

Revenue Maximization in Stackelberg Pricing Games: Beyond the Combinatorial Setting

Toni Böhnlein¹, Stefan Kratsch², and Oliver Schaudt³

¹ Universität zu Köln, Institut für Informatik, Cologne, Germany

boehnlein@zpr.uni-koeln.de

² Universität Bonn, Institut für Informatik, Bonn, Germany

kratsch@cs.uni-bonn.de

³ Universität zu Köln, Institut für Informatik, Cologne, Germany

schaudto@uni-koeln.de

Abstract

In a Stackelberg Pricing Game a distinguished player, the *leader*, chooses prices for a set of items, and the other players, the *followers*, each seeks to buy a minimum cost feasible subset of the items. The goal of the leader is to maximize her revenue, which is determined by the sold items and their prices. Most previously studied cases of such games can be captured by a combinatorial model where we have a base set of items, some with fixed prices, some priceable, and constraints on the subsets that are feasible for each follower. In this combinatorial setting, Briest et al. and Balcan et al. independently showed that the maximum revenue can be approximated to a factor of $H_k \sim \log k$, where k is the number of priceable items.

Our results are twofold. First, we strongly generalize the model by letting the follower minimize any continuous function plus a linear term over any compact subset of $\mathbb{R}_{\geq 0}^n$; the coefficients (or *prices*) in the linear term are chosen by the leader and determine her revenue. In particular, this includes the fundamental case of linear programs. We give a tight lower bound on the revenue of the leader, generalizing the results of Briest et al. and Balcan et al. Besides, we prove that it is strongly NP-hard to decide whether the optimum revenue exceeds the lower bound by an arbitrarily small factor. Second, we study the parameterized complexity of computing the optimal revenue with respect to the number k of priceable items. In the combinatorial setting, given an efficient algorithm for optimal follower solutions, the maximum revenue can be found by enumerating the 2^k subsets of priceable items and computing optimal prices via a result of Briest et al., giving time $O(2^k |I|^c)$ where $|I|$ is the input size. Our main result here is a W[1]-hardness proof for the case where the followers minimize a linear program, ruling out running time $f(k)|I|^c$ unless $\text{FPT} = \text{W}[1]$ and ruling out time $|I|^{o(k)}$ under the Exponential-Time Hypothesis.

1998 ACM Subject Classification G.2.0 [Discrete Mathematics] General

Keywords and phrases Algorithmic pricing, Stackelberg games, Approximation algorithms, Revenue maximization, Parameterized complexity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.46

1 Introduction

Pricing problems are fundamental in both economics and mathematical optimization. In this paper we study such pricing problems formulated as games, which are usually called *Stackelberg Pricing Games* [18]. In our setting, in order to maximize her revenue one player chooses prices for a number of items and one or several other players are interested in buying these items. Following the standard terminology, the player to choose the prices is called the *leader* while the other players are called *followers*. Depending on the follower's preferences,



© Toni Böhnlein, Stefan Kratsch, and Oliver Schaudt;
licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;


Article No. 46; pp. 46:1–46:13



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





■ **Figure 1** An instance of the Stackelberg Minimum Spanning Tree Game.

computing optimal prices can be a computational non-trivial problem. In a setting where followers have valuations over individual items only, the problem is simple. If, however, valuations become more complex, e.g., over whole subsets of items, pricing problems become much harder—also in a formal sense.

Largely the literature has focused on what we call the *combinatorial setting*: there is a set Y of items and one follower seeks to buy a feasible subset. Some of the items have fixed costs, the others have prices that are chosen by the leader. If the follower buys a feasible subset $S \subseteq Y$ of the items, he has to pay the sum of the fixed costs of the elements of S , plus the leader's prices of the bought elements. The leader's revenue is the sum of the prices of the priceable items in S . This can also be captured by defining a solution space X containing 0/1-vectors corresponding to the feasible subsets S of Y . The goal of the follower is then to minimize a given additive function $f: X \rightarrow \mathbb{R}$ that depends on both fixed and leader-chosen prices.

So-called Stackelberg Network Pricing Games became popular when Labbé et al. [15] used them to model road toll setting problems. In this game, the leader chooses prices for a subset of *priceable* edges in a network graph while the remaining edges have fixed costs. Each follower has a pair of vertices (s, t) and wants to buy a minimum cost path from s to t , taking into account both the fixed costs and the prices chosen by the leader. The work of Labbé et al. led to a series of studies of the Stackelberg Shortest Path Game. Roche et al. [16] showed that the problem is NP-hard, even if there is only one follower, and it has later been shown to be APX-hard [5, 14]. More recently, other combinatorial optimization problems were studied in their Stackelberg pricing version. For example, Cardinal et al. [9, 10] studied the Stackelberg Minimum Spanning Tree Game, proving APX-hardness and giving some approximation results. Moreover, a special case of the Stackelberg Vertex Cover Game in bipartite graphs has been shown to be polynomially solvable by Briest et al. [7].

To get more familiar with the setting, we briefly discuss an example of the Stackelberg Minimum Spanning Tree Game.

