expected revenue of the optimal bounded differences pricing is $k\epsilon^k$, for any graph $G(V, E)$. Hence the ratio is at least $k(1 - \epsilon)$.*

PROOF. Let D^k be the following distribution. The support includes ϵ^i , for $1 \leq i \leq k$ and 0. The probability of ϵ^i is ϵ^{k-i} , and the remaining probability is on the value 0.

For any fixed price the expected revenue is at most $\epsilon^k/(1 - \epsilon)$. More specifically, the probability of valuations above ϵ^i is $\sum_{j=i}^k \epsilon^{k-j} = \epsilon^{k-i} \sum_{j=0}^{k-1} \epsilon^j$, and hence the expected revenue is at most $\epsilon^k/(1 - \epsilon)$.

Since the support of the distribution D^k is $[0, \epsilon]$, using bounded differences pricing we can get the sum of nodes valuation. This implies that the expected revenue is $k\epsilon^k$. \square

The following theorem shows the difference between the best fixed price and the optimal bounded differences pricing is small.

THEOREM 5.5. *For any distribution D over $[0, 1]$, for the line graph, the difference between the expected best fixed price and the expected optimal bounded differences pricing per node is at most $O(\epsilon^{1/3})$.*

The proof is very similar to the proof of Theorem 5.2 and is omitted.

5.4. Fixed prices with discontinuities: line graph

In this section we present an exact computation of the expected revenue for a line graph, for the fixed prices with discontinuities model. We consider a model with two equally likely values $\{1, 2\}$. Essentially we are back to the discrete model with $\ell = 2$ and $\Pr[\text{val}(v) = 1] = \Pr[\text{val}(v) = 2] = 1/2$. When considering a single fixed price, then the revenue of $p = 1$ is 1 (always sells at price 1) and the expected revenue of $p = 2$ is also 1 (sells with probability $1/2$ and gets revenue 2). The main challenge in this section is to compute the exact expected revenue of the optimal fixed price with discontinuities. The main result of the section would be to establish that the expected revenue with discontinuities is exactly $7/6$ per node, which means that it is significantly higher than the optimal fixed price with no discontinuities. This shows, once again, the significant difference between pricing with discontinuities and without them.

Before we compute the exact expectation let us observe that it has to be strictly larger than 1. Consider pricing all nodes at price 2, this has an expected revenue of 1 and with high probability the realized revenue is close to the expected one. Given a realization of values, consider any consecutive sequence of three or more values of 1. When we price all nodes at price $p = 2$ we are not receiving any revenue from any of those nodes. We can strictly increase our revenue by doing a discontinuity at both ends of the interval (offering a price of \perp) and selling to the other nodes in the interval at price 1. Since the probability of having an interval of size exactly k is 2^{-k-2} , and we increase the gain on it by $k - 2$, we have an expected revenue strictly more than 1. (Specifically, this strategy would give an expected revenue of $1 + \sum_{k=3}^{\infty} (k - 2)2^{-k-2} = 9/8$.) As we will see, however, this is not *the* optimal strategy of offering prices.

One can compute the optimal prices given a value realization using dynamic programming. For each node t and each possible price, we maintain the optimal revenue of the first t nodes, assuming node t is offered price p_t . More specifically, our set of prices is $\{1, 2, \perp\}$. It is easy to see that the optimal pricing will use only prices from this set. For each node t consider the triplet $(r_1^t, r_2^t, r_{\perp}^t)$ where value of r_p^t is the optimal revenue after t nodes and assuming that the price of node t is $p_t = p$. Given $\text{val}(v_{t+1})$ and $(r_1^t, r_2^t, r_{\perp}^t)$ we can compute $(r_1^{t+1}, r_2^{t+1}, r_{\perp}^{t+1})$ as follows. Let $r_{\perp}^{t+1} = \max_p \{r_p^t\}$, since we can offer node $t + 1$ price \perp and take any best pricing up to node t . Let $r_1^{t+1} = \max\{r_1^t, r_{\perp}^t\} + 1$ since when we offer price $r_{t+1} = 1$ we always get a revenue 1 from node $t + 1$, and we can select the best sequence leading to it, ending either in $p_t = 1$ or $p_t = \perp$. In case $\text{val}(v_{t+1}) = 2$ let $r_2^{t+1} = \max\{r_2^t, r_{\perp}^t\} + 2$, and otherwise (i.e., $\text{val}(v_{t+1}) = 1$) $r_2^{t+1} = \max\{r_2^t, r_{\perp}^t\}$. Clearly, given the realizations of $\text{val}(v_t)$ we can run this dynamic program in time $O(n)$ and

