

Revenue Maximization with Nonexcludable Goods

MohammadHossein Bateni, Google Research

Nima Haghpanah, Dept. of Electrical Engineering and Computer Science, Northwestern University

Balasubramanian Sivan, Computer Sciences Dept., University of Wisconsin-Madison

Morteza Zadimoghaddam, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology

We consider the design of revenue maximizing auctions for nonexcludable public goods. In particular, we consider auctions in Bayesian settings for facility location problems on graphs. Edges represent sites for facilities (e.g., roads, communication links) to be constructed. Each agent represents a pair of nodes in the graph, and he derives some nonnegative private value drawn from a known distribution if the set of facilities constructed connects his pair. The designer chooses a subset of the edges in the graph to construct facilities on. The nonexcludability arises from the following fact: once a facility is constructed on an edge, every agent is allowed to utilize it for his connectivity purposes.

We show that the pointwise optimization problem involved in implementing the revenue optimal mechanism for this problem is $O(n^{2-\epsilon})$ -hard to approximate even in star graphs, and that even in expectation over the valuation profiles, the problem is APX-hard. In the interesting special case when all the agent pairs share a common endpoint, we construct polynomial time truthful mechanisms that approximate the optimal expected revenue within a constant factor.

We also study the effect of partially mitigating nonexcludability in the form of collecting tolls for using the roads. A posted price mechanism first chooses edges to construct facilities, and posts prices on these edges. Any agent is allowed to use an edge but only if he pays for that edge, thus partially mitigating the effect of nonexcludability. We show that such posted price mechanisms get much higher revenue than auctions subject to the nonexcludability constraint and often approach the optimal revenue obtainable with full excludability.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: nonexcludability, revenue maximization, pricing

ACM Reference Format:

MohammadHossein Bateni, Nima Haghpanah, Balasubramanian Sivan and Morteza Zadimoghaddam, 2013. Revenue Maximization with Nonexcludable Goods *ACM X*, X, Article X (February 2013), 13 pages. DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

How should a seller maximize his revenue while selling nonexcludable services? As a representative example, consider the following facility location problem: firms in an industrial town want to connect their respective warehouses to their outlets with good road links. But the seller (in this case the construction company that lays the roads)

Part of the work was done while the second and fourth coauthors were visiting Google Research, New York. Coauthors' email addresses: bateni@google.com, nima.haghpanah@gmail.com, balu2901@cs.wisc.edu, morteza@mit.edu.

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.

© 2013 ACM 0000-0000/2013/02-ARTX \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

cannot enforce exclusivity on the roads it lays: once a firm buys a connection between its warehouse and outlet, the seller cannot exclude others from using those road links. We study the following two questions in this work. Which potential sites should the seller choose for laying roads to maximize his own revenue in such nonexcludable settings? If the seller could enforce partial excludability by collecting tolls for roads, where only those firms that pay the toll for a road can use it, how much more revenue can he earn?¹

The auction design problem. We assume that each agent/firm has a nonnegative private value, drawn from a known distribution, for getting a connection. We aim for the design of optimal auctions which, given the reports of the agents, selects a set of edges to build roads on and payments of agents in order to maximize the expected revenue. We invoke Myerson's [1981] characterization which states that the expected revenue of any auction in a Bayesian Nash Equilibrium (BNE) is equal to the expected *virtual surplus* of the agents *served*, i.e., the sum of virtual value of those agents whose warehouse and outlet are connected in the solution.² The virtual value of an agent is a function of her value and her distribution of values, and can be negative. This reduces the problem into solving an interesting graph theoretic problem, which to our knowledge has not been studied before. Model the problem with a graph where every unordered pair of vertices denotes an agent and every edge denotes the *possibility* of a road to be laid between two vertices. For each agent we are given a distribution over her possible values, which has nonnegative support. An algorithm selects a subset of edges given the agents' profile of types. We say that an agent is *served* if she can use the selected edges to satisfy her connectivity requirement. Our goal is to design a polynomial time algorithm that, given any graph and distribution over values, selects a subset of edges that (approximately) maximizes the virtual values of served agents *in expectation* (over the randomization of the values and of the algorithm). If the virtual values were all positive, clearly the optimal choice for the seller would be to pick all the edges in the graph. However, the problem is made complicated by the fact that while the value of an agent is always nonnegative, the virtual values could be negative. In particular, while aiming to serve some subset of agents, another subset of agents with negative virtual values automatically get served due to nonexcludability. Deciding on the set of edges to select for maximizing virtual surplus is the nontrivial underlying graph theoretic problem.

The rooted and general versions. We study two versions of the graph theoretic problem described above. In the general version, any unordered pair of vertices in the graph is an agent. In the rooted version (which is a special case of the general version), a single vertex is designated as root, and only those unordered pairs which have root as one of their vertices are agents. This corresponds to the root being a central location where all warehouses are located, and other nodes being outlet locations.

Results for general version (Section 3). As a first attempt, one might hope to design an algorithm that approximates the maximum virtual value *given any profile of values*, which we call the *pointwise* optimization problem. Clearly, a pointwise approximation guarantee also implies optimization in expectation. However, we show that except for very special cases, pointwise optimization is not possible unless $P = NP$. In particular, we show that the problem is NP-hard to approximate to within a factor $O(n^{2-\epsilon})$ even on star graphs. For the special case where the underlying graph is a path, we

¹Note that collecting tolls only enforces partial excludability because *every* firm that pays the toll is allowed to use the road.

