Strategic Recommendation: Revenue Optimal Matching for Online Platforms (Student Abstract)

Luca D'Amico-Wong*, Gary Qiurui Ma*, David C. Parkes

Computer Science Department, Harvard University, 29 Oxford Street, Cambridge, USA
ldamicowong@college.harvard.edu (806-543-3462), qiurui_ma@g.harvard.edu, parkes@eecs.harvard.edu

Abstract

We consider a platform in a *two-sided market* with *unit-supply* sellers and *unit-demand* buyers. Each buyer can transact with a subset of sellers it knows off platform and another seller that the platform recommends. Given the choice of sellers, transactions and prices form a *competitive equilibrium*. The platform selects one seller for each buyer, and takes a fixed percentage of the prices of all transactions that it recommends. The platform seeks to maximize total revenue.

We show that the platform's problem is NP*-hard*, even when each buyer knows at most two buyers off platform. Finally, when each buyer values all sellers equally and knows only one buyer off platform, we provide a polynomial time algorithm that optimally solves the problem.

Introduction

In 2019, Amazon was sued in Europe for favoring some sellers over others at the expense of consumers. It was claimed to have used the "Buy Box", a key feature at the top right of the product page, to draw buyers to Amazon's own products or third party sellers who pay hefty delivery and storage fees to Amazon, obscuring better deals elsewhere (Veljanovski 2022).

For a platform, it is hard to decide how to recommend sellers to buyers. Recommending a high price product to a buyer risks losing the trade to a competitor off platform. Recommending a low price product forgoes a possible higher commission fee. In this paper, we model this for the platform, and characterize the computational complexity when the platform solves for a revenue optimal strategy.

We formulate the problem in a two-sided market modeled by a *bipartite graph*. Buyers and sellers are vertices on either side of the graph, and edges indicate transaction opportunities between buyers and sellers off platform. The platform adds at most one more edge for each buyer and seller. The market clears according to a competitive equilibrium, subject to transaction constraints represented by the edges. The platform's revenue is proportional to the total price of transactions through the edges it adds. We show that the platform's problem of selecting which set of edges to add is computationally hard.

*Equal contribution.

Related Work

Most related to our work is the study of competitive equilibrium prices on network-formation games (Even-Dar, Kearns, and Suri 2007). Kranton and Minehart (2000) and Elliott (2015) leverage the *network decomposition theorem* to relate prices in a network to the *opportunity path* of the trading agent. We make use of their structural result in the analysis of this paper. These works, however, typically do not assume an intermediary or platform that facilitates trade to gain revenue.

More recent works in the computer science literature model the platform explicitly. Condorelli, Galeotti, and Renou (2017) treat the platform as a liquidity provider that buys and sells as a part of the trading network. The closest to ours is Eden, Ma, and Parkes (2023). They model sellers' and buyers' incentives to join the platform and analyze the social welfare when the platform chooses the commission fee strategically. We instead analyze the platform's matching problem.

From another perspective, recommender systems give personalized suggestions to each user independently, while maximizing overall welfare (Mladenov et al. 2020). We help the recommender maximize revenue, while accounting for the effect that a recommendation has on other buyers within the competitive equilibrium.

Our Model

We adopt the buyer-seller network model used by Kranton and Minehart (2000). Formally, the two-sided market is defined by a bipartite graph $G = (B, S, E)$, where $B = \{b_1, \dots, b_n\}$ represents the set of n unit-demand buyers, $S = \{s_1, \dots, s_m\}$ the set of m unit-supply sellers, and E the possible transaction opportunities. That is, a buyer b_i and seller s_j can only transact if $(b_i, s_j) \in E$. Each buyer b_i values seller s_j's item at value v_{ij}, and each edge $e_{ij} \in E$ has corresponding weight v_{ij}.

The market clears according to a competitive equilibrium, subject to transaction constraints. A competitive equilibrium is defined by a set of transactions that correspond to a maximum weight matching on G, along with a set of item prices p_j that supports the equilibrium. Further assume sellers are able to extract the maximum amount of surplus from the market, charging the highest prices that still yield a competitive equilibrium. Gul and Stacchetti (1999) showed that

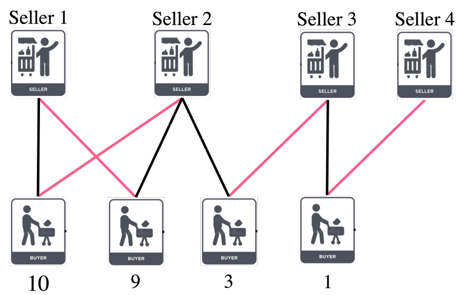

Figure 1: World edges are shown in black, and the revenue-optimal platform edges are in pink. Buyers have homogeneous valuations (i.e. buyer 1 has value 10 for all items). In the revenue-optimal platform strategy, all transactions occur through platform edges, with item prices given by $p_1 = p_2 = 9, p_3 = 3, p_4 = 1$.

these prices are defined by $p_j = W(B, S, E) - W(B, S \setminus \{s_j\}, E)$, where $W(G)$ denotes the weight of the maximum weight matching on G.

To model the platform's role in the market, we consider an existing set of *world edges* E_w, representing transaction opportunities available to buyers and sellers off platform. The world bipartite graph is thus given by $G_w = (B, S, E_w)$. Seeing G_w and possessing knowledge of all v_{ij}, the platform chooses to add a set of *platform edges* E_p (with $E_p \cap E_w = \emptyset$) between buyers and sellers, recommending further transactions that can occur on-platform. [1] The platform problem is to find the set of platform edges, along with a maximum weight matching, that maximizes the sum of item prices sold via platform edges.