compute the optimal pricing for the realization. The optimal revenue of a given realization is $\max\{r_1^n, r_2^n, r_\perp^n\}$. Our main challenge would be to compute analytically the expected value of this dynamic program, i.e., $E[\max\{r_1^n, r_2^n, r_\perp^n\}]$.

When considering the dynamic program, one can consider the triplet $(r_1^t, r_2^t, r_\perp^t)$ as the state of the system. Given the realization $val(v_{t+1})$ we move to $(r_1^{t+1}, r_2^{t+1}, r_\perp^{t+1})$. This implies that from each state we have two possible next states, one for $val(v_{t+1}) = 1$ and one for $val(v_{t+1}) = 2$. This defines a Markov chain, where in each state, the two next states are equally likely. The problem is that we never get back to the same state, since the revenue is always increasing. For a line of length n the Markov chain has $\Theta(n)$ states, and we need to compute its expected revenue. We would like to reduce this large Markov chain to an equivalent finite one.

It is not hard to see that any two different revenues among the three quantities in $(r_1^t, r_2^t, r_\perp^t)$ cannot differ by more than 2. This observation follows, since we can modify any price sequence and end the sequence with a price \perp , reducing the revenue by at most 2. Given this observation, we can normalize the states by doing the following. Given a state $(r_1^t, r_2^t, r_\perp^t)$, we renormalize the state by having $R^t = \min\{r_1^t, r_2^t, r_\perp^t\}$, and the new state is $(r_1^t - R^t, r_2^t - R^t, r_\perp^t - R^t)$. (With an abuse of notation, we will denote the new states also by $(r_1^t, r_2^t, r_\perp^t)$.) Now we have a finite Markov chain, where the states are a subset of $\{0, 1, 2\}^3$ and in a transition we receive a reward R (which is also in $\{0, 1, 2\}$). This is now a much more manageable Markov chain. The specific resulting Markov chain is the following.

current state	next state $v = 1$	reward $v = 1$	next state $v = 2$	reward $v = 2$	average reward	steady distribution
0 0 0	1 0 0	0	1 2 0	0	0.0	0
1 0 0	2 0 1	0	1 1 0	1	0.5	1/12
1 1 0	2 0 1	0	1 2 0	1	0.5	2/12
1 2 0	2 0 2	0	0 2 0	2	1.0	2/12
2 0 1	2 0 1	1	1 1 0	2	1.5	3/12
2 0 2	1 0 0	2	1 2 0	2	2.0	1/12
0 2 0	1 0 2	0	0 3 1	1	0.5	1/12
0 3 1	1 0 2	1	0 3 1	2	1.5	1/12
1 0 2	1 0 0	2	1 2 0	2	2.0	1/12

The initial state of the Markov chain is $(0, 0, 0)$, and not all possible 27 states are reachable from it. In the table, for each state, we show the two next states, for each possible node value, and the reward.

Given a finite Markov chain, we can solve for the steady state distribution, which appears in the last column. One can easily verify that the claimed steady state is correct (checking the incoming probability to each state).

Now we are ready to compute the expected optimal revenue. This is simply the expected reward of the Markov chain, which we can compute, given the steady state distribution, as follows,

$$E[R] = 0.5 \frac{1}{12} + 0.5 \frac{2}{12} + 1.0 \frac{2}{12} + 1.5 \frac{3}{12} + 2.0 \frac{1}{12} + 0.5 \frac{1}{12} + 1.5 \frac{1}{12} + 2.0 \frac{1}{12} = \frac{14}{12} = \frac{7}{6}$$

We have established the following theorem.

