Sharpness of Approximation Boundary in the Bipartite Pricing Problem

Kyle Woodward

December 16, 2013

Abstract

It is known that the bipartite demand problem has an easy approximation which generates at least half the maximum profits for the selling side of the market. We demonstrate via an example that this approximation is tight, and that the results of the approximation algorithm cannot necessarily be improved upon any subsequent algorithm which takes the approximation as given.

1 Introduction

Consider the pricing problem faced by a monopolist who produces heterogeneous goods at 0 cost, where consumer demand is unitary over pairs of the goods. That is, a consumer might be willing to pay \$5 for an apple and an orange together, but is willing to pay \$0 for any other combination; another consumer is willing to pay \$3 for an apple and a pear together, and \$0 for any other bundle. Demands are assumed to be commonly known. In a market of this kind, the optimal product-pricing scheme is well-defined (if nonunique), however for large numbers of products and agents it may be computationally intractable, hence the ability to approximate optimal pricing is desirable.

Balcan and Blum [1] give a simple algorithm which is guaranteed to obtain at least half of the maximum profits: optimally price only one "half" of the market, and price all other products at 0. Since the priced half of the market is unconnected, there are no computational concerns regarding forward effects on the demand network.

In the given example, we can view the sides of the network as $\{(\text{Apple}), (\text{Orange}, \text{Pear})\}$; the optimal price for the "Apple" side of the market is $p_{\text{Apple}} = 3$ (for a profit of 6), the optimal prices for the other side of the market are $p_{\text{Orange}} = 5$, $p_{\text{Pear}} = 3$ (for a profit of 8). In this setting, the total social gains (8) equal the total profit

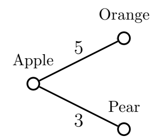

Figure 1: the fruit-demand network

available from optimal pricing, which in turn equals the total profit available from pricing one side of the market.

Khandekar et al. [2] demonstrate that the $o(2)$ pricing bound — the one-sided pricing algorithm guarantees at least half the total surplus — is sharp; we contribute a simple example of the sharpness of this bound. In this same example, we demonstrate that not only is the bound sharp, but it cannot be improved by any algorithm which follows one-sided pricing by optimal re-pricing of the unpriced side of the market so long as prices are constrained to be positive. The case where prices are possibly negative is left undiscussed.

2 Sharpness

Let $k \in \mathbb{N}_{++}$ represent the number of consumers in an economy. Let $V^k = \{v_{\text{root}}\} \cup [\cup_{i=1}^{k}\{v_{1/i}\}]$ be a set of products and $E^k = \{(v_{\text{root}}, v_{1/i}) : 1 \leq i \leq k\}$ be the set of pairwise demands; together, these form the graph $G_j(V, E)$. We associate the set E with the k consumers, where consumer i has demand represented by the edge $(v_{\text{root}}, v_{1/i})$. Let consumer i have value $1/i$ for this bundle.[1] Denote the economy described by these features by $\mathcal{E}^k$.

Given k, choose two such economies $\mathcal{E}_1^k$ and $\mathcal{E}_2^k$; let $\mathcal{E}(k) = \mathcal{E}_1^k \cup \mathcal{E}_2^k$ — where the union of two economies is the union over their vertices and edges, where each is labeled according to the original economy from which it is drawn — together with an additional consumer $i = (0, 0)$ who has value 1 for the bundle $(v_{\text{root},1}, v_{\text{root},2})$.[2] The graph of this economy is $\mathcal{G}(k)$, pictured in Figure 2.

[1] We can restrict this model to integer valuations by multiplying by $k!$; for ease of notation we will retain the convention of fractional valuations.

[2] There is a fairly broad range of values for the "connector" agent $(v_{\text{root},1}, v_{1/k,1})$ such that this example will still hold. In particular, so long as he exists and has value less than 2, our argument requires no modification.

