Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Efficiency and Complexity of Price Competition among Single-Product Vendors

Ioannis Caragiannis, Xenophon Chatzigeorgiou, Panagiotis Kanellopoulos
George A. Krimpas, Nikos Protopapas, Alexandros A. Voudouris
University of Patras & CTI "Diophantus"

Abstract

Motivated by recent progress on pricing in the AI literature, we study marketplaces that contain multiple vendors offering identical or similar products and unit-demand buyers with different valuations on these vendors. The objective of each vendor is to set the price of its product to a fixed value so that its profit is maximized. The profit depends on the vendor's price itself and the total volume of buyers that find the particular price more attractive than the price of the vendor's competitors. We model the behaviour of buyers and vendors as a two-stage full-information game and study a series of questions related to the existence, efficiency (price of anarchy) and computational complexity of equilibria in this game. To overcome situations where equilibria do not exist or exist but are highly inefficient, we consider the scenario where some of the vendors are subsidized in order to keep prices low and buyers highly satisfied.

1 Introduction

We focus on marketplaces that contain multiple vendors offering a single product and unit-demand buyers. For example, we may think of software development companies, each offering an operating system. Each potential user is interested in buying one operating system from some software company and has preferences over the different options available in the market. Her final choice depends not only on her preferences but also on the prices of the available products; eventually, each user will choose the product with the best value for money, or will simply abstain from purchasing a product if the available options are not satisfactory for her. In turn, vendors are aware of this buyer behaviour and aim to set the price of their product to a value that will maximize their profit. In particular, the dilemma a vendor faces is to select between a very small price that will guarantee a large market share or a huge price that will be attractive only to a few buyers. Of course, there are usually many options in between, and coming up with a pricing that will maximize profits in such an environment is rather challenging.

We model the above scenario as a two-stage full-information game (with both the vendors and the buyers as players) which we call a *price competition game*. In the first stage, each vendor selects the price of its product among a set of viable price values (i.e., the price values that are above a fixed production cost per unit of product). Buyers have unit demands and (possibly different) valuations for vendors. Together with the valuations of buyers, a vector of prices (with one price per vendor) determines in a second stage the most attractive vendor for each buyer. Each vendor has full information about the valuations of buyers and can predict their behaviour. The objective of each vendor is to set its price so that its profit (i.e., volume of buyers it attracts times the difference of price and production cost) is maximized given the prices of the other vendors.

We present a list of results for these price competition games. Our starting point is the observation that equilibria (i.e., buyers-to-vendors assignments and corresponding prices so that all vendors and all buyers are satisfied) are guaranteed to exist only when all buyers have the same valuations; price competition games with buyers belonging to at least two different types (with respect to their valuations) may not have equilibria. Even when equilibria exist, they may be highly suboptimal. We use the notion of the price of anarchy (introduced by Koutsoupias and Papadimitriou [2009]; see also Papadimitriou [2001]) to quantify how low the social welfare of equilibria can be compared to the optimal one. The social welfare is essentially the sum of buyer utilities and vendor profits. We also formulate several variations of equilibrium computation problems and present complexity results about them. These range from polynomial-time algorithms (e.g., for the problem of determining prices that form an equilibrium together with a given buyers-to-vendors assignment) to hardness results (e.g., for the general problem of deciding whether a given price competition game admits an equilibrium). Motivated by the negative results on the existence and quality of equilibria, we investigate whether efficient buyers-to-vendors assignments can be enforced as equilibria by *subsidizing* the vendors. Our main contribution here is conceptual: subsidies can indeed overcome the drawbacks of price competition. Our technical contributions include tight bounds on the amount of subsidies sufficient to enforce a social welfare-maximizing buyers-to-vendors assignment as an equilibrium, as well as inapproximability results for the problem of minimizing the amount of subsidies sufficient to do so.

Our model is very similar to (and actually inspired from) the one considered by Meir et al. [2014] who focus on the impact of discounts (i.e., prices that are decreasing functions of demand) on vendors' profit compared to fixed prices. After observing that price discounts have no impact at all in the full-information setting, they mostly focus on a Bayesian setting with uncertainty on buyers' valuations. In contrast, we restrict our attention to the full-information model and consider only fixed prices. As we will see, this simple setting is very rich from the computational point of view. With the work of Meir et al. [2014] as an exception, our assumptions differ significantly from most of the literature on price competition. For example, unlike early models such as the ones proposed by Cournot and Bertrand (see the book of Mas-Colell et al. [1995]) as well as very recent refinements (e.g., the work of Babaioff et al. [2013]), we assume that all vendors have unlimited supply. Also, contrary to other recent models that consider buyers with combinatorial valuations for bundles of different products as in the papers of Guruswami et al. [2005], Chawla and Roughgarden [2008], Babaioff et al. [2014], Lev et al. [2015] and more, we specifically assume that each buyer is interested in obtaining just a single product. In this way, the decision each buyer faces is rather trivial and this allows us to concentrate on the competition between the vendors. On a more technical level, we implicitly assume an infinite number of buyers and use the notion of buyer types to distinguish between sets of buyers; this is a less important difference of our model to previous work on pricing.