²Note that Myerson's mechanism remains revenue optimal in all single-parameter settings regardless of whether or not the good is excludable.

give a dynamic program to solve the pointwise optimization. Thus, for paths, we get the expected optimal revenue with nonexcludability for any product distribution over agents' values.

Results for rooted version (Section 4). Our results for the rooted version are three-fold.

- (1) First, we give a simple polynomial time algorithm for pointwise optimization in trees—this contrasts with our result that for the general version pointwise optimization is hard to approximate beyond a factor of $O(n^{2-\epsilon})$ even for star graphs.
- (2) Second, as our main result, we give a polynomial time algorithm for optimization in expectation for general graphs, when the agents' valuations are drawn i.i.d.
- (3) Third, we establish APX-hardness of optimization in expectation for general graphs, when the valuations of all agents are independent but not necessarily identical.

Our main result, which is the second point above, is that we design an algorithm that guarantees a constant factor approximation to the optimal virtual surplus in expectation. To this end, we extensively use the key (yet simple) property that even though virtual values can be negative, the expected virtual value of an agent is nonnegative. This suggests the following high-level approach to the problem. Partition agents into a constant number of sets. Pick a target set at random, and run an algorithm that is a constant factor approximation for the agents of that set (ignoring other agents). The nonnegativity of the expected virtual value implies that the contribution from non-targeted sets is nonnegative. Since each set is targeted with constant probability, this implies a constant factor approximation. We use this approach to solve the abstract edge-weighted version of the problem (in which edges are agents who derive value from being connected to the root), and then reducing the original problem to the edge-weighted version. In particular, we present an algorithm that partitions the edges of any graph into two sets that, loosely speaking, correspond to edges in well-connected parts of the graph and edges in sparse cuts of the graph. We show that once we contract the edges in a set (that corresponds to ignoring their value in the above high level approach), the remaining graph has a nice structure that allows for constant approximations.

Partial excludability via pricing (Section 5). If the seller were allowed to set prices on edges (i.e., collect tolls), can he get close to the optimal revenue when excludability is allowed? Typically the optimal revenue with full excludability is much higher than the optimal revenue with nonexcludability, and it is worthwhile for a profit maximizing firm to explore this option. In this problem the seller first has to decide which edges to pick to lay roads on, and decide how to price roads to maximize revenue. We construct a pricing scheme so that when the value distributions are i.i.d., the seller obtains (for the rooted version) the optimal revenue when full excludability is allowed. Further, if the distributions are independent but not identical and satisfy a technical MHR condition³, we show how to price edges to obtain a $\frac{1}{\log n}$ fraction of the optimal revenue possible when full excludability is allowed (again for the rooted version). An implication of these results is that mitigating externalities via tolls is much more remunerative than aiming for the optimal solution with nonexcludability.

³Roughly speaking, this means that the tail of the distribution is no heavier than the exponential distribution. Many natural classes of distributions like the uniform, exponential, Gaussian (Normal) distributions satisfy the MHR condition. See Section 2 for a formal definition.

Related Work. Approximating the expected version of a problem that is hard-to-approximate in the worst case, via exploiting the fact that the expected virtual value for any distribution is nonnegative is a relatively new idea. To our knowledge, it has been used only in Hagpanah et al. [2011]. Auctions with externality (a notion related to nonexcludability where an agent’s utility doesn’t just depend on the services he received but also on the outcomes for the other agents) have been studied in multiple flavors before. Settings with positive externality include increased value for having a telephone or a music player or a new technology if more of your neighbors have them [Rohlf’s 1974; Hartline et al. 2008]. A prime and well-studied example for a setting with negative externality is the sale of contiguous ad slots to two competing businesses [Aggarwal et al. 2008; Athey and Ellison 2011; Giotis and Karlin 2008; Gomes et al. 2009; Jeziorski et al. 2009; Kempe and Mahdian 2008]. Equilibria which are surprising at first sight, like no allocation equilibria can result in large revenue for the auctioneer in settings with negative externality [Deng and Pekec 2011]. Posted pricings have often been a mechanism of choice for settings with externalities [Anari et al. 2010; Akhlaghpour et al. 2010; Candogan et al. 2010; Hartline et al. 2008]. The main difference between these and the posted prices studied in our work (apart from the presence of nonexcludability in our settings) is that these works allow agent-specific prices, whereas our setting is more constrained: we place prices on edges that are common for all agents.

2. MODEL AND NOTATION

General version. We consider a universe of n potential sites, located on the vertices of a graph, $G = (V, E)$, with m undirected edges. An undirected edge $(i, j) \in E$ means that a link/road connecting sites i and j is allowed (but need not necessarily be constructed). An agent is an unordered pair of sites (i, j) , interested in having some path constructed between i and j . Thus, there could be up to $\binom{n}{2}$ agents in a mechanism. An instance $\mathcal{I} = (G, A, F)$ of the problem consists of an undirected graph G , a set A of agents (represented by a set of pairs of vertices), and a distribution F_i associated with agent i (distributions are explained below). Sometimes we use i to denote a single site and sometimes to denote an agent, who is actually a pair of vertices. The context will make it clear whether i is a single site or a pair of sites.

Rooted version. The rooted version is a special case of the setting described above. A special vertex r is designated as the root. Only vertex pairs of the form (i, r) are agents, i.e., the root r is one of the end points of the path desired for every agent.