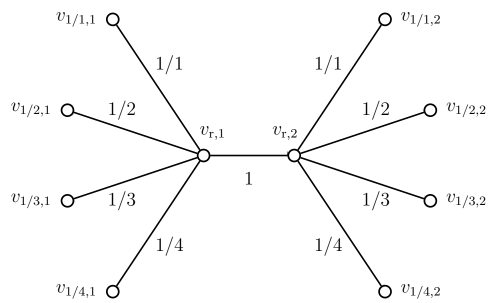

Figure 2: the market network G, with $k = 4$.

2.1 Profit baseline

A perfectly-discriminating monopolist could extract all rents from consumers; its surplus would then be

$$\pi^{\text{disc}} = 1 + 2\sum_{i=1}^{k} 1/i.$$

Since the monopolist is not able to price-discriminate, it must price vertices (goods) individually so as to maximize the profit captured from consumer willingness to pay. We claim that the monopolist maximizes profits by setting prices

$$p_{1/i,j} = 1/i - 1/k, \quad p_{\text{root},j} = 1/k \quad \implies \quad \pi^{\text{mon}} = 2/k + 2\sum_{i=1}^{k} 1/i.$$

Note that raising the price of any "leaf" good will cause the demanding consumer to drop out of the market, reducing profits. Additional rents are available from the connecting consumer, but raising the price of either of his goods entails other consumers dropping out of the market[3], in such a fashion as to reduce total profits: raising the root price by ε eliminates all of that market's leaf consumers, unless accompanied by a similar reduction in price for all of the leaf goods, for nonpositive profits. Since prices are bounded below by 0, this profit differential is occasionally strictly negative; hence the monopolist is profit-maximizing at these prices.

[3]Note that if we allow prices to be negative the monopolist can capture all consumer surplus in this market.

2.2 Approximation

Notice that $\mathcal{G}(k)$ is bipartite: there is a natural partition of vertices (products) such that there is no edge between any two vertices in a particular element of the partition; that is, there is no consumer who demands more than one good in any element of the partition. This partition is demonstrated in Figure 3, and is given by

$$\mathcal{P}(k) = \left\{ \{v_{\mathrm{root},1}\} \cup \left[\cup_{i=1}^{k} \{v_{1/i,2}\} \right], \{v_{\mathrm{root},2}\} \cup \left[\cup_{i=1}^{k} \{v_{1/i,1}\} \right] \right\}.$$

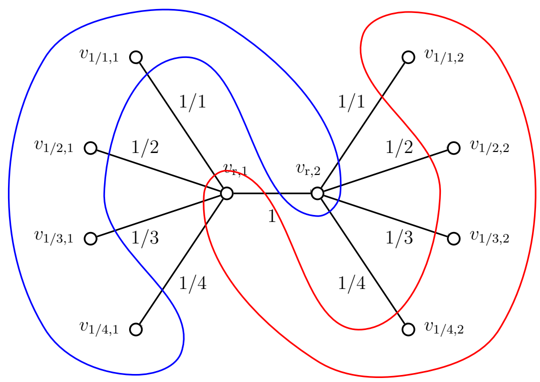

Figure 3: a partition of the market into two internally-disconnected submarkets, $\mathcal{P}(k)_1$ in red and $\mathcal{P}(k)_2$ in blue.

Now consider the profit available from pricing only one side (i.e., one partitional element) of $\mathcal{G}(k)$; by symmetry of $\mathcal{G}(k)$, we may constrain ourselves to pricing the vertices $V' = V^1 \setminus \{v_{\mathrm{root},1}\}$. It is clear that we may extract full surplus from all agents $i \leq k$ with $j = 1$: their vertices in V' have degree 1. Hence profit of $\sum_{i=1}^{k} 1/i$ is obtained from these agents. As for the agents in $\mathcal{E}_2^k$ and the connector agent, all surplus to be extracted from them must come from the price assigned to $v_{\mathrm{root},2}$.

By construction, the profit available from agents in $\mathcal{E}_2$ from pricing $p = v_{\mathrm{root},2}$ is

$$p \cdot \#(\{1/i \geq p\}) = (1/p')p' = 1, \quad p' \in \mathbb{N} \cap [1, k].$$

There is additional surplus of p available from the connecting agent when $p \leq 1$, hence the profit-maximizing price of $v_{\mathrm{root},2}$ is $p = 1$, yielding a total profit of 2 on this partitional element.