The use of subsidies in price competition games suggests yet another way of introducing external monetary incentives in games; such incentives (or disincentives) have been considered in many different contexts. Much of the work in *mechanism design* uses such incentives to motivate players to act truthfully (see Nisan [2007] for an introduction to the field). The (apparently non-exhaustive) list also includes their use in *cooperative game theory* in order to encourage coalitions of players to reach stability [Bachrach *et al.*, 2009] and as a means to stabilize normal form games [Monderer and Tennenholtz, 2004]. As in [Augustine *et al.*, 2012] and [Buchbinder *et al.*, 2010], the use of monetary incentives in the current paper aims to improve efficiency. Monetary disincentives like *taxes* have been used to improve the efficiency of network routing (see Cole et al. [2006] and the references therein for a relatively recent approach that extends early developments in the literature of the *economics of transportation*) and, in the recent AI literature, in boolean games [Wooldridge *et al.*, 2013].

The rest of the paper is structured as follows. We begin with preliminary definitions in Section 2. Then, we consider questions about the existence of equilibria in price competition games and their price of anarchy in Section 3. We formulate computational problems for equilibria and study related complexity questions in Section 4. We investigate the potential of subsidizing specific vendor prices in Section 5 and, finally, we conclude with open questions in Section 6.

2 Preliminaries

Our setting includes a set M containing m vendors targeting a large population of buyers. The buyers are classified into n buyer types from a set N. We denote by μ_i the volume of buyer type i. Each of these buyers has a non-negative valuation v_{ij} for vendor j (representing the satisfaction each buyer of type i has when buying the product of vendor j). Each vendor j has a non-negative cost c_j per unit of product; we refer to c_j as the *production cost* of vendor j. The objective of each vendor j is to determine a price p_j for its product; naturally, $p_j \geq c_j$ so that the vendor always has non-negative profit. A *price vector* $\mathbf{p} = (p_1, ..., p_m)$ (containing a price per vendor) defines a *demand set* $D_i(\mathbf{p})$ which, for each buyer type i, denotes the set of vendors that maximize the utility of the buyers of type i, i.e., $D_i(\mathbf{p}) = \arg\max_{j\in M}\{v_{ij} - p_j\}$. Intuitively, the demand set for buyers of type i consists of the most attractive vendors for these buyers. We assume that the operator $\arg\max_{j\in M}$ returns (a set containing) an artificial vendor which represents an "abstain" option that a buyer has when its maximum utility (over all vendors) is non-positive. With a small abuse of notation, we introduce an extra vendor into M in order to represent this abstain option for buyers; this vendor has production cost of 0, it always has a price of 0, and the valuations of buyers for it are 0. A *buyers-to-vendors assignment* (or, simply, an *assignment*) is represented by an $n \times (m+1)$ matrix $\mathbf{x}$ which denotes how the volume of the buyers of each type is split among different vendors. In particular, the entry x_{ij} denotes the volume of buyers of type i that are assigned to vendor j and it must be $\sum_{j\in M} x_{ij} = \mu_i$ for every buyer type i. An assignment $\mathbf{x}$ is *consistent* to a price vector $\mathbf{p}$ if $x_{ij} > 0$ implies $j \in D_i(\mathbf{p})$. We can interpret such an assignment as maximizing the utility of buyers given the price vector $\mathbf{p}$. We will denote by $t_i(\mathbf{x}, \mathbf{p}) = \sum_{j\in M} x_{ij}(v_{ij} - p_j)$ the total utility of buyers of type i given a price vector $\mathbf{p}$ and a consistent assignment $\mathbf{x}$.

We study the game induced among vendors and buyers and use the term *price competition game* to refer to it. This can be thought of as a two stage game. At a first stage, the strategy of each vendor is its price. At a second stage, the buyers respond to these prices as described above. The utility of vendor j, when the vendors use a price vector $\mathbf{p}$ and the buyers are assigned to vendors according to an assignment $\mathbf{x}$ that is consistent to $\mathbf{p}$, is defined as $u_j(\mathbf{x}, \mathbf{p}) = (p_j - c_j)\sum_{i\in N} x_{ij}$. Vendors are utility-maximizers. A price vector $\mathbf{p}$ and a consistent assignment $\mathbf{x}$ form a *(pure Nash) equilibrium* when for every vendor j, the price p_j maximizes the utility $u_j(\mathbf{y}, (p'_j, \mathbf{p}_{-j}))$ among all prices $p'_j \geq c_j$ and all assignments $\mathbf{y}$ that are consistent to $(p'_j, \mathbf{p}_{-j})$. Here, the notation $(p'_j, \mathbf{p}_{-j})$ is used to represent the price vector where all vendors besides j use the prices in $\mathbf{p}$ and vendor j has deviated to price p'_j.