► **Example 1.** The left hand side of Figure 1 depicts an instance of the problem. Here the leader can choose the prices p_1 , p_2 and p_3 for the dashed edges, while the solid edges have fixed costs as displayed. To motivate the problem, think of the vertices as hubs in a network and of the edges as data connections. In this scenario, the followers are Internet Service Providers and want to connect all the hubs at minimum cost, thus want to compute a minimum spanning tree. The leader owns the dashed connections and wants to set prices, that yield a large revenue. Furthermore, there are competitors who own the solid connections and it is known how much they charge for their usage.

On the right hand side, an optimal pricing ($p_1 = p_2 = 5$ and $p_3 = 4$) and a corresponding minimum spanning tree are depicted. Thus the leader's revenue amounts to 9, which can be verified to be maximum. Observe that we can compute a minimum spanning tree without

any priceable edges; otherwise the leader's revenue is unbounded. In this example, the total cost of such a minimum spanning tree is 17. In contrast, if we set all prices to 0 and let the follower compute a minimum spanning tree, it has a total cost of 8. The difference of these two values, $17 - 8 = 9$, is an upper bound on the revenue of the leader, as explained later. This upper bound, which we denote by R , is sometimes called the optimal *social welfare* and will be important for our approximation result.

An important contribution to the study of Stackelberg Games was the discovery by Briest et al. [7]. They show that the optimal revenue can be approximated surprisingly well using a single-price strategy. For a single-price strategy the leader sets the same price for all of her priceable items. Basically, their result says the following: In any Network Pricing Game with k priceable items, there is some $\lambda \in \mathbb{R}_{\geq 0}$ such that, when assigning the price of λ to all priceable items at once, the obtained revenue is only a factor of H_k away from the optimal revenue. Here, $H_k = \sum_{i=1}^k 1/i$ denotes the k -th harmonic number. This discovery has been made independently, in a slightly different model, by Balcan et al. [2]. Actually, in both papers [2, 7] a stronger fact is proven: The single-price strategy yields a revenue that is at least R/H_k , where R is a natural upper bound on the optimal revenue. The definition of R was sketched in the example above, and is formally laid out later.

Our results. Our work focuses on pushing the knowledge on Stackelberg Pricing Games beyond the well-studied combinatorial setting, in order to capture more complex problems of the leader. This is motivated by the simple fact that the combinatorial setting is too limited to even model, e.g., a follower that has a minimum cost flow problem—a crucial problem in both, combinatorial optimization and algorithmic game theory. More generally, we might want to be able to give bounds and algorithms in the case when the follower has an arbitrary linear or even convex program. For example, the follower might have a production problem in which he needs to buy certain materials from the leader, but such pricing problems haven't been discussed in the literature so far.

We prove an approximation result that applies even to a setting generalizing linear and convex programs. In our model, the follower minimizes a continuous function f over a compact set of feasible solutions $x \in X \subseteq \mathbb{R}_{\geq 0}^n$. For some of the variables, say x_1 up to x_k , the leader can choose a price vector $p \in \mathbb{R}^k$. Now the follower chooses a vector $x \in X$ that minimizes his objective function $f(x) + \sum_{i=1}^k p_i x_i$. We remark that if X is a set containing 0/1-vectors only, then we are back to the classical combinatorial setting. The result of Briest et al. can be transferred to the case when f is non-additive, in view of their original proof. Moreover if X is a polytope and f is additive, the follower minimizes a linear program, which is an important special case.

In Section 2, we formally introduce this more general model and prove the following results.

- (i) The maximum revenue obtainable by the leader can be approximated to a logarithmic factor using a single-price strategy. This generalizes the above mentioned result of Briest et al. [7] not only to linear programs but to any kind of follower that is captured by our model.
- (ii) The analysis of point (i) is tight. There is a family of instances for which the single-price strategy yields maximum revenue. And this revenue meets the bound of point (i).
- (iii) It is strongly NP-hard to decide whether one can achieve a revenue that is only slightly larger than the one guaranteed by the single-price strategy. This holds true even in a very restricted combinatorial setting.

The second part of the paper deals with the parameterized complexity of Stackelberg Pricing Games (Section 3). To the best of our knowledge, the only result in this direction is

an XP-algorithm by Cardinal et al. [10] for the Stackelberg Minimum Spanning Tree Game in graphs of bounded treewidth.

In contrast to structural parameters like the treewidth of the input graph, we consider the complexity of the pricing problem when parameterized by the number of priceable variables (or items in the combinatorial setting). Our main result in this part is a $W[1]$ -hardness proof for the case that the optimization problem of the follower is a linear program, which is arguably one of the most interesting cases that does not fit into the combinatorial setting. This rules out algorithms of running time $f(k)|I|^c$ unless $FPT = W[1]$ for any function f and polynomial $|I|^c$ of the input size; it also rules out running time $|I|^{o(k)}$ under the Exponential-Time Hypothesis of Impagliazzo et al. [13]. This intractability result is complemented by a fairly simple FPT-algorithm with running time $O(2^k |I|^c)$ for any Stackelberg Game that fits into the combinatorial model, when provided with an efficient algorithm for finding optimal follower solutions. The algorithm enumerates all subsets of priceable items and applies a separation argument of Briest et al. [7] to compute optimal leader prices and revenue.

Related work. Most important for our work are the approximation results due to Briest et al. [7] and Balcan et al. [2], which were discussed above.