An outcome $o \in \Omega = \{0, 1\}^m$ is the set of edges constructed. Agent i has a valuation function $v_i : \Omega \rightarrow \mathbf{R}^+ \cup \{0\}$, which maps outcomes to nonnegative real numbers. We study mechanisms in the single-parameter setting where the function $v_i(\cdot)$ takes only two values: v_i and zero. An agent i has a nonzero value v_i if and only if the set of edges selected contains a path between the two sites she represents, and in this case we say that agent i was served. Let S denote the set of all feasible sets of agents, i.e., the set of all sets of agents that can be simultaneously served. Note that this set system S , is not downward closed: $S \in \mathcal{S}$ does not necessarily mean that $S' \in \mathcal{S}$ for all $S' \subset S$. Clearly, the non-downward closedness stems from the inability to exclude agents from using the roads.

We study mechanisms for this problem in a Bayesian setting, i.e., for every i , the single parameter v_i is assumed to be drawn independently from a publicly known distribution function F_i . Thus $F = F_1 \times F_2 \times \dots \times F_n$ denotes the product distribution from which the vector of types \mathbf{v} is drawn. The mechanisms in this paper assume the availability of any expectation defined with respect to F .

A direct revelation mechanism or an auction solicits sealed bids (b_1, b_2, \dots, b_n) from all the agents, and determines the outcome $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and payments $\mathbf{p} = (p_1, p_2, \dots, p_n)$. Each agent i is risk-neutral and has a linear utility $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$.

Let $(\mathbf{x}(\cdot), \mathbf{p}(\cdot))$ denote a mechanism. When agent i is bidding in the auction, she knows only her own value v_i . A mechanism $(\mathbf{x}(\cdot), \mathbf{p}(\cdot))$ is incentive compatible (IC) if $v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq v_i x_i(z, v_{-i}) - p_i(z, v_{-i})$ for all i, z . Since we focus on the class of IC mechanisms, we have $\mathbf{b} = \mathbf{v}$.

Optimal auctions. To solve for optimal auctions, Myerson [1981] defined *virtual valuations* for agents as $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ and proved that the expected payment of an agent, $\mathbf{E}_{v_i}[p_i(v_i)]$ is equal to her expected virtual value $\mathbf{E}_{v_i}[\phi_i(v_i)x_i(v_i)]$ ⁴. The distribution F_i is said to be *regular* if the virtual valuation function is monotone. For regular distributions, maximizing virtual values pointwise results in a Bayesian incentive compatible (BIC) allocation rule and therefore the corresponding revenue-optimal auction serves that feasible set of agents who maximize virtual value. For irregular distributions, where maximizing virtual values may result in non-BIC allocation rules, Myerson applies a fix by describing a general *ironing* technique. The ironing procedure converts any virtual valuation function $\phi_i(\cdot)$ to an *ironed virtual value function* $\bar{\phi}_i(\cdot)$ such that maximizing $\bar{\phi}_i(\cdot)$ pointwise results in a BIC allocation rule.

THEOREM 2.1. [Myerson 1981] *The revenue optimal auction in single-parameter Bayesian settings serves the set S of agents where $S = \operatorname{argmax}_{S \in \mathcal{S}} \sum_{i \in S} \bar{\phi}_i(v_i)$.*

As a corollary, when S contains all possible 2^n sets, the revenue optimal auction puts a price of $\phi_i^{-1}(0)$ for agent i and makes a take-it-or-leave-it offer to each of them.

Monotone Hazard Rate. A class of distributions that always has the regularity condition described above are the ones which satisfy the monotone hazard rate condition. A distribution with cdf F satisfies the MHR condition if $\frac{f(x)}{1-F(x)}$ is nondecreasing in x . The MHR condition holds for many natural classes of distributions like the uniform, exponential and normal distributions.

nonnegative virtual valuations. An important property of virtual valuations (and ironed too) is that when the support of the distribution is nonnegative (which is true in our case since values are nonnegative), the virtual valuation in expectation over an agent's distribution is always nonnegative. This property crucially helps us in providing approximations in expectation, for problems which are very hard to approximate for every single realization of values.

3. THE GENERAL VERSION

3.1. Pointwise optimization in paths

In this section, we show that in paths, a polynomial-time dynamic program is all that is necessary to implement the pointwise optimization involved in implementing Myerson's mechanism.

Consider a path with vertices 1 to n . For $1 \leq i \leq j \leq n$, let $S(i, j)$ be the sum of virtual values of the set of all agents (i.e., vertex pairs) whose both endpoints are in the interval $[i, j]$. We set $S(i, i) = 0$ for all i —assuming that there is no trivial agent whose vertex pairs are the same. Define $OPT(i)$ to be optimal solution when we are restricted to choose only among edges connecting vertices 1 through i . The unrestricted

⁴In fact the equality holds even for a specific \mathbf{v}_{-i} , i.e., $\mathbf{E}_{v_i}[p_i(\mathbf{v})] = \mathbf{E}_{v_i}[\phi_i(v_i)x_i(\mathbf{v})]$.

optimal solution OPT is therefore $OPT(n)$. We set $OPT(0) = 0$. The following recursive formula can be used in the dynamic program to solve for OPT .

$$\forall i \leq n, OPT(i) = \max_{0 \leq k < i} OPT(k) + S(k+1, i).$$

In the above recursion, $k+1$ is the leftmost vertex that is connected to i . All vertices $k+1, \dots, i$ are connected, and $S(k+1, i)$ is by definition their contribution to the objective. Since the edge connecting k to $k+1$ is not included, the set of edges chosen among the first k vertices must be equal to $OPT(k)$.