The *social welfare* of an assignment $\mathbf{x}$ is defined as

$$\text{SW}(\mathbf{x}) = \sum_{i\in N}\sum_{j\in M} x_{ij}(v_{ij} - c_j).$$

This definition does not require the assignment $\mathbf{x}$ to be consistent to a price vector and can be used to define the *optimal social welfare* as

$$\text{SW}^* = \sum_{i\in N} \mu_i \max_{j\in M}\{v_{ij} - c_j\}.$$

When the assignment $\mathbf{x}$ is consistent to a price vector $\mathbf{p}$, the social welfare can be equivalently seen as the total utility of vendors and buyers since

$$\begin{aligned}\mathrm{SW}(\mathbf{x}) &= \sum_{i\in N}\sum_{j\in M} x_{ij}(v_{ij}-c_j)\\ &= \sum_{i\in N}\sum_{j\in M} x_{ij}(v_{ij}-p_j) + \sum_{j\in M}(p_j-c_j)\sum_{i\in N} x_{ij}\\ &= \sum_{i\in N} t_i(\mathbf{x},\mathbf{p}) + \sum_{j\in M} u_j(\mathbf{x},\mathbf{p}).\end{aligned}$$

The *price of anarchy* of a price competition game is the ratio of the optimal social welfare over the minimum social welfare among all equilibria. Of course, this is well-defined only for price competition games that do have equilibria.

In the following, we sometimes use the abbreviation x^+ instead of $\max\{0,x\}$ and write $[\ell]$ instead of the set $\{1,2,...,\ell\}$ for an integer $\ell \geq 1$.

3 Existence and quality of equilibria

As a warm up, we present a negative result that reveals a strong relation of the price of anarchy of price competition games to the number of buyer types.

Lemma 1. *There are one-vendor price competition games with price of anarchy that is arbitrarily close to n.*

Proof. Consider a price competition game with n buyer types and one vendor with a production cost of 0. Let $\alpha \in (0,1)$. The volume of a buyer type $i \in N$ is $\mu_i = \alpha^{i-1}$. The valuation of buyers of type i is $v_i = \left(\sum_{j=i}^{n}\mu_j\right)^{-1}$ for $i \in [n-1]$ and $v_n = (1+\alpha)/\mu_n$. By setting its price up to v_i for $i \in [n-1]$, the vendor can only get a utility of at most $v_i\sum_{j=i}^{n}\mu_j$ by attracting the buyers of type $i, i+1, ..., n$. By the definition of v_i, this utility is at most 1. This is smaller than the utility the vendor would have by selecting a price of v_n and attracting only the buyers of type n (the remaining buyers simply abstain). This is an equilibrium in which the utility of the vendor (as well as the social welfare) is $1+\alpha$. In contrast, the social welfare of the assignment in which all buyers are assigned to the vendor is $\sum_{i=1}^{n}\mu_i v_i \geq (1-\alpha)n$; the inequality holds by the definition of v_i and since $\sum_{j=i}^{n}\mu_j \leq \mu_i\sum_{j=0}^{\infty}\alpha^j = \mu_i(1-\alpha)^{-1}$. The price of anarchy is then at least $(1-\alpha)n/(1+\alpha)$ which can become arbitrarily close to n by selecting α appropriately. □

Interestingly, the price of anarchy does not depend on any other quantity and the lower bound of Lemma 1 is tight.

Theorem 2. *The price of anarchy of any price competition game with n buyer types is at most n.*

Proof. Consider an equilibrium $(\mathbf{x},\mathbf{p})$ of a price competition game. We first claim that if buyers of some type i are split between two vendors j and j', then it must be $p_j = c_j$ and $p_{j'} = c_{j'}$ (hence, the two vendors have zero utility) and the assignment in which all these buyers are assigned to vendor j without changing the prices is still an equilibrium and has the same social welfare. This is due to the fact that, at equilibrium, the utilities of buyers assigned to j and j' should be the same. Hence, if one of the two vendors had a price strictly higher than its production cost, it could increase its utility by negligibly decreasing its price; this would result in attracting all buyers of type i previously assigned to j and j'. So, by moving all buyers of type i from vendor j' to vendor j, we still have an assignment that is consistent to $\mathbf{p}$ in which the utilities of buyers and vendors do not change. Clearly, this new assignment is an equilibrium with a social welfare equal to the initial one.

So, without loss of generality, we consider an equilibrium $(\mathbf{x},\mathbf{p})$ such that, for every i, all buyers of type i are assigned to the same vendor j, i.e., $x_{ij} = \mu_i$. We denote by $\eta(i)$ the vendor where the buyers of type i are assigned in $\mathbf{x}$. Also, we denote by $o(i)$ the vendor to which the buyers of type i are assigned in an optimal assignment. We can further assume that when $\eta(i) \neq o(i)$, this implies that $v_{i,o(i)} - c_{o(i)} > v_{i,\eta(i)} - c_{\eta(i)}$. If this is not the case and it is $v_{i,o(i)} - c_{o(i)} = v_{i,\eta(i)} - c_{\eta(i)}$, we can consider the optimal assignment that assigns the buyers of type i to vendor $\eta(i)$.