A larger body of work focuses on specific network problems in their Stackelberg Game version. Briest et al. [7] give a polynomial time algorithm for a special case of the Stackelberg Bipartite Vertex Cover Game. An algorithm with improved running time was later given by Baïou and Barahona [1]. As mentioned, Labbé et al. [15] use the Stackelberg Shortest Path Game to model road toll setting problems. They establish NP-hardness and use LP bilevel formulations to solve small instances. A combinatorial approximation algorithm with the same logarithmic approximation guarantee as the single-price strategy was given by Roch et al. [16]. Moreover, a lower bound on the approximability is due to Briest et al. [5]: they show that the Stackelberg Shortest Path Game is NP-hard to approximate within a factor of less than 2. This is an improvement over previous results by Joret [14] showing APX-hardness. Further research on the Stackelberg Shortest Path Game can be found in a survey by van Hoesel [17]. A similar problem, the Stackelberg Shortest Path Tree Game, is studied by Bilo et al. [4]. They give an NP-hardness proof and develop an efficient algorithm assuming that the number of priceable edges is constant. Later their algorithm was improved by Cabello [8].

Cardinal et al. [9] proved several positive approximation results for the Stackelberg Minimum Spanning Tree Game. In the same paper, they proved that the revenue maximization for this game is APX-hard and strongly NP-hard. We make use of their reduction in the proof of Theorem 7. Furthermore, Cardinal et al. [10] prove that this game remains NP-hard if the instances are planar graphs. However, the problem becomes polynomial-time solvable on graphs of bounded treewidth. Bilo et al. [3] consider the Stackelberg Minimum Spanning Tree Game for complete graphs.

Briest et al. [6] consider Stackelberg Games where the follower's optimization problem cannot be solved to optimality. Instead the follower uses a known approximation algorithm. They show that the Stackelberg Knapsack Game is NP-hard if the follower uses a greedy 2-approximate algorithm, and derive a $2 + \epsilon$ approximation algorithm. Furthermore, the revenue maximization problem can be solved efficiently in the Stackelberg Vertex Cover Game if the follower implements a primal-dual approximation.

2 Approximability of Stackelberg Pricing Games

In this section we first introduce the model in its full generality. Then we give a tight approximation result on the maximum revenue using a single-price strategy. We complement

this with a hardness proof by showing that deciding whether one can achieve a revenue that is only slightly larger than the one guaranteed by the single-price strategy is strongly NP-hard.

Our model. Let k and n be some natural numbers with $k \leq n$. The optimization problem of the follower is the following: He minimizes a continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ over his set of feasible solutions $X \subseteq \mathbb{R}_{\geq 0}^k \times \mathbb{R}_{\geq 0}^{n-k}$. The only restriction we put on X is that we require it to be a compact set, i.e., bounded and closed under limits.

The first move of the Stackelberg Pricing Game is made by the *leader*: She chooses a *price vector* $p \in \mathbb{R}^k$. Now the second move is made by the *follower*: He chooses an optimal solution (x^*, y^*) of the program

$$\begin{aligned} \min \quad & p^T x + f(x, y) \\ \text{s.t.} \quad & (x, y) \in X. \end{aligned}$$

The *revenue* of the leader is then given by the value $p^T x^*$. This value is her objective function and it is to be maximized. We remark that this problem has a bilevel structure.

To avoid technicalities, we make the following optimistic assumption: If the follower has several optimal solutions in X , we assume that the solution which is most profitable for the leader is chosen. That is the solution, which maximizes the value $p^T x^*$. Moreover, we assume that there is a point $(x, y) \in X$ with x being the k -dimensional all-zeroes vector. This simply means that the follower has a solution that does not give any revenue to the leader. Otherwise the revenue maximization problem would be unbounded which is obviously not an interesting case.

Before we can state our results for the new model, we need to introduce a number of technical notions. Given a feasible solution $(x, y) \in X$ of the follower, we call the value $\mathbf{1}^T x = \sum_{i=1}^k x_i$ the *mass* of (x, y) . A *single-price* is a price vector p of the form $p = \lambda \mathbf{1}$ where λ is some real number. Slightly abusing notation, we sometimes call λ the single-price. Note that when the leader uses a single-price the revenue is simply the mass of the follower's solution times the single-price.

Let M be the maximum mass the follower buys if the leader sets all her prices to 0. Formally,

$$\begin{aligned} M := \max \quad & \mathbf{1}^T x \\ \text{s.t.} \quad & \exists y \in \mathbb{R}^{n-k} : (x, y) = \arg \min \{ f(x', y') : (x', y') \in X \}. \end{aligned}$$

This value M exists since X is a compact set.

Consider, for example, the case where the follower seeks to buy a shortest s - t -path in a network. Then M is the maximum number of priceable edges of a shortest s - t -path in the network, when the priceable edges all have a price of 0 and thus can be bought for free by the follower.

Since X is a compact set, there exists a largest single-price at which the follower buys a non-zero mass from the leader. Let μ be the maximum mass the follower buys at this price. Consider again the case where the follower searches for a shortest s - t -path in a network. Then μ is the maximum number of priceable edges contained in a shortest s - t -path, under the largest single-price for which a shortest path exists that contains a priceable edge.

46:6 Revenue Maximization in Stackelberg Pricing Games

For all $m \in [0, M]$, let $\Delta(m)$ be the minimum price the follower has to pay if he buys a mass of at most m from the leader. More formally

$$\begin{aligned} \Delta(m) := \min & f(x, y) \\ \text{s.t.} & \mathbf{1}^t x \leq m \\ & (x, y) \in X, \end{aligned}$$

where $\mathbf{1}^t x$ is the mass bought by the follower. This minimum price $\Delta(m)$ exists, because X with the additional constraint of $\mathbf{1}^t x \leq m$ is again a compact set.