Our Contributions

We begin with the following theorem, which shows that the platform's problem is NP-hard, even when buyers know at most two sellers off platform (i.e., $\deg(b_i) \leq 2$ in G_w).

Theorem 1. *The decision version of the platform problem is NP-hard, even when $v_{ij} \in \{0, v_i\}$ and each buyer has at most two existing world edges.*

Proof Sketch. The proof modifies that of Chen and Deng (2014) for revenue-maximizing envy-free pricing and reduces from a version of 3-SAT. Unlike a competitive equilibrium, envy-free pricing does not require unsold items to have price 0; to address this, we introduce dummy buyers D_i whose values are maximal among all buyers, which effectively acts to mimic the optimal envy-free pricing mechanism. □

Theorem 1 requires that some buyers know at least two sellers via world edges; one might then ask if the platform problem is still hard when we assume that buyers know at most a single buyer off-platform. If one additionally restricts buyers to have homogeneous valuations ($v_{ij} = v_i$ for all j), we show that the platform problem becomes tractable.

Denote a seller subgraph S_j as the set of edges and buyers that seller s_j connects to – note that the S_j are disjoint in this setting. Sort and rank all seller subgraphs by the highest-value buyer in the subgraph $S_1, \ldots S_m$. Figure 1 provides such an example market. With any set of platform edges that transact, seller subgraphs are connected into *cycles* and *chains*. We show an example of the structure with Figure 1 here.

$$\underbrace{S_4 \to S_3}_{\text{Chain}} \to \underbrace{S_2 \to S_1 \to S_2}_{\text{Cycle}}$$

Lemma 1. *When buyers have homogeneous valuations and each buyer has at most one existing world edge, there exists a platform-optimal strategy, where the resulting transactions connect seller subgraphs into cycles and at most one chain. All cycles and chains connect contiguous (in sorted order) subgraphs. Cycles are of length at most three. The chain connects to a cycle.*

Note that the revenue-optimal strategy indeed satisfies this lemma for Figure 1. We now prove the following theorem.

Theorem 2. *When buyers have homogeneous valuations and each buyer has at most one existing world edge, the platform problem can be solved in $O(n^2)$ time.*

Proof. Fix an arbitrary contiguous cycle (of length at most 3) for the chain to attach to. Consider the graph G' obtained by removing all the subgraphs in the cycle from G_w.

Let $DP[i]$ denote the maximum obtainable revenue from connecting the first i subgraphs in G' via cycles, obtainable in linear-time via *Dynamic Programming*. Let $R[i]$ be the revenue from linking all subgraphs after i in a chain. It follows that the optimal revenue, given that a chain must attach to the fixed cycle, is given by $\max_i DP[i] + R[i]$.

The maximum revenue can be obtained by repeating this process for all $O(n)$ possible fixed cycles. Going through each fixed cycle takes linear time, so the whole algorithm takes time $O(n^2)$. Note that this exhaustively searches all choices of the fixed cycle and thus finds the optimal configuration satisfying the properties of Lemma 1. By Lemma 1, there exists an optimal solution that satisfies these properties and thus we find it, concluding the proof. □

Conclusion and Future Work

In this work, we model the possible transaction relationships between buyers and sellers, and we analyze how a platform strategically matches buyers to sellers to maximize its revenue. We develop an efficient algorithm for homogeneous good market where each buyer only knows one buyer in the world, and provide hardness results for the general case.

There are a number of promising directions for future work. While the general problem is NP-hard, we would like to design efficient approximation algorithms. Additionally, we assume that the platform possesses complete knowledge of buyer valuations; what can the platform achieve in the case of partial information? Finally, one might want to analyze the welfare properties of the platform-optimal strategy.

[1] WLOG, all platform edges transact. If not, one can show that the platform can drop the non-transacting platform edges and weakly increase revenue.

References

Chen, N.; and Deng, X. 2014. Envy-free pricing in multi-item markets. *ACM Transactions on Algorithms (TALG)*, 10(2): 1–15.

Condorelli, D.; Galeotti, A.; and Renou, L. 2017. Bilateral trading in networks. *The Review of Economic Studies*, 84(1): 82–105.

Eden, A.; Ma, G. Q.; and Parkes, D. C. 2023. Platform Equilibrium: Analyzing Social Welfare in Online Market Places. Forthcoming.

Elliott, M. 2015. Inefficiencies in networked markets. *American Economic Journal: Microeconomics*, 7(4): 43–82.

Even-Dar, E.; Kearns, M. J.; and Suri, S. 2007. A network formation game for bipartite exchange economies. In *SODA*, 697–706.

Gul, F.; and Stacchetti, E. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic theory*, 87(1): 95–124.

Kranton, R. E.; and Minehart, D. F. 2000. Competition for goods in buyer-seller networks. *Review of Economic Design*, 5: 301–331.

Mladenov, M.; Creager, E.; Ben-Porat, O.; Swersky, K.; Zemel, R.; and Boutilier, C. 2020. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *International Conference on Machine Learning*, 6987–6998. PMLR.

Veljanovski, C. 2022. Algorithmic antitrust: a critical overview. *Algorithmic Antitrust*, 39–64.