3.2. $O(n^{2-\epsilon})$ hardness of pointwise optimization in stars

In this section, we show that if we move even slightly beyond paths, to star graphs, pointwise optimization becomes NP-hard. We show that given a star graph, $(0, 1, 2, \dots, n)$ with edges $(i, 0)$ for all i , and agents being (i, j) for $i, j \neq 0$, there exists an assignment of virtual values to agents for which the virtual value maximizing subset cannot be found. To demonstrate this, it is enough to show that there is some set of $\binom{n}{2}$ virtual values (positive or negative), one for each agent, for which finding the optimal set of agents is inapproximable.

THEOREM 3.1. *Given an arbitrary set of virtual values, the problem of finding the virtual value maximizing set of agents in a star graph is $O(n^{2-\epsilon})$ hard to approximate.*

PROOF. The reduction is from the $O(n^{1-\epsilon})$ hardness of approximation of independent set [Hastad 1996]. Given an arbitrary graph $G = (V, E)$, assign virtual values on the star graph G' with n spokes and a center 0 as follows. For every edge $(i, j) \in E$, let the virtual value of agent (i, j) in the star graph be $-n^2$. For every pair $(i, j) \notin E$, let the virtual value of agent (i, j) in the star graph be 1. It is easy to see that the optimal solution in the star graph will pick an independent set of vertices from G and connect them to the center. Thus if the independent set is of size k , the value of the optimal solution is $\binom{k}{2}$. We know that independent set is inapproximable to a factor of $O(n^{1-\epsilon})$ in general graphs. Thus, the value of the optimal set in star graphs is inapproximable to a factor of $O(n^{2-\epsilon})$. \square

3.3. APX-hardness of optimization in expectation in general graphs

Given Theorem 3.1, the only possible approximation we can hope for is in expectation over realized values, and not for every realized value. We show in this section the APX-hardness of this problem in general graphs. The reduction resembles the one given by Haghpanah et al. [2011]. We note that our hardness result holds even for the special case of rooted version, and hence the proof is presented in Section 4.3 along with other results for the rooted version.

4. THE ROOTED VERSION

Given the hardness results in Section 3 even for undirected graphs, we consider an important special case of our problem in this section: there is a designated root node in an undirected graph, and each nonroot vertex is an agent. An agent derives value only if there is a path constructed from his node to the root node. As before, construction of an edge e can happen only if $e \in G$.

4.1. Pointwise optimization in trees

As a subroutine for optimization in trees, we first show how to perform pointwise optimization in paths (much easier and faster than the pointwise optimization for paths in the general version)

4.1.1. Pointwise optimization in paths. Consider an undirected path (it is easy to see that the directed path problem can be reduced to this undirected case) of length $n + 1$ with one end of the path, namely node 0, as the root. If the root is not one of the end vertices, such a path can be broken into two at the root, and each path will have root at one of its ends. (These problems can be solved independently.) Thus root being at the end of the path is without loss of generality. Once the agents submit their values, the auction has to decide on the set of edges to pick. In the path case, this task reduces to picking the best among paths of the form $0 - i$ for $i = 1 \dots n$. For every i , the mechanism computes the sum of the virtual values of agents from 1 to i and picks the i with the maximum virtual value.

4.1.2. Pointwise optimization in trees. Consider an undirected tree with a designated root. Since there is a unique path between every node and the root, it is easy to see that the directed tree can be reduced to the undirected tree case by simply eliminating few nodes from the tree which can never reach the root. The revenue optimal mechanism in trees is constructed using the optimal mechanism in paths as a subroutine. The mechanism proceeds from the bottom-up. It chooses the bottom-most node, say b , in the tree which has branches below. By definition, each such branch must be a path. The mechanism computes the optimal solution and the corresponding optimal virtual value for each path, and adds them up and assigns them to b . All the descendants of b are now deleted, with the virtual value of b now adjusted to the new value. This procedure is now repeated in the new tree till we exhaust all nodes. By induction on the depth of the tree, it is easy to establish the revenue optimality of this mechanism.

4.2. Main result: Optimization in expectation in general graphs

Next we present a constant-factor approximation algorithm for optimizing the expected revenue for the rooted, node-weighted version of the problem in general graphs (the virtual value of an agent is the weight of the nonroot node of that agent). In Section 4.2.1 we explain how the problem can be solved on edge-weighted graphs. This is later used in Section 4.2.2 as a subroutine to tackle the vertex-weighted problem.

4.2.1. Edge-Connectivity for the Rooted Version. Given a connected undirected graph, we show how to partition its edges into two parts such that contracting the former edge set yields a 3-edge connected subgraph whereas contracting the latter edge set results in a *roulette* subgraph—a special series-parallel graph to be defined below. We then demonstrate that it is possible to solve the problem (approximately) on each of the two subgraphs, and finally argue that this suffices to obtain a constant-factor approximation for the general rooted edge-weighted case.

Let us define some notations first. For a set S of edges in a graph G , subgraph G/S is obtained from G after contracting all edges S one at a time, where contracting an edge simply refers to removing the edge and identifying its endpoint vertices. We recursively define the class of roulette graphs as follows. A simple cycle (possibly of length one) is a roulette, and so is a cycle each of whose vertices is replaced by a roulette (with one or two vertices of the inner roulette taking the place of an original vertex of the cycle). Finally, any graph whose 2-edge connected components are roulettes is itself a roulette.

Now we can present the main structural lemma that reduces our general problem into two tractable subproblems.