As observed by several authors (cf. [2, 5]), an upper bound on the optimum revenue is $R := \Delta(0) - \Delta(M)$. To see this, let r^* be the maximum revenue, and let (x^*, y^*) be the corresponding follower's solution. We have

$$r^* + \Delta(M) \leq r^* + \Delta(\mathbf{1}^T x^*) \leq r^* + f(x^*, y^*) \leq \Delta(0),$$

because $\Delta(0)$ is an upper bound on the objective value of the follower and Δ is non-increasing. We remark that R is indeed a tight upper bound, in the sense that there are examples of games where the maximum revenue equals R , e.g., the minimum spanning tree pricing problem described in the introduction.

As our first result shows, the maximum revenue of the leader is always reasonably close to R , unless the ratio M/μ is large. This is true even if the leader uses a single-price strategy.

► **Theorem 2.** *There is a single-price for the Stackelberg Pricing Game over X whose revenue is at least*

$$\frac{R}{1 + \ln\left(\frac{M}{\mu}\right)}.$$

This result extends previous work of Briest et al. [7] and Balcan et al. [2], who proved the above theorem in the combinatorial setting, i.e., for $X \subseteq \{0, 1\}^n$. We give the main idea of the proof of Theorem 2 in the following sketch and defer the full proof to the appendix.

Proof Sketch for Theorem 2. Our proof makes use of the following concept. For each $m \in (0, M]$, let $P(m)$ be the supremum of all single-prices for which the follower has an optimal solution with a mass of at least m from the leader. Using the compactness of X and the continuity of f one can prove that at the single-price of $P(m)$, the follower has an optimal solution of a mass of at least m . Thus, the supremum is indeed a maximum here.

It turns out that there are certain mass values which dictate the value of the function P . Let $T \subseteq (0, M]$ be the set of mass values t for which the follower does not have an optimal solution, at a single-price $P(t)$, with a mass more than t . The set T plays a key role in our proof, as the following claim indicates.

► **Claim 3.** *It holds that $\mu \in T$, $M \in T$ and $\mu = \min T$. Moreover, for all $m \in (0, M]$ it holds that $P(m) = \max_{t \in T \cap [m, M]} P(t)$.*

Recall that $\Delta(m)$ is defined as the price the follower has to pay for the non-priceable variables if he buys a mass of at most m from the leader. Similar to the proof of Briest et al. [7] for the combinatorial setting, we next show that the functions P and Δ are closely related. In our case, however, we have to deal with several difficulties that arise because we allow for non-discrete optimization problems of the follower.

Consider the lower convex hull H of the point set $\{(m, \Delta(m)) : 0 \leq m \leq M\}$. Let ∂H be the lower border of H , and let $\hat{\Delta} : [0, M] \rightarrow \mathbb{R}$ be the function for which $(m, \hat{\Delta}(m)) \in \partial H$ for all $m \in [0, M]$. We remark that, since $\hat{\Delta}$ is convex and decreasing,

$$D_- \hat{\Delta}(m) = \sup_{\ell < m} \frac{\hat{\Delta}(m) - \hat{\Delta}(\ell)}{m - \ell}, \text{ for each } m \in (0, M].$$

Here, D_- denotes the *lower left Dini derivative* of $\hat{\Delta}$. It is defined, for all $m \in (0, M]$, by

$$D_- \hat{\Delta}(m) = \liminf_{h \rightarrow 0^-} \frac{\hat{\Delta}(m) - \hat{\Delta}(m+h)}{h}.$$

As our main claim shows, the values of $P(m)$ and $D_- \hat{\Delta}(m)$ are essentially equal.

► **Claim 4.** *Except for a set of measure 0, it holds for all $m \in (0, M)$ that $D_- \hat{\Delta}(m) = -P(m)$.*

To establish Claim 4 is indeed the difficult part of the whole proof. We skip the details due to space constraints.

In the calculation that finishes the proof we make use of the inverse operation of the lower left Dini derivative, the so-called *lower left Dini integral*, denoted $(LD) \int$. For more background we refer to the article of Hagood and Thomson [11].

We have

$$R = \Delta(0) - \Delta(M) = \hat{\Delta}(0) - \hat{\Delta}(M) = \lim_{\epsilon \rightarrow 0^+} \hat{\Delta}(\epsilon) - \hat{\Delta}(M) = \lim_{\epsilon \rightarrow 0^+} (LD) \int_{\epsilon}^M -D_- \hat{\Delta}(m) dm.$$

Due to Claim 4, the integral

$$\lim_{\epsilon \rightarrow 0^+} (LD) \int_{\epsilon}^M P(m) dm$$

is well defined and equals

$$\lim_{\epsilon \rightarrow 0^+} (LD) \int_{\epsilon}^M -D_- \hat{\Delta}(m) dm.$$

Recall that, by Claim 3, $\mu = \min T$ and so $P(m) = P(\mu)$ for all $m \in (0, \mu]$. Hence,