We will show that

$$t_i(\mathbf{x},\mathbf{p}) + u_{o(i)}(\mathbf{x},\mathbf{p}) \geq \mu_i(v_{i,o(i)} - c_{o(i)}) \tag{1}$$

for every $i \in N$. Then, the following derivation can prove the theorem:

$$\begin{aligned} n\cdot\mathrm{SW}(\mathbf{x}) &\geq \sum_i t_i(\mathbf{x},\mathbf{p}) + n\cdot\sum_j u_j(\mathbf{x},\mathbf{p})\\ &\geq \sum_i \left(t_i(\mathbf{x},\mathbf{p}) + u_{o(i)}(\mathbf{x},\mathbf{p})\right)\\ &\geq \sum_i \mu_i\left(v_{i,o(i)} - c_{o(i)}\right)\\ &= \mathrm{SW}^*.\end{aligned}$$

The first inequality follows by the definition of the social welfare, the second one follows from the fact that the function $o(\cdot)$ can assign at most all n buyer types to the same vendor, and the third one follows from (1).

It remains to prove inequality (1). If $\eta(i) = o(i)$, we use the fact that vendor $o(i)$ attracts (at least) the buyers of type i at equilibrium. Hence,

$$\begin{aligned} t_i(\mathbf{x},\mathbf{p}) + u_{o(i)}(\mathbf{x},\mathbf{p}) &\geq \mu_i(v_{i,o(i)} - p_{o(i)})\\ &\quad + \mu_i(p_{o(i)} - c_{o(i)})\\ &= \mu_i(v_{i,o(i)} - c_{o(i)}).\end{aligned}$$

If $\eta(i) \neq o(i)$, let $q_{o(i)} = v_{i,o(i)} - v_{i,\eta(i)} + p_{\eta(i)}$. By our assumption $v_{i,o(i)} - c_{o(i)} > v_{i,\eta(i)} - c_{\eta(i)}$ above and since $p_{\eta(i)} \geq c_{\eta(i)}$, we have

$$\begin{aligned} q_{o(i)} &= v_{i,o(i)} - v_{i,\eta(i)} + p_{\eta(i)}\\ &> c_{o(i)} - c_{\eta(i)} + p_{\eta(i)}\\ &\geq c_{o(i)}.\end{aligned}$$

This means that vendor $o(i)$ can consider deviating to any price value δ from the non-empty interval $[c_{o(i)}, q_{o(i)})$. Since $v_{i,o(i)} - \delta > v_{i,\eta(i)} - p_{\eta(i)}$, with this deviation, vendor $o(i)$

attracts the buyers of type i from vendor $\eta(i)$. Using the equilibrium condition (and denoting by $\mathbf{x}'$ the resulting assignment when vendor $o(i)$ deviates to price δ), we have that

$$\begin{aligned} u_{o(i)}(\mathbf{x},\mathbf{p}) &\geq u_{o(i)}(\mathbf{x}',(\delta,\mathbf{p}_{-o(i)})) \\ &\geq \mu_i(\delta - c_{o(i)}). \end{aligned}$$

Since the above inequality holds for any $\delta < q_{o(i)}$, it must also be

$$\begin{aligned} u_{o(i)}(\mathbf{x},\mathbf{p}) &\geq \mu_i(q_{o(i)} - c_{o(i)}) \\ &= \mu_i(v_{i,o(i)} - c_{o(i)}) - \mu_i(v_{i,\eta(i)} - p_{\eta(i)}) \\ &= \mu_i(v_{i,o(i)} - c_{o(i)}) - t_i(\mathbf{x},\mathbf{p}). \end{aligned}$$

This completes the proof of the theorem. □

Recall that the upper bound on the price of anarchy is meaningful only for games that admit equilibria. In the following, we show that games with one buyer type always have equilibria (and, by Theorem 2, they also have optimal social welfare) while the existence of a second buyer type may result to instability.

Lemma 3. *Price competition games with one buyer type always have at least one equilibrium.*

Proof. Consider a price competition game with one buyer type. We use the simplified notation v_j to denote the valuation of the buyers for vendor j. Let $j^* \in \arg\max_{j\in M}\{v_j - c_j\}$ and $j' \in \arg\max_{j\in M\setminus j^*}\{v_j - c_j\}$ be two vendors with the highest values for the difference $v_j - c_j$. Set $p_{j^*} = v_{j^*} - v_{j'} + c_{j'}$ and $p_j = c_j$ for any vendor $j \neq j^*$. We claim that this price vector together with the consistent assignment $\mathbf{x}$ that assigns the buyers to vendor j^* is an equilibrium. Indeed, no vendor $j \neq j^*$ has any incentive to change its price; a decrease would result in negative utility while an increase would not change the assignment. Moreover, vendor j^* has no incentive to change its price; a decrease can only lower its utility while an increase would result in a new assignment in which all buyers are attracted by vendor j'. □

Lemma 4. *There exists a price competition game with two buyer types that admits no equilibrium.*

Sketch of proof. The proof of the lemma uses a price competition game with two buyer types of unit volume each and two vendors with a production cost of 0. We use the terms left and right to refer to the vendors and buyer types. The valuation of the left buyers is $v_{\ell\ell} = 5$ for the left vendor and $v_{\ell r} = 3$ for the right vendor; the valuation of the right buyers is $v_{r\ell} = 3$ for the left vendor and $v_{rr} = 5$ for the right vendor. Through a case analysis, we can show that no pair of a price vector and consistent assignment can be an equilibrium. □

We remark that Meir et al. [2014] also present a two-vendor four-buyer-type price competition game that does not admit any equilibrium; the game in the proof of Lemma 4 is the simplest one with this property.