LEMMA 4.1. *There exists a polynomial-time algorithm that, given a graph $G(V, E)$, partitions the edge set E into two sets S_1 and S_2 such that graphs $G_1 = G/S_1$ and $G_2 = G/S_2$ are respectively 3-edge connected and roulette.*

PROOF. Let S_1 be the set of all edges in G that belong to a cut of size at most 2. We note that all cuts of size at most 2 can be found in polynomial time, e.g., using a naïve

brute-force search. Since G_1 is obtained by a series of edge contractions, every cut in G_1 represents a cut in G as well. Therefore, since all edges of S_1 are contracted in G_1 , no cut of size at most 2 is present in G_1 , hence G_1 is 3-edge connected.

On the other hand, every edge in G_2 belongs to a cut of size at most 2 in G , hence in G_2 . The bridges in G_2 do not hurt the roulette structure if the 2-edge connected subgraphs of G_2 are roulettes. Thus we assume that the graph G_2 is 2-edge connected. We claim that if each of $\{e_1, e_2\}$ and $\{e_1, e_3\}$ is a cut in G_2 , then so is $\{e_2, e_3\}$. Therefore, the edges of G_2 form an equivalence class. To see this equivalence relation, it suffices to focus on the following alternative definition of edge cuts of size 2. In a 2-edge connected graph, two edges e and e' form a cut of size 2 if and only if the set of cycles in the graph that contain e is the same as the set of cycles in the graph that contain e' . Based on this new definition, the two sets of cycles containing e_2 and e_3 respectively are both equal to the set of cycles containing e_1 , and therefore they are equal to each other as well which means that e_2 and e_3 form a cut of size 2.

Therefore, the graph looks like a cycle of this equivalence class (the equivalence class of e_1) where some vertices are replaced by another structure; the same argument applies to each of these smaller structures, giving rise to the inductive definition of roulette graphs. We should note that two edges from two of these smaller structures do not belong to the same equivalence class (they cannot form a cut of size 2), and therefore each of the other equivalence classes belong to one smaller structure and is not split between different structures. Thus we can inductively claim each smaller structure is a roulette graph. \square

The following decomposition lemma serves as the starting point for our 3-approximation algorithm of the 3-edge connected graph G_1 .

LEMMA 4.2. *There exists a polynomial-time algorithm that finds 3 spanning trees T_1, T_2 , and T_3 in a 3-edge connected graph G such that every edge of G is missing in at least one of these spanning trees.*

PROOF. We replace each edge of G by two parallel edges to get the 6-edge connected graph G^2 with the same vertex set. Catlin et al. [2009] show, among other things, that any $2k$ -edge connected graph has k edge-disjoint spanning trees, while Roskind and Tarjan [1985] show how to find k edge-disjoint spanning trees in a graph (if they exist) in quadratic time. Therefore, we can find 3 edge-disjoint spanning trees T_1, T_2 , and T_3 in G^2 . Edge-disjointness guarantees that each edge of G can belong to at most two of these spanning trees. \square

On the other hand, the problem can be solved optimally for roulette graphs. The intuition behind the algorithm is that roulettes can be shown to have treewidth of at most two, hence, as are many similar problems on bounded-treewidth graphs, our problem can be solved the dynamic-programming method.

LEMMA 4.3. *There exists a polynomial-time algorithm that finds the optimal solution for roulette graphs.*

PROOF. Let us say we have a *cycle decomposition* of our roulette graph, which describes the recursive structure of its 2-edge connected components. Each cycle representing one equivalence class is called an *essential cycle* of the graph. Instead of solving the rooted problem, we consider a slightly more general problem where up to two vertices s, t on an essential cycle are specified, and these vertices should both appear in the connected subgraph of the output. At the beginning we are going to have only one vertex $s = t$ which is the root vertex.

Let us ignore the bridges at this point and assume the graph is 2-edge connected. Focus on the essential cycle, and imagine that all the recursive structures are con-

tracted for now. In the resulting graph, the optimal solution looks like a path or it is the entire cycle. There are polynomially many cases to consider, and we will output the best solution among them. In each case we have a subproblem for each contracted piece if we decide that the optimal solution passes through or ends at the contracted vertex. Since the base cases of the induction (i.e., vertices or cycles) are easily solvable, we can argue inductively that our slightly more general problem can be solved using a dynamic program.

For each bridge connected to the essential cycle, we compute the best solution for the rooted problem on the other side of the bridge (whose root is the endpoint of the bridge) and add that to the main solution if its weight plus the weight of the bridge turns out positive. \square

We conclude this section by putting together the above ideas to obtain a 4-approximation algorithm for the rooted edge-weighted problem.

LEMMA 4.4. *There exists a polynomial-time algorithm that achieves an approximation factor of 4 for any graph G .*

PROOF. If the graph is not connected, we can focus on the connected component that contains the root and disregard the rest of the connected components. Using Lemma 4.1, we partition the edges of G into two parts S_1 and S_2 . We also use Lemma 4.2 to find three spanning trees T_1, T_2 , and T_3 in graph G_1 .