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} (LD) \int_{\epsilon}^M P(m) dm &= \lim_{\epsilon \rightarrow 0^+} (LD) \int_{\epsilon}^{\mu} P(m) dm + (LD) \int_{\mu}^M P(m) dm \\ &= \mu \cdot P(\mu) + (LD) \int_{\mu}^M P(m) dm. \end{aligned}$$

Let r be the maximum revenue achieved by the single-price strategy. Note that r is at least the revenue at the single-price $P(m)$, for each $m \in (0, M]$, which is in turn at least $m \cdot P(m)$. We thus have

$$\begin{aligned} \mu \cdot P(\mu) + (LD) \int_{\mu}^M P(m) dm &= \mu \cdot P(\mu) + (LD) \int_{\mu}^M \frac{m \cdot P(m)}{m} dm \\ &\leq r + (LD) \int_{\mu}^M \frac{r}{m} dm \\ &= r + r \cdot (\ln(M) - \ln(\mu)) \\ &= r \left(1 + \ln \left(\frac{M}{\mu} \right) \right). \end{aligned}$$

This shows that $R \leq r(1 + \ln(M/\mu))$, as desired. ◀

Balcan et al. [2] and Briest et al. [7] extend their result to the situation when there are several followers.¹ The same approach works in our generalized model.

Assume there are ℓ followers and each follower has its own optimization problem. Formally, the i -th follower minimizes his objective function $p^T x^i + f_i(x^i, y^i)$ where (x^i, y^i) belongs to the set $X_i \subseteq \mathbb{R}_{\geq 0}^k \times \mathbb{R}_{\geq 0}^{n-k}$ of his feasible solutions, $i = 1, \dots, \ell$. The pricing vector p appears in the objective function of every follower and is again set by the leader in order to maximize her revenue $\sum_{i=1}^{\ell} p^T x^i$. The difficulty here is that, while each follower has an individual optimization problem, the leader can set only one price vector for all followers at once. There is, however, a canonical way of reducing the pricing game to the case of a single follower. To this end, we consider the pricing game with respect to follower i only, and let μ_i (resp. M_i) be the minimum non-zero mass (resp. the maximum mass) bought by follower i . Moreover, let R_i be the upper bound on the revenue with respect to follower i .

► **Corollary 5.** *There is a single-price for the Stackelberg Pricing Game with ℓ followers whose revenue is at least*

$$\frac{\sum_{i=1}^{\ell} R_i}{1 + \ln \left(\frac{\sum_{i=1}^{\ell} M_i}{\min_{i=1}^{\ell} \mu_i} \right)}.$$

To see this, consider a single follower with the feasible subset $X = X_1 \times X_2 \times \dots \times X_{\ell}$. It is easy to see that we have $M = \sum_{i=1}^{\ell} M_i$ and $R = \sum_{i=1}^{\ell} R_i$ in this game. Moreover, the smallest non-zero mass μ bought by the newly defined single follower is at least the minimum smallest non-zero mass bought by one of the ℓ followers, that is $\min_{i=1}^{\ell} \mu_i$. Now applying Theorem 2 to the single follower yields Corollary 5.

Theorem 2 is tight in the following sense.

► **Proposition 6.** *There are Stackelberg Pricing Games of arbitrarily large R and M in which the optimum revenue equals*

$$(1 + o(1)) \cdot \frac{R}{1 + \ln \left(\frac{M}{\mu} \right)}.$$

This holds true even for games in which the follower minimizes a linear objective function of the form $p^T x + c^T y$ over a uniform matroid.

Note that, in the above statement, every possible pricing is considered and not just single-price strategies. In other words, the lower bound in Theorem 2 is tight not only for single-price strategies, but for arbitrary pricings. So far, it was known that there are combinatorial pricing games where the optimum revenue is in $O(R/\log k)$, where k is the number of priceable elements (cf. [5]). The merit of Proposition 6 is that it shows tightness of Theorem 2 up to a factor of $1 + o(1)$, which is best possible. This fact, and the construction given in the proof of Proposition 6, enable us to prove the following hardness result.

► **Theorem 7.** *Fix a sufficiently small rational number $\epsilon > 0$, and consider a Stackelberg Pricing Game where the follower minimizes an objective function of the form $p^T x + c^T y$ over a matroid. It is strongly NP-hard to decide whether there is some pricing of revenue at least*

$$(1 + \epsilon) \cdot \frac{R}{1 + \ln \left(\frac{M}{\mu} \right)}.$$

¹ In this paper we consider the case of *unlimited supply*, meaning that the followers buy their favorite solution independently of each other.

Due to space constraints, the proof of Theorem 7 is deferred to the appendix. In the statement of the above theorem, we assume that the matroid is given by its ground set and a membership oracle. Thus, μ , M , and R can be computed in polynomial time. Moreover, it will be clear from the proof that the natural logarithm can be computed in polynomial time with the precision needed to decide the problem.

Let us remark that the assumption of ϵ being *sufficiently small* is merely technical, and can be dropped by giving a more careful hardness reduction. The message of the above theorem is, however, that there is a sharp contrast between the revenue guaranteed by the simple single-price strategy given in Theorem 2 and anything more than that.

3 Parameterized complexity of Stackelberg Pricing Games with few priceable objects

In this section we study the parameterized complexity of Stackelberg Pricing Games when parameterized by the number of priceable objects. Intuitively, this addresses the question of whether there are improved algorithms for the case that only few objects are priceable. We give a general positive result for the combinatorial model. Our main result in this part, however, is a hardness proof for the linear programming case; we begin with the latter.