4 Complexity of equilibria

We begin the discussion of this section by formulating some concrete computational problems related to equilibria of price competition games.

> VERIFYEQUILIBRIUM: Given a price vector $\mathbf{p}$ and a buyers-to-vendors assignment $\mathbf{x}$ in a price competition game $\mathcal{G}$, decide whether $(\mathbf{x},\mathbf{p})$ is an equilibrium of $\mathcal{G}$.
>
> COMPUTEPRICE: Given a buyers-to-vendors assignment $\mathbf{x}$ in a price competition game $\mathcal{G}$, decide whether there exists a price vector $\mathbf{p}$ to which $\mathbf{x}$ is consistent so that $(\mathbf{x},\mathbf{p})$ is an equilibrium of $\mathcal{G}$.
>
> PRICECOMPETITION: Decide whether a given price competition game has any equilibrium or not.

VERIFYEQUILIBRIUM can be easily seen to be solvable in time $\mathcal{O}(nm)$. First, one needs to check whether $\mathbf{x}$ is consistent to $\mathbf{p}$, i.e., whether the utility of each buyer type is maximized at the vendor(s) used in $\mathbf{x}$; this can be done by computing at most $\mathcal{O}(nm)$ buyer utilities. Then, for every vendor j and every buyer type i, it suffices to compute the maximum price level which is sufficient so that vendor j attracts buyers of type i and the vendor's utility when deviating to this price level (equal to the volume of buyers it attracts times the difference of the price level from the production cost). The final decision is YES if $\mathbf{x}$ is consistent to $\mathbf{p}$ and the utility of all vendors in $(\mathbf{x},\mathbf{p})$ is equal to the maximum utility over all the deviations considered; otherwise, it is NO. In the following, we call this algorithm `Verify`.

The problem COMPUTEPRICE looks significantly more difficult at first glance since there are too many price vectors (to which $\mathbf{x}$ is consistent) that have to be considered. Interestingly, we will present a polynomial-time algorithm (henceforth called `CandidatePrice`) which, given a price competition game and a buyers-to-vendors assignment $\mathbf{x}$, comes up with a single candidate price vector $\mathbf{p}$ that can in turn easily be checked whether it forms an equilibrium together with $\mathbf{x}$ using `Verify`.

`CandidatePrice` works as follows. It first computes a set Z of *seed* vendors which will have a price equal to their production cost. In order to define Z, it is convenient to consider the directed graph H that has a node for each vendor and a directed edge from node j to node j' labelled by i if buyers of type i are assigned to vendor j in $\mathbf{x}$ and, furthermore, $v_{ij} - c_j \leq v_{ij'} - c_{j'}$. Now, the set Z is defined recursively as follows:

1. Any vendor that is not assigned any buyer in $\mathbf{x}$ belongs to Z; such a vendor is called *empty*.
2. Any vendor j such that $\min_{i:x_{ij}>0} v_{ij} = c_j$ belongs to Z.
3. Any vendor that is part of a directed cycle in H belongs to Z.
4. Any vendor that has a directed edge to a vendor of Z also belongs to Z.

`CandidatePrice` returns the price vector $\mathbf{p}$ with $p_j = c_j$ for each seed vendor j and

$$p_j = \begin{cases} \min\limits_{i:x_{ij}>0} v_{ij}, & \text{if } Z = \emptyset \\ \min\limits_{i:x_{ij}>0} \left\{ v_{ij} - \max\limits_{j' \in Z} \{v_{ij'} - c_{j'}\}^+ \right\}, & \text{otherwise} \end{cases}$$

for each non-seed vendor.

The correctness of the algorithm is given by the following lemma; the proof is omitted.

Lemma 5. *Let* $\mathbf{p}$ *be the price vector returned by* `CandidatePrice` *on input a price competition game* $\mathcal{G}$ *and a buyers-to-vendors assignment* $\mathbf{x}$*. If* $\mathcal{G}$ *admits an equilibrium* $(\mathbf{x}, \mathbf{q})$*, then* $q_j = p_j$ *for every non-empty vendor* j*.*

We now turn our attention to PRICECOMPETITION and first consider the cases where either the number of buyer types or the number of vendors is constant. If the number of buyer types is constant, then there are at most $(m+1)^n$ different buyers-to-vendors assignments and corresponding instances of COMPUTEPRICE that we need to consider.[1] The case of a constant number of vendors (this can be thought of as an *oligopoly*) is considerably more involved but still computable in polynomial time as we show in the following.