We consider four candidate solutions, and in each realization of edge weights we pick the candidate solution with the maximum revenue. One solution is the union of S_1 and the edges in tree T_1 . Since T_1 is a spanning tree in $G_1 = G/S_1$, the union of T_1 and S_1 is a connected spanning subgraph of G . Serving this connected spanning subgraph allows us to choose any subset of the remaining edges to serve. Among the remaining edges (edges not in $S_1 \cup T_1$), we serve those with positive realized virtual values. In a similar fashion we construct two other candidate solutions based on T_2 and T_3 . The fourth and last candidate solution is to contract set S_2 of edges, and solve the problem optimally in the roulette graph G_2 using Lemma 4.3. The optimal solution we find in G_2 is a connected subgraph of G_2 , however, it might not be a connected subgraph of G that includes the root. However, it is always possible to serve a subset of edges $S'_2 \subseteq S_2$ to make the whole solution not only a connected subgraph of G but also one that includes the root vertex as well. Our fourth candidate solution consists of the edges in S'_2 and the optimal solution for G_2 .

Rather than showing that picking the best solution in each realization leads to a 4-approximate solution, we prove a stronger claim—for any instance there is one of these candidate solutions that guarantees a 4-approximate solution on its own. For $1 \leq i \leq 3$, if we stick to solution i all the time, our gain from edges in S_1 and T_i is nonnegative (i.e., the sum of their expected virtual values), and in addition we get all edges with positive virtual value outside $S_1 \cup T_i$. Since each edge of S_2 is missing in at least one $S_1 \cup T_i$, the total value we get from the first three solutions is at least the projection of optimal solution on set S_2 of edges. On the other hand in the fourth solution, our expected revenue from edges added from S_2 is nonnegative (once again since the expectation of the virtual values are positive), and our expected revenue from S_1 is at least the amount that the optimal solution gains from them. Therefore, these four solutions together achieve no less than the optimal revenue. Thus, one candidate solution has expected revenue at least a quarter of the optimum. \square

4.2.2. Vertex-Connectivity for the Rooted Version. In this section we provide a constant-factor approximation for the vertex-connectivity problem using the algorithm for the edge-connectivity version of the problem described above. Let V_2 be the set of degree-2 vertices in graph G . Similar to the edge-connectivity approach, we describe two algo-

rithms (one using a reduction to the edge-connectivity problem) that achieve constant-factor approximations to the problems where the values of $V \setminus V_2$ and V_2 are replaced by zero, respectively. Again, since the expected value of each vertex is nonnegative, this implies a constant approximation for the vertex-connectivity problem.

First consider the instance in which the values of vertices in $V \setminus V_2$ are replaced by zero. Notice that this results in an instance in which any vertex with nonzero value has degree 2. Construct another instance in which each vertex of degree 2 is replaced by an edge with the same value. We can solve this instance using our edge-connectivity algorithm.

Next consider the instance in which the values of vertices in V_2 are replaced by zero. We convert the instance to one without any degree-2 vertices via replacing all paths consisting only of degree-2 vertices by an edge. The following lemma shows that this graph has a spanning tree at least $\frac{1}{7}$ of its vertices are leaves. The algorithm uses the *internal* nodes (i.e., nonleaves) of this tree to connect the leaves, whenever they are positive, to the root. This gives a 7-approximation to the problem.

Putting these two algorithms together, we obtain a constant-factor approximation algorithm for the vertex-connectivity rooted revenue maximization in expectation. In particular, a balancing argument puts a bound of 11 on its approximation ratio.

LEMMA 4.5. *Given a graph with no degree-2 vertices, a spanning tree can be constructed in polynomial-time where at least $\frac{1}{7}$ of the vertices are leaves.*

PROOF. Start from an arbitrary spanning tree T of G . Let T_2 be the set of vertices of degree 2 in T , and let $\hat{T}_2 \subseteq T_2$ be those vertices in T_2 both whose neighbors, too, are in T_2 . Modify T as long as any of the following two rules apply.

- (1) If there exists a vertex $v \in \hat{T}_2$ that has an edge in G to an internal vertex u of T , update T by adding the edge (v, u) to T , and removing the edge incident to v in the unique cycle formed after adding (v, u) . Since both neighbors of v in T had degree 2, this process generates a new leaf without removing any of the old leaves.
- (2) If two vertices $v, u \in \hat{T}_2$ have edges in G to the same leaf l of T , add edges from v and u to that leaf, and remove two edges from T as follows. The addition of edge (u, l) to T produces a cycle that passes through exactly one of the two neighbors of u ; call it u' . Note that u' has degree 2 in T by definition of \hat{T}_2 ; let u, u'' be its neighbors. Remove the edge (u', u'') from T , removing the cycle u and maintaining the connectivity of T . We carry out a similar operation, *mutatis mutandis*, for v . The result will be a tree T on the same set of vertices with one more leaf (increasing the degree of l but turning two other internal vertices into leaves).

The process terminates in a linear number of iterations since the number of leaves increases in each step. We end up with a tree T for which neither of the rules applies. Let T_1 and $T_{\geq 3}$, respectively, denote subsets of vertices of degrees one and at least three in T . We argue below that $|T_1|$ is at least a constant fraction of $|T_2| + |T_{\geq 3}|$. As no vertex of G has degree two, any vertex in \hat{T}_2 is bound to have degree at least three in G , hence an edge not in T . This edge cannot be to an internal vertex of T because Rule (1) no longer applies to T . Rule (2), on the other hand, implies that these leaves are distinct for different vertices of \hat{T}_2 . Therefore, we have

$$|T_1| \geq |\hat{T}_2|. \quad (1)$$

As trees have average degree less than two, we know

$$|T_1| > |T_{\geq 3}|. \quad (2)$$

To bound $|T_2| - |\hat{T}_2|$, if this quantity is not zero, orient T from an arbitrary vertex in $T_2 \setminus \hat{T}_2$ towards the leaves. Assign each vertex $v \in T_2 \setminus \hat{T}_2$ to its closest descendant in the oriented tree that is in $T_1 \cup T_{\geq 3}$. Such an assignment is always possible since no vertex in the former group is a leaf of the (oriented) tree. Each vertex in the latter group is assigned to at most twice, otherwise there should be a path of vertices of degree two with more than two vertices in $T_2 \setminus \hat{T}_2$ —a contradiction. As a result, we get

$$|T_2| - |\hat{T}_2| \leq 2(|T_1| + |T_{\geq 3}|) \leq 4|T_1|, \quad (3)$$

where the last inequality is due to (2).