In the LP-PRICING problem there is a linear program over which the follower minimizes. The leader may choose the price, i.e. target function coefficient, of k specified variables. Her revenue is determined by the corresponding (weighted) sum over these variables.

LP-PRICING

Input: A linear program with k priceable variables and $\lambda \in \mathbb{Q}$.

Question: Is there a price vector whose revenue is at least λ ?

We prove that this problem is at least as hard to solve as the parameterized k -clique problem. The hardness proof creates linear programs with only non-negative variables and non-negative target function over which the follower seeks to minimize. As such it proves hardness also for our more general model parameterized by number of priceable variables.

► **Theorem 8.** LP-PRICING is $W[1]$ -hard when parameterized by the number k of priceable variables.

The theorem is proven by a reduction from the well-known (parameterized) MULTI-COLORED CLIQUE problem. Therein, we are given a k -partite graph G (or, equivalently, a properly k -colored graph) and have to determine whether G contains a clique on k vertices; the problem is $W[1]$ -hard with respect to parameter k . Thus, unless $FPT = W[1]$, there is no algorithm running in time $f(k)n^c$ for instances of size n . Moreover, under the Exponential-Time Hypothesis [13] the reduction implies that there is no $O(n^{o(k)})$ time algorithm for LP-PRICING. In instances created by the reduction the leader can effectively enforce the choice of the k clique vertices by setting appropriate prices for k variables; the remaining variables are used to verify the choice and a certain revenue threshold can only be attained if there is indeed a k -clique. The behavior of these k priceable variables is quite similar to integer variables, as they can be shown to only take specific values from a finite set in solutions meeting the threshold (one value corresponding to each vertex of G). This arguably gives our parameterized LP PRICING problem some similarity to the MIXED ILP FEASIBILITY problem parameterized by the number of integer variables. Interestingly, the latter problem is FPT due to a classic result of Lenstra [12]. Due to space constraints we will restrict ourselves to an informal description of the reduction complemented by some intuition.

Proof Sketch for Theorem 8. We give a parameterized reduction from the $W[1]$ -hard MULTICOLORED CLIQUE(k) problem. Therein, we are given a k -partite graph $G = (V_1, \dots, V_k, E)$ and have to determine whether it contains a clique of size k , i.e., a clique containing exactly one vertex from each partite set. The reduction is polynomial-time computable and creates an instance of LP-PRICING with $k + 1$ priceable variables, proving $W[1]$ -hardness of LP-PRICING.

Much effort in the design of the linear program and the correctness proof goes into setting up constraints on k pairs of variables (x_i, y_i) and forcing each pair to take a value in $\{(p_1, q_1), \dots, (p_n, q_n)\}$; each x_i is priceable and the y_i have fixed prices. We have $p_1 > p_2 > \dots > p_n$ and $q_1 < q_2 < \dots < q_n$ with the high-level idea that changing the price of x_i will make a particular point (p_j, q_j) optimal. Each choice of $(x_i, y_i) = (p_j, q_j)$ corresponds to the selection of one vertex (i, v_i) from the i th partite set V_i of G . Additionally, there are variables $z_{i,j} \geq |y_i - q_j|$ that serve as indicators for whether $y_i = q_j$, and clique-testing constraints on pairs of variables $z_{i,u}, z_{j,v}$ for every non-edge $\{(i, u), (j, v)\}$ of G to prevent choosing both (i, u) and (j, v) for the clique. Notably, there is a single slack variable y_0 present in all clique-testing constraints that is connected to a priceable variable x_0 in such a way that the maximum revenue c_0 for x_0 is only possible if the slack y_0 is not needed, i.e., if all clique-testing constraints are active.

The reader will have noticed that there are no constraints that we could pose in a linear program that would enforce $(x_i, y_i) \in \{(p_1, q_1), \dots, (p_n, q_n)\}$. Instead, the points (p_j, q_j) are the extremal points (when projected to two dimensions) of certain constraints on each pair x_i and y_i ; we call these the *core constraints*. There are intended prices r_1, \dots, r_n with the goal of showing that setting the leader price d_i of x_i to r_j leads to $(x_i, y_i) = (p_j, q_j)$ being uniquely optimal. This would probably be easier if we could restrict to $d_i \in \{r_1, \dots, r_n\}$, but we are not allowed to do so. Instead, the strategy is roughly as follows:

1. Consecutive pairs of points (p_j, q_j) and (p_{j+1}, q_{j+1}) are chosen in such a way that for $d_i = r_j$ all choices of (x_i, y_i) on the line segment defined by the two points give the same follower cost including updates to indicator variables $z_{i,j}$. (This is rather helpful for replacement arguments, which we use to disprove solutions not following intended behavior, but does not cover the variable y_0 .) The leader payoff in this case is highest for $x_i = p_j$ as $p_j > p_{j+1}$, and we have $r_j := \frac{1}{p_j}$ so that the payoff is exactly 1 in this case.
2. A critical issue is that we must make sure that the leader cannot make a profit larger than 1 with any variable x_i since that may be more beneficial than attempting to gain the revenue for having the clique-testing constraints fulfilled without increasing the slack variable y_0 above 0.
3. A significant part of the proof hence goes towards a technical claim that allows to rule out both profit greater than one for any variable x_i and the case that $p_s > x_i > p_{s+1}$, i.e., that x_i lies between two intended coordinates. To this end, the core constraints on pairs (x_i, y_i) contain varying numbers of slack variables $w_{i,1}, \dots, w_{i,n}$ of fixed prices r_1, \dots, r_n . These are placed in such a way that they provide alternative ways of satisfying the core constraints when $d_i \notin \{r_1, \dots, r_n\}$, without keeping the follower from actually paying the full price of x_i in other cases.