For a fixed price vector $\mathbf{p}$, we define the induced preference, denoted by $\succ_i$, of buyers of type i as $j \succ_i j'$ if $v_{ij} - p_j > v_{ij'} - p_{j'}$ and $j \simeq_i j'$ if $v_{ij} - p_j = v_{ij'} - p_{j'}$. We use the term preference profile to refer to a combination of buyer preferences. The main idea is to enumerate the different preference profiles that are defined for all price vectors $\mathbf{p} \in \mathbb{R}^m$. Observe that the sign (from $\{-, 0, +\}$) of the expression $v_{ij} - v_{ij'} - p_j + p_{j'}$ indicates whether buyers of type i prefer vendor j to vendor j' (i.e., $j \succ_i j'$), are indifferent between the two (i.e., $j \simeq j'$), or prefer vendor j' to j (i.e., $j' \succ_i j$). The number of different preferences of buyers between the two specific vendors is given by the number of different sign patterns for the expressions $v_{ij} - v_{ij'} - p_j + p_{j'}$ for $i = 1, ..., n$ as the difference $p_j - p_{j'}$ runs from $-\infty$ to ∞. Since there are at most n different values of the difference $v_{ij} - v_{ij'}$, this number is at most $2n+1$. In total, the number of distinct sign patterns we need to enumerate in order to consider all distinct preference profiles is at most $(2n+1)^{\binom{m+1}{2}}$; this is polynomial in n when m is constant.

When considering a preference profile $\succeq$, we compute the following assignment $\mathbf{x}$ which should be given to `CandidatePrice` in order to return a price vector $\mathbf{p}$; the pair $(\mathbf{x}, \mathbf{p})$ will in turn be given to `Verify` to detect whether it corresponds to an equilibrium or not. For each buyer type i with a unique top preference (i.e., strictly preferring a particular vendor to all others), $\mathbf{x}$ assigns the buyers of type i to their most preferred vendor. For each buyer type i that has a set T of at least two vendors tied as its top preference, $\mathbf{x}$ assigns i to a(ny) vendor j of T maximizing $v_{ij} - c_j$. We call this algorithm `Enumerate`.

Clearly, on input a price competition game that does not admit an equilibrium, `Enumerate` will not find any. The next lemma completes the proof of correctness of `Enumerate`.

Lemma 6. *On input a price competition game,* `Enumerate` *returns an equilibrium if one exists.*

Proof. Assume that `Enumerate` is applied on input a price competition game $\mathcal{G}$ that admits an equilibrium $(\mathbf{x}, \mathbf{q})$. If $\mathbf{x}$ is the unique assignment that is consistent to $\mathbf{q}$, `Enumerate` will consider the preference profile $\succeq$ corresponding to vector $\mathbf{q}$ and will pass the uniquely defined assignment $\mathbf{x}$ as input to `CandidatePrice` to compute a price vector $\mathbf{p}$; by Lemma 5, $(\mathbf{x}, \mathbf{p})$ will form an equilibrium of $\mathcal{G}$.

Now, assume that $\mathbf{x}$ is not the unique assignment that is consistent to $\mathbf{q}$. Denote by $\mathbf{x}'$ the assignment computed by `Enumerate` (notice that $\mathbf{x}'$ is consistent to $\mathbf{q}$ as well) when considering the preference profile that corresponds to the price vector $\mathbf{q}$. We will show that $(\mathbf{x}', \mathbf{q})$ is an equilibrium as well; then, Lemma 5 guarantees that an equilibrium will be found when `Enumerate` will run `CandidatePrice` with input assignment $\mathbf{x}'$.

Consider a buyer type i with $x_{ij} > 0$ and $x'_{ij'} > 0$ for two different vendors j and j'. By our assumptions, we have $v_{ij} - q_j = v_{ij'} - q_{j'}$ (since buyers of type i are indifferent between vendors j and j' in $\mathbf{q}$) and $v_{ij} - c_j \leq v_{ij'} - c_{j'}$ (since `Enumerate` set $x'_{ij'} > 0$). We will show that $q_j = c_j$ and $q_{j'} = c_{j'}$. Indeed, assume that $q_{j'} > c_{j'}$. Then, by negligibly decreasing its price in $(\mathbf{x}, \mathbf{q})$, vendor j' could increase its utility by attracting (in addition to the buyers it gets in $\mathbf{x}$) all the buyers of type i. Hence, $q_{j'} = c_{j'}$ and $v_{ij} - q_j = v_{ij'} - c_{j'}$. By the inequality $v_{ij} - c_j \leq v_{ij'} - c_{j'}$, we obtain that $q_j = c_j$. The lemma follows since the different assignment of buyers in $\mathbf{x}$ and $\mathbf{x}'$ does not affect the utility of the corresponding vendors (which is zero). □

The restrictions on the numbers of vendors or buyer types are necessary in order to come up with efficient algorithms for PRICECOMPETITION (unless P = NP).

Theorem 7. PRICECOMPETITION *is NP-hard.*

5 Enforcing equilibria using subsidies

We now consider the option to use subsidies. A subsidy given to a vendor aims to compensate it for setting its price at a particular value. In this way, subsidies can be used to enforce a particular pair of price vector and consistent buyers-to-vendors assignment. Formally, given a price vector $\mathbf{p}$ and a consistent assignment $\mathbf{x}$, denote by $\theta_j(\mathbf{x}, \mathbf{p})$ the maximum utility of vendor j over all deviations p'_j and all assignments $\mathbf{y}$ that are consistent to $(p'_j, \mathbf{p}_{-j})$. Vendor j has no incentive to follow any such deviation when it is given an amount of subsidies $s_j \geq \theta_j(\mathbf{x}, \mathbf{p}) - u_j(\mathbf{x}, \mathbf{p})$. If this inequality holds for every vendor j, we say that the pair $(\mathbf{x}, \mathbf{p})$ is enforced as an equilibrium. We denote by $\mathbf{s}(\mathbf{x}, \mathbf{p})$ the entry-wise minimum subsidy vector that enforces $(\mathbf{x}, \mathbf{p})$ as an equilibrium, i.e., $s_j(\mathbf{x}, \mathbf{p}) = \theta_j(\mathbf{x}, \mathbf{p}) - u_j(\mathbf{x}, \mathbf{p})$. We use the terms "total amount" and "cost" to refer to the sum of all entries of a subsidy vector.