Summing up (1), (2) and (3) with $|T_1| \geq |T_1|$, we obtain $7|T_1| \geq |T_1| + |T_2| + |T_{\geq 3}|$ as desired. \square

4.3. APX-hardness of optimization in expectation in general graphs

In this section we prove the APX-hardness of the rooted version in general graphs when the valuations of agents are independent but not necessarily identical.

Definition 4.6. The prize-collecting set cover problem (PCSCP) consists of a collection of sets S_1, S_2, \dots, S_n over a universe U . For a collection C of sets, let $Q_C = \cup_{i \in C} S_i$. The goal is to find a collection C^* that maximizes $\alpha|Q_{C^*}| + n - |C^*|$ for some given $\alpha > 0$.

We first show that there is an approximation preserving reduction from PCSCP to our problem, and then use the result from [Haghpanah et al. 2011] that establishes the APX-hardness of PCSCP, and that its approximation ratio is at least $\frac{530}{529} = 1.002$.

LEMMA 4.7. *There is an approximation preserving reduction from PCSCP to our problem*

PROOF. Given an instance of the PCSCP, where the sets are denoted S_1, S_2, \dots, S_n and the elements are denoted e_1, e_2, \dots, e_m , we construct an instance of our problem as follows.

Vertices. We start by putting a root vertex r . For every element e_i , we construct one vertex with the same name e_i . For every set S_i , we construct one vertex denoted by S_i as well.

Edges. For each i , set vertex S_i is connected by an edge to root vertex r and all e_j for which $e_j \in S_i$.

Agents. The agents are (r, S_i) for each set i , called set-agents, and (r, e_i) for each element i , called element-agents.

Distribution. The value distribution of element agents is deterministic α . For set agents, the value is drawn from the distribution Bernoulli($L - 1, 1/L$)—i.e., the value is equal to $L - 1$ with probability $1/L$, and zero otherwise—where $L \gg mn\alpha$. Thus, the virtual value for these agents is -1 w.p. $1 - 1/L$ and $L - 1$ w.p. $1/L$.

The optimal revenue in our problem, as $L \rightarrow \infty$, can be analyzed in two cases.

- (1) If at least one set-agent has positive virtual valuation (which happens w.p. approximately $n/L \rightarrow 0$), the solution chooses all the edges incident on those set agents (with positive virtual valuation) to obtain expected revenue n . The expected revenue from the remaining agents (set-agents with negative virtual valuation, and element-agents) is at most $\alpha mn/L \ll 1$. Therefore, the optimal solution has contribution n from this event as $L \rightarrow \infty$, and this solution is trivial to compute.

- (2) If no set has positive virtual valuation (which happens w.p. $1 - n/L \rightarrow 1$), the value of the solution is precisely $\alpha|Q_{C^*}| - |C^*|$. This is because once a set-agent is chosen, clearly all the edge-agents that is contained by this set must be chosen since they all have deterministic positive virtual value α . We should also note that you can obtain the α virtual value of an element agents only if you connect it to root using one of the set agents it belongs to, and for each set you pick you have -1 virtual value in this case.

Therefore the value of the optimal solution is $\alpha|Q_{C^*}| + n - C^*$. \square

5. POSTED-PRICING RESULTS

As mentioned in the introduction, one goal of this work to compare to natural mechanisms for nonexcludable public goods—a direct revelation mechanism, which has to deal with externalities, and a posted price mechanism which mitigates externality to a certain extent by ensuring that only those who pay for an edge can use that edge. In this section we show that for the rooted version, a posted price mechanism approximates the optimal mechanism’s revenue even without any externality constraints—i.e., suppose we were allowed to run Myerson’s mechanism without being subjected to the externality that arises through the nonexcludability, clearly our revenue would be higher. We show that a posted price mechanism approximates even this superior benchmark. This partly explains why having tolls on freeways is far more remunerative than raising funds once and for all at the beginning.

5.1. Rooted version with i.i.d. agents

Consider a directed graph G with a designated root r . Without any externality constraints, Myerson’s mechanism’s revenue is $\sum_{i \neq r} \phi^{-1}(0)(1 - F(\phi^{-1}(0)))$, where $1 - F(\phi^{-1}(0))$ is the probability that agent i ’s value is above $\phi^{-1}(0)$. The posted price mechanism places a price of $\phi^{-1}(0)$ on every edge incident on the root (directed either way). Every agent who can reach the root must use one such edge, and thus pays the price if his value is higher than his price. This gives a revenue of precisely $\sum_{i \neq r} \phi^{-1}(0)(1 - F(\phi^{-1}(0)))$, which is the optimal revenue possible.