We remark that the correctness proof of course also requires the reverse direction of ensuring leader payoff at least $k + c_0$ if G has a clique of size k . Here we are in the easier situation of being able to select leader prices and fixing a suggested solution. We then prove feasibility, which is straightforward, and optimality, which requires combining most of the constraints into a lower bound on the optimal follower cost that matches that of our suggested solution. ◀

Theorem 8 is in sharp contrast to the combinatorial setting, where under mild assumptions one can see the problem to be fixed-parameter tractable. Here we assume that $X \subseteq \{0, 1\}^n$ but put no further restriction on the follower's objective function $f : X \rightarrow \mathbb{R}$. In particular, this model covers the classical setting where each item has a fixed cost and if the follower buys a set S of items, he has to pay the sum of the fixed costs of the elements of S , plus the leader's price of the bought elements.

► **Theorem 9.** *Assume that $X \subseteq \{0, 1\}^n$, and that we are given a polynomial-time algorithm to compute an optimal solution of the follower for given leader prices $p \in \mathbb{R}^k$. Then the computation of optimal prices and optimal leader revenue is fixed-parameter tractable, with running time $\mathcal{O}(2^k n^c)$, when parameterized by the number of priceable items.*

In the above statement we make the natural assumption that the input size is at least $n + \text{size}(f)$, where $\text{size}(f)$ denotes the maximum length of the binary encoding of any value f can take.

Proof Sketch for Theorem 9. Due to space constraints the proofs of the claims in this section are deferred to the appendix. We need the following proposition by Briest et al. [7]. It says that if the leader wants to force the follower to pick a certain solution, she can compute suitable prices in polynomial time. We state their result in a slightly more general fashion than in the original paper. Indeed, Briest et al. were only concerned with the case when f is additive but the proof did not make use of the additivity of f at all.

► **Proposition 10** (Briest et al. [7]). *Given a vector $z \in X$, one can compute an optimal price vector p such that z is an optimal solution of the follower with respect to p , or decide that such a price vector does not exist, in polynomial time.*

Our algorithm works as follows. For each vector $x \in \{0, 1\}^k$ we compute a price vector p_x , if exists, such that

- (a) there is some $y \in \{0, 1\}^{n-k}$ such that the vector (x, y) is an optimal solution of the follower with respect to the price vector p_x ,
- (b) subject to (a) the revenue $p_x^T x$ is maximum.

When this procedure is finished we choose the vector $\hat{x} \in \{0, 1\}^k$ with maximum value of $p_{\hat{x}}^T \hat{x}$, and output the price vector $p_{\hat{x}}$. As the next claim shows, this price vector is the optimum solution.

► **Claim 11.** *The output $p_{\hat{x}}$ is an optimal price vector for the leader.*

In the remainder of the proof we show how to compute p_x for a fixed candidate vector $x \in \{0, 1\}^k$. First we aim to find a vector $y_x \in \{0, 1\}^{n-k}$ such that $(x, y_x) \in X$ and, subject to this, $f(x, y_x)$ is minimum. Note that possibly such a vector y_x does not exist. To find a vector y_x , or decide that none exists, we define a price vector p by setting

$$p_i = \begin{cases} -M & \text{if } x_i = 1, \\ M & \text{if } x_i = 0, \end{cases}$$

for all $i \in [k]$. Here, M is a number that is large enough to ensure that

$$\text{for all } (x', y'), (x'', y'') \in X, \text{ it holds that } f(x', y') - f(x'', y'') < M. \quad (1)$$

As both values $|f(x', y')|$ and $|f(x'', y'')|$ are bounded by $2^{\text{size}(f)}$, we may simply put $M = 2^{\text{size}(f)+1} + 1$.

Now we compute an optimal solution of the follower with respect to the price vector p , say (x^*, y^*) , using the assumed polynomial-time algorithm.

► **Claim 12.** *If for some $y \in \{0, 1\}^{n-k}$ it holds that $(x, y) \in X$, then $x^* = x$.*

If $x^* \neq x$, Claim 12 implies that there is no price vector satisfying (a). Thus, we may safely abort the process and go over to the next candidate vector x . Otherwise if $x^* = x$, we put $y_x = y^*$.

► **Claim 13.** *$f(x, y_x) \leq f(x, y)$ holds for all $y \in \{0, 1\}^{n-k}$ with $(x, y) \in X$.*

Next, we use Proposition 10 to compute a price vector p_x such that the vector (x, y_x) is optimal for the follower and, subject to that, $p_x^T x$ is maximum. Note that this price vector does exist since, e.g., the vector p is a feasible price vector. By construction, p_x has the property (a). So far, we only know that $p_x^T x$ is maximum subject to the condition that under the price vector p_x the vector (x, y_x) is an optimal solution of the follower.