[1] At first glance, we have to consider all possible ways to split the volume of the buyers among different vendors. A naive implementation could require exponential time but, fortunately, using the same argument as in the first paragraph of the proof of Theorem 2, we can safely conclude that we only need to consider non-fractional assignments.

Our first observation is that a large amount of subsidies may inherently be necessary to enforce *any* equilibrium.

Theorem 8. *For every $\delta > 0$, there exists a price competition game, in which no subsidy assignment of cost smaller than $(1/4 - \delta)SW^*$ can enforce any pair of price vector and consistent buyers-to-vendors assignment as an equilibrium.*

Sketch of proof. As in the proof of Lemma 4, we use a price competition game with two buyer types of unit volume each and two vendors. Let $\epsilon > 0$; then the valuations are $v_{\ell\ell} = v_{rr} = 4 - \epsilon$ and $v_{r\ell} = v_{\ell r} = 3$. Through a case analysis, we can show that no price vector and consistent buyers-to-vendors assignment can be enforced as an equilibrium using an amount of subsidies less than $2 - 2\epsilon$. The optimal social welfare is $8 - 2\epsilon$ and the theorem will follow by setting ϵ to a sufficiently small positive value. □

We now restrict ourselves to optimal assignments and show tight bounds on the cost of subsidies that are necessary and sufficient to enforce these assignments as equilibria.

Theorem 9. *In every price competition game, the optimal assignment can be enforced as an equilibrium using an amount of subsidies that is at most SW*. This bound is tight. In particular, for every $\epsilon > 0$, there exists a price competition game in which the optimal assignment cannot be enforced as an equilibrium with a total amount of subsidies less than $(1-\epsilon)SW^*$.*

Proof. We first prove the upper bound. Consider an optimal assignment $\mathbf{x}$ with $x_{ij} \in \{0, \mu_i\}$ for every buyer type/vendor pair and the price vector $\mathbf{p}$ with $p_j = c_j$; clearly, $\mathbf{x}$ is consistent to $\mathbf{p}$. By deviating to any other price, vendor j cannot get any buyers that are not assigned to it in the optimal assignment. Hence, it suffices to assign a subsidy of $s_j(\mathbf{x}, \mathbf{p}) = \sum_i x_{ij}(v_{ij} - c_j)$ to each vendor j; this obviously yields a total amount of subsidies equal to SW*.

For the lower bound, let $\chi > 2$ and consider the price competition game with two buyer types of unit volume and valuations χ and 1 for a single vendor of production cost of 0. The optimal social welfare is SW* $= \chi + 1$. Observe that the utility of the vendor is maximized to χ by setting its price to χ while any price that is consistent to assigning both buyer types to the vendor is at most 1 for a vendor utility of at most 2. Hence, the amount of subsidies required to enforce the optimal assignment as an equilibrium is at least $\chi - 2$ which becomes at least $(1 - \epsilon)$SW* by setting χ sufficiently large. □

Even though the minimum amount of subsidies that is sufficient to enforce the optimal assignment as an equilibrium can be large in terms of the optimal social welfare, one might hope that it could be efficiently computable. Unfortunately, this is far from true as we show below in Theorem 10. Before presenting the theorem, let us formally define the corresponding optimization problem:

> MINSUBSIDIES: Given a price competition game $\mathcal{G}$ with an optimal assignment $\mathbf{x}$, compute a price vector $\mathbf{p}$ that minimizes the cost $\mathbf{s}(\mathbf{x}, \mathbf{p})$ over all price vectors to which $\mathbf{x}$ is consistent.

Recall that MINSUBSIDIES should return an equilibrium $(\mathbf{x}, \mathbf{p})$ when one exists. This can be efficiently decided using algorithm `CandidatePrice`. The hardness of the problem manifests itself in instances that do not admit equilibria; we exploit such instances in the proof of Theorem 10.

Theorem 10. *Approximating* MINSUBSIDIES *within any constant is NP-hard.*

Sketch of proof. We will use an approximation-preserving reduction from the NODE COVER problem in k-uniform hypergraphs (i.e., hypergraphs in which every edge consists of $k \geq 2$ nodes), which is formally described as follows.

> NODE COVER: Given a k-uniform hypergraph G, find a node subset C of minimum size so that every (hyper)edge e has at least one of its nodes in C.