On the other hand, if we use a direct revelation mechanism, the optimal revenue achievable, even for paths, is a factor of \sqrt{n} away from the externality-free Myerson’s mechanism’s revenue. Consider the simple example of a path in which the value of each vertex is 1 with probability $1/2$, and is -1 otherwise. If we are allowed to select vertices regardless of the externality, we will select each vertex when its value is positive, achieving the expected value of $n/2$. On the other hand, the optimal value considering externalities is obtained by selecting the sub-path with maximum value that is connected to the root. This can be seen as a random walk with n steps, in which we are interested in the expected maximum distance during the process. This expectation is well known to be $O(\sqrt{n})$.

5.2. Rooted version with non-i.i.d. agents

When the distributions are independent but not necessarily identical, we design a pricing which obtains a $O(\log n)$ approximation when the distributions follow the MHR property (i.e., the hazard function $\frac{f(x)}{1-F(x)}$ is monotonically nondecreasing). When the agents are non-iid, the monopoly price of all agents are not necessarily the same—so the previous mechanism clearly fails. We need two facts to give a $O(\log n)$ approximation.

- (1) When a distribution F follows MHR property, we have $F(\phi^{-1}(0)) \leq 1 - 1/e$. This follows from

$$F(x) = 1 - e^{-\int_0^x h(t) dt} \leq 1 - e^{-\int_0^x h(x) dt} = 1 - e^{-xh(x)},$$

where the inequality follows from the monotonicity of $h(x)$. Since $xh(x) = 1$ at $x = \phi^{-1}(0)$, we have that $F(\phi^{-1}(0)) \leq 1 - 1/e$.

- (2) Given n numbers v_1, v_2, \dots, v_n , we have $\max_i iv_i \geq \frac{1}{\log n} \sum_i v_i$. Let $t = \operatorname{argmax}_i iv_i$. So for all i , $tv_t \geq iv_i$, and thus $v_i \leq tv_t/i$ for all i . Summing over this inequality for all i , we get $\sum_i v_i \leq tv_t(\log n)$.

Without loss of generality let $\phi_1^{-1}(0) \geq \phi_2^{-1}(0) \geq \dots \geq \phi_n^{-1}(0)$. Let $t = \operatorname{argmax}_i i\phi_i^{-1}(0)$. Now suppose we post a price of $t\phi_t^{-1}(0)$. All agents with monopoly price above $\phi_t^{-1}(0)$ will pay the toll with probability at least $1/e$ (by point 1 above). Thus we get a revenue of $\frac{1}{e} \max_i i\phi_i^{-1}(0) \geq \frac{1}{e \log n} \sum_i i\phi_i^{-1}(0)$ (by point 2 above). Thus we get a $O(\log n)$ approximation to the optimal revenue with full excludability.

REFERENCES

- AGGARWAL, G., FELDMAN, J., MUTHUKRISHNAN, S., AND PÁL, M. 2008. Sponsored search auctions with markovian users. *Internet and Network Economics*, 621–628.
- AKHLAGHPOUR, H., GHODSI, M., HAGHPANAH, N., MIRROKNI, V., MAHINI, H., AND NIKZAD, A. 2010. Optimal iterative pricing over social networks. *Internet and Network Economics*, 415–423.
- ANARI, N., EHSANI, S., GHODSI, M., HAGHPANAH, N., IMMORLICA, N., MAHINI, H., AND MIRROKNI, V. 2010. Equilibrium pricing with positive externalities. *Internet and Network Economics*, 424–431.
- ATHEY, S. AND ELLISON, G. 2011. Position auctions with consumer search. *The Quarterly Journal of Economics* 126, 3, 1213–1270.
- CANDOGAN, O., BIMPIKIS, K., AND OZDAGLAR, A. 2010. Optimal pricing in the presence of local network effects. *Internet and Network Economics*, 118–132.
- CATLIN, P. A., LAI, H.-J., AND SHAO, Y. 2009. Edge-connectivity and edge-disjoint spanning trees. *Discrete Mathematics* 309, 5, 1033–1040.
- DENG, C. AND PEKEC, S. 2011. Money for nothing: exploiting negative externalities. In *Proceedings of the ACM Conference on Electronic commerce (EC)*. 361–370.
- GIOTIS, I. AND KARLIN, A. 2008. On the equilibria and efficiency of the gsp mechanism in keyword auctions with externalities. *Internet and Network Economics*, 629–638.
- GOMES, R., IMMORLICA, N., AND MARKAKIS, E. 2009. Externalities in keyword auctions: An empirical and theoretical assessment. *Internet and Network Economics*, 172–183.
- HAGHPANAH, N., IMMORLICA, N., MIRROKNI, V., AND MUNAGALA, K. 2011. Optimal auctions with positive network externalities. In *Proceedings of the 12th ACM conference on Electronic commerce*. EC '11. 11–20.
- HARTLINE, J., MIRROKNI, V., AND SUNDARARAJAN, M. 2008. Optimal marketing strategies over social networks. In *Proceedings of the 17th international conference on World Wide Web*. ACM, 189–198.
- HASTAD, J. 1996. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*.
- JEZIORSKI, P., SEGAL, I., ET AL. 2009. What makes them click: Empirical analysis of consumer demand for search advertising. *SSRN eLibrary* 33.
- KEMPE, D. AND MAHDIAN, M. 2008. A cascade model for externalities in sponsored search. *Internet and Network Economics*, 585–596.

- MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1, 58–73.
- ROHLFS, J. 1974. A theory of interdependent demand for a communications service. *The Bell Journal of Economics and Management Science*, 16–37.
- ROSKIND, J. AND TARJAN, R. E. 1985. A note on finding minimum-cost edge-disjoint spanning trees. *Mathematics of Operations Research* 10, 4, 701–708.