THEOREM 5.6. *The expected revenue of the optimal fixed prices with discontinuities for a line graph is $7/6$ per node.*

An intriguing open problem is to study the case where the value is uniform in $[0, 1]$. While simulations show that the expected revenue is close to $1/e$, we were not able to analytically analyze this case, and leave it as an open problem.

6. CONCLUDING REMARKS AND OPEN PROBLEMS

We have introduced and studied the algorithmic problem of maximizing revenue in a network using differential pricing, where the prices offered to neighboring vertices cannot be substantially different. It turns out that in the basic model the optimal pricing can be computed efficiently, even for arbitrary revenue functions. In contrast, when discontinuities are allowed the optimization problem becomes computationally hard, though it is still tractable for some special classes of graphs. A similar contrast exists in our stochastic model. We show that the revenue of the best fixed price is similar to the revenue of the optimal pricing, in the basic model, while this is not the case when discontinuities are allowed.

The techniques used in the paper combine combinatorial and graph theoretic tools with properties of totally unimodular matrices and Markov chains. Such techniques may have additional related applications.

Several problems remain open. One of these is the question of deciding whether or not the problem of maximizing the revenue with discontinuities is computationally hard or hard to approximate for single value revenue functions. Another question is that of determining more accurately the correct behavior of the error term in the question discussed in Section 5.1.

Another challenging problem is computing the expected revenue per node for the line graph when the values are random i.i.d. uniform in $[0, 1]$ and discontinuities are allowed.

Besides these technical problems, we believe that the general problem of revenue maximization using differential pricing with local constraints is natural and deserves further study.

REFERENCES

- Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.*, 3(4):801–808, 1990.
- Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., Hoboken, NJ, third edition, 2008. With an appendix on the life and work of Paul Erdős.
- M. Armstrong. Recent developments in the economics of price discrimination. In R. Blundell, W. Newey, and T. Persson, editors, *Advances in Economics and Econometrics: Theory and Applications: Ninth World Congress of the Econometric Society*. Cambridge University Press, Cambridge, UK, 2006.
- G.E. Bolton and A. Ockenfels. A theory of equity, reciprocity and competition. *American Economic Review*, 100:166–193, 2000.
- Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004.
- F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- Gary Chartrand and S. F. Kapoor. The cube of every connected graph is 1-hamiltonian. *J. Res. Nat. Bur. Standards Sect. B*, 73B:47–48, 1969.
- E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM J. Comput.*, 23(4):864–894, 1994.
- E. Fehr and K.M. Schmidt. A theory of fairness, competition and co-operation. *Quarterly Journal of Economics*, 114:817–868, 1999.
- William W. Fisher. When should we permit differential pricing of information? *UCLA Law Review*, 55(1), 2007.
- Alain Ghouila-Houri. Caractérisation des matrices totalement unimodulaires. *C. R. Acad. Sci. Paris*, 254:1192–1194, 1962.

Proceedings Article

- Nima Haghpanah, Nicole Immorlica, Vahab Mirrokni, and Kamesh Munagala. Optimal auctions with positive network externalities. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 11–20, New York, NY, USA, 2011. ACM.
- J. D. Hartline, V. S. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *WWW*, pages 189–198, 2008.
- Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.*, 182(1):105–142, 1999.
- D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.
- J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis and Policy*, 2(1):1102–1127, 2003.
- R. Schmalensee. Output and welfare implications of monopolistic third-degree price discrimination. *American Economic Review*, 71(1):242–247, 1981.
- H. Varian. Price discrimination and social welfare. *American Economic Review*, 75(4):870–875, 1985.
- H. Varian. Price discrimination. In R. Schmalensee and R. Willig, editors, *Handbook of Industrial Organization: Volume I*, pages 597–654. North Holland, Amsterdam, 1989.