The quantity k in the definition of NODE COVER is a constant. It is known that, for every constant $\epsilon > 0$, approximating NODE COVER within $k - 1 - \epsilon$ is NP-hard [Dinur *et al.*, 2005]. Given a k-uniform hypergraph G, we construct the following price competition game:

- for every edge e of G, there is an *edge* vendor e and a buyer type b_e with volume 1 and valuation $(k + 1)^2$ for vendor e;
- for every node j of G, there are: one *node* vendor j, one *auxiliary* vendor j^*, one buyer type b_j with volume 1 and valuations $k + 3$ for vendor j and $k + 2$ for every vendor e such that $j \in e$, and a buyer type b_j^* with volume $1/(k + 2)$ and valuations $k + 2$ for vendor j and $k + 3$ for vendor j^*.
- all valuations not mentioned above as well as all production costs are zero.

In the optimal assignment $\hat{\mathbf{x}}$, for every node j of G, buyers of type b_j are assigned to vendor j, and buyers of type b_j^* are assigned to vendor j^* and, for every edge e of G, buyers of type b_e are assigned to vendor e. We can show that the minimum amount of subsidies required to enforce this optimal assignment as an equilibrium is *equal* to the size of a minimum node cover of G. More details are omitted. □

6 Open problems

In this work, we have posed and answered a long list of questions about price competition games. Of course, our work reveals a lot more open problems. We mention a few here. First, observe that we have made no particular attempt to optimize the running time of our algorithms. We believe that there is much room for improvement on the running time of `CandidatePrice` and `Enumerate`. In particular, it would be interesting to come up with FPT algorithms (see Downey and Fellows [1999]) for PRICECOMPETITION with respect to different parameters. Second, in spite of our inapproximability result (Theorem 10), we believe that it is important to design polynomial-time approximation algorithms for MINSUBSIDIES. For example, is there a logarithmic approximation algorithm? What about additive approximations using an amount of subsidies that exceeds the minimum by at most $\rho \cdot$ SW* for some small $\rho > 0$?

Another set of open problems comes from introducing constraints to price competition games such as supply limitations. For example, consider additional input parameters that indicate the maximum volume of buyers each vendor can support. We believe that this subtle difference in the definition makes the setting even richer from the computational point of view. Another question concerns mixed equilibria. Do such equilibria always exist? Observe that the strategy spaces of vendors have infinite size in this case. Can they be computed efficiently? What is their price of anarchy? What about generalizations of our model that include uncertainty for buyer valuations? It is our firm belief that these questions deserve investigation.

Acknowledgements

This work was partially supported by the European Social Fund and Greek national funds through the research funding program Thales on "Algorithmic Game Theory" and by the Caratheodory research grant E.114 from the University of Patras.

References

[Augustine *et al.*, 2012] J. Augustine, I. Caragiannis, A. Fanelli, and C. Kalaitzis. Enforcing efficient equilibria in network design games via subsidies. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 277–286, 2012.

[Babaioff *et al.*, 2013] M. Babaioff, B. Lucier, and N. Nisan. Bertrand networks. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 33–34, 2013.

[Babaioff *et al.*, 2014] M. Babaioff, N. Nisan, and R. Paes Leme. Price competition in online combinatorial markets. In *Proceedings of the 23rd International Conference on World Wide Web (WWW)*, pages 711–722, 2014.

[Bachrach *et al.*, 2009] Y. Bachrach, E. Elkind, R. Meir, D. Pasechnik, M. Zuckerman, J. Rothe, and J. S. Rosenschein. The cost of stability in coalitional games. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*, pages 122–134, 2009.

[Buchbinder *et al.*, 2010] N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative cost sharing games via subsidies. *Theory of Computing Systems*, 47(1):15–37, 2010.

[Chawla and Roughgarden, 2008] S. Chawla and T. Roughgarden. Bertrand competition in networks. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT)*, pages 70–82, 2008.

[Cole *et al.*, 2006] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72(3):444–467, 2006.

[Dinur *et al.*, 2005] I. Dinur, V. Guruswami, S. Khot, and O. Regev. A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM Journal on Computing*, 34(5):1129–1146, 2005.

[Downey and Fellows, 1999] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer-Verlag, 1999.

[Guruswami *et al.*, 2005] V. Guruswami, J. D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173, 2005.

[Koutsoupias and Papadimitriou, 2009] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.

[Lev *et al.*, 2015] O. Lev, J. Oren, C. Boutilier, and J. S. Rosenschein. The pricing war continues: on competitive multi-item pricing. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[Mas-Colell *et al.*, 1995] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[Meir *et al.*, 2014] R. Meir, T. Lu, M. Tennenholtz, and C. Boutilier. On the value of using group discounts under price competition. *Journal of Artificial Intelligence*, 216:163–178, 2014.

[Monderer and Tennenholtz, 2004] D. Monderer and M. Tennenholtz. K-implementation. *Journal of Artificial Intelligence Research*, 21:37–62, 2004.

[Nisan, 2007] N. Nisan. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, chapter 9, pages 209–241. Cambridge University Press, 2007.

[Papadimitriou, 2001] C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.

[Wooldridge *et al.*, 2013] M. Wooldridge, U. Endriss, S. Kraus, and J. Lang. Incentive engineering for boolean games. *Artificial Intelligence*, 195:418–439, 2013.