The one-sided pricing algorithm uses these prices as the final prices for the entire graph. The total profit obtained in the larger market from the one-sided approximation is therefore

$$\pi^{\text{appx}} = 2 + \sum_{i=1}^{k} 1/i.$$

The ratio of optimal profit to one-sided optimal profit is given by

$$\frac{2 + \sum_{i=1}^{k} 1/i}{2/k + 2\sum_{i=1}^{k} 1/i} = \frac{1}{2} + (2 - 1/k)\left(2/k + \sum_{i=1}^{k} 1/i\right)^{-1}.$$

Since k may be chosen arbitrarily large and the sum $\sum_{i=1}^{k} 1/i$ diverges, the limit of the right-hand term is 0. It follows that the bound $o(2)$ is sharp for the Balcan-Blum pricing algorithm.

3 Conditional unimprovability

The question arises: fixing this approximation, is there a simple algorithm which obtains a non-trivial amount of the remaining surplus? Our example demonstrates that capturing *any* remaining surplus can be impossible.

In particular, following the approximation above there is no additional surplus to be extracted from consumers in $\mathcal{E}_1^k$: their valuation is being fully extracted at one of the goods they desire. Turning to $\mathcal{E}_2^k$, we recall that the price of $v_{\text{root},2}$ is 1; this price is (weakly) above the valuation of any bundle on this side of the market, hence positive prices cannot affect either intensive or extensive market activity. Applying a pricing algorithm to the unpriced side of the market therefore results in zero additional profits. It follows that the $o(2)$ bound is tight even when a positive-repricing scheme is applied.

4 Negative pricing

Allowing negative prices would not only allow for full surplus extraction in the non-approximation case, but it would also allow for full surplus extraction in the repricing case. Although negative prices are never optimal when pricing the first side of the market — there are no knock-on effects to be mitigated with negative pricing, hence such prices are only lossy — they can affect a consumer's willingness to enter if used in the two-stage repricing scheme. In this case, by applying negative prices to the unpriced side of the market such that all formerly-outpriced consumers are just willing to purchase, full surplus may be obtained by the monopolist.

It is not clear how general this “subsidy principle” is; intuitively, it seems clear that full extraction through negative repricing is an edge case, and in our example this is a result of the relative sparseness of demand. More work is necessary to determine the bounds on approximation profits if the monopolist can follow the Balcan-Blum approximation with subsidies for those priced out of the market.

5 Discussion

We have provided a constructive example to prove the sharpness not only of the Balcan-Blum pricing approximation on bipartite demand networks, but also of any algorithm which takes the output of this approximation as fixed and attempts to improve monopoly profits by positive pricing. It is assumed that this construction generalizes to higher dimensions, as does Balcan-Blum's approximation when demand is over bundles with more than two goods.

Although the construction is surprisingly general — the key features are that the sum of agent valuations diverges while a single price can only extract a finite amount of surplus[4] — the statement should be taken with a grain of salt: intuitively, the likelihood that valuations over homogeneous goods take this form seems fairly low. In particular, any assumption that values must be distributed, say, normally will break this example. This does not devalue the result, but it should provide context as to the relevance of a worst-case analysis. Moreover, in a sense the introduction of a connecting agent results in a pathological demand graph; this should be interpreted more as a weakness of the bipartite approach than of our particular example.

References

[1] Maria-Florina Balcan and Avrim Blum. Approximation algorithms and online mechanisms for item pricing. *Theory of Computing*, 3:179–195, 2007.

[2] Konstantin Makarychev Rohit Khandekar, Tracy kimbrel and Maxim Sviridenko. On hardness of pricing items for single-minded bidders. In *Approximation, Randomization, and Combinatorial Optimization*, Lecture Notes in Computer Science.

[4] Additionally, the connecting consumer cannot be too valuable (from the firm's perspective).