► **Claim 14.** *Subject to (a), the revenue $p_x^T x$ is maximum.*

As the running time of the whole algorithm is $O(2^k \cdot (n + \text{size}(f))^{O(1)})$, the proof is complete. ◀

4 Conclusion and future work

The basis for the first part of this paper were the results of Briest et al. [7] and Balcan et al. [2] who gave a lower bound on the optimal revenue in Stackelberg Network Pricing Games. We proved that this bound carries over to a much more general setting, where, basically, the follower minimizes a continuous function over a compact set of points. This model captures important settings that are not covered by the classical combinatorial model. For example, the case when the follower is minimizing a linear program, e.g., a minimum cost flow problem.

The proven lower bound also holds if a single-price strategy is applied, and it is tight up to a factor of $(1 + o(1))$. Moreover, we used this tightness example to show that it is strongly NP-hard to decide whether the revenue of an optimal pricing exceeds the lower bound by an arbitrarily small linear factor.

In the second part of the paper we studied the parameterized complexity of the revenue maximization problem. It turned out that in the combinatorial setting (i.e., when the follower only has 0/1-valued solutions) there is an elegant FPT algorithm. Once we leave this regime, however, things become more difficult. Indeed, if the follower has an optimization problem in the form of a linear program, the revenue maximization problem becomes W[1]-hard and is thus most likely not FPT.

Several central questions remain. Most importantly, one should consider multiple-follower scenarios. An intriguing model is when the particular resources have a limited supply. In the combinatorial setting, a limited supply means that every item to be sold is available only a limited number of times. Now the followers come one by one, in a certain order, and buy according to their preferences and the prices set by the leader. Balcan et al. [2] prove a tight lower bound on the revenue obtained by the single-price strategy. In the non-combinatorial model, the limited supply might be translated to a constraint of the form $(x, y) \leq s$, where $s \in \mathbb{R}_{\geq 0}^n$ is a fixed vector that is added to the usual constraint $(x, y) \in X$ in the optimization problem of the followers.

References

- 1 Mourad Baïou and Francisco Barahona. Stackelberg Bipartite Vertex Cover and the Preflow Algorithm. to appear in *Algorithmica*, 2016.

- 2 Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item Pricing for Revenue Maximization. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008)*, Chicago, IL, USA, June 8-12, 2008, pages 50–59, 2008. doi:10.1145/1386790.1386802.
- 3 Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Specializations and Generalizations of the Stackelberg Minimum Spanning Tree Game. *Theoretical Computer Science*, 562:643–657, 2015.
- 4 Davide Bilò, Luciano Gualà, Guido Proietti, and Peter Widmayer. Computational Aspects of a 2-player Stackelberg Shortest Paths Tree Game. In *Internet and Network Economics*, pages 251–262. Springer, 2008.
- 5 Patrick Briest, Parinya Chalermsook, Sanjeev Khanna, Bundit Laekhanukit, and Danupon Nanongkai. Improved Hardness of Approximation for Stackelberg Shortest-Path Pricing. In *Internet and Network Economics*, pages 444–454. Springer, 2010.
- 6 Patrick Briest, Luciano Gualà, Martin Hoefer, and Carmine Ventre. On stackelberg pricing with computationally bounded customers. *Networks*, 60(1):31–44, 2012. doi:10.1002/net.20457.
- 7 Patrick Briest, Martin Hoefer, and Piotr Krysta. Stackelberg Network Pricing Games. *Algorithmica*, 62(3-4):733–753, 2012. doi:10.1007/s00453-010-9480-3.
- 8 Sergio Cabello. Stackelberg Shortest Path Tree Game, Revisited. *arXiv preprint arXiv:1207.2317*, 2012.
- 9 J. Cardinal, E. D. Demaine, S. Fiorini, G. Joret, S. Langerman, I. Newman, and O. Weimann. The Stackelberg Minimum Spanning Tree Game. *Algorithmica*, 59:129–144, 2011.
- 10 Jean Cardinal, Erik D Demaine, Samuel Fiorini, Gwenaël Joret, Ilan Newman, and Oren Weimann. The Stackelberg Minimum Spanning Tree Game on Planar and Bounded-Treewidth Graphs. *Journal of combinatorial optimization*, 25(1):19–46, 2013.
- 11 J. W. Hagoood and B. S. Thomson. Recovering a Function from a Dini Derivative. *The American Mathematical Monthly*, 113:34–46, 2006.
- 12 Jr. H. W. Lenstra. Integer Programming with a Fixed Number of Variables. *Mathematics of Operations Research*, 8(4):538–548, 1983. doi:10.1287/moor.8.4.538.
- 13 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- 14 Gwenaël Joret. Stackelberg Network Pricing is Hard to Approximate. *Networks*, 57(2):117–120, 2011.
- 15 M. Labbé, P. Marcotte, and G. Savard. A Bilevel Model of Taxation and Its Application to Optimal Highway Pricing. *Management Science*, 44:1608–1622, 1998.
- 16 S. Roche, G. Savard, and P. Marcotte. An approximation algorithm for Stackelberg network pricing. *Networks*, 46:57–67, 2005.
- 17 S. van Hoesel. An overview of Stackelberg pricing in networks. *European Journal of Operational Research*, 189:1393–1402, 2008.
- 18 H. von Stackelberg. *Marktform und Gleichgewicht*. Springer, 1934.