Strategic Recommendation: Revenue Optimal Matching for Online Platforms

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.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

*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); Kotowski and Leister (2019) 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

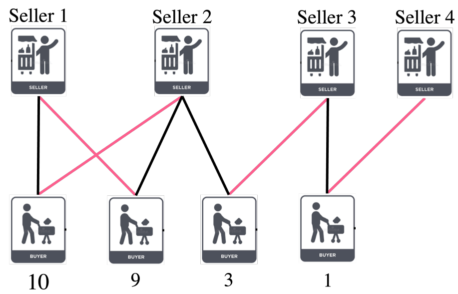

Figure 1: World edges are shown in black, and the revenue-optimal platform edges are in pink. Buyers have homogeneous valuations (i.e. buyer 1 has value 10 for all items). In the revenue-optimal platform strategy, all transactions occur through platform edges, with item prices given by $p_1 = p_2 = 9,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

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.

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

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. Unpublished.

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.

Kotowski, M. H.; and Leister, C. M. 2019. Trading networks and equilibrium intermediation.

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.

Supplementary Material for "Strategic Recommendation: Revenue Optimal Matching for Online Platforms"

1 Competitive Equilibria

As a preliminary, this section formally defines competitive equilibria on a buyer-seller network model and is taken from the authors' unpublished work (Eden, Ma, and Parkes 2023).

We adopt the buyer-seller network model of (Kranton and Minehart 2000). There is a set of n buyers, $B = \{b_1, \dots, b_n\}$, and m differentiated sellers, $S = \{s_1, \dots, s_m\}$. Each seller has a single product to sell, and each buyer b_i has a *unit-demand valuation*. Buyer b_i's value for seller s_j's product is $v_{ij} \geq 0$. A bipartite graph $G = \{g_{ij}\}_{b_i \in B s_j \in S}$ models which buyers and sellers can directly transact, where

$$g_{ij} = \begin{cases} 1 & \text{Buyer } i \text{ can directly transact with seller } j \\ 0 & \text{Otherwise} \end{cases}. \quad (1)$$

We define a competitive (Walrasian) equilibrium in the buyer-seller network model.

Definition 1.1 (Competitive Equilibrium)**.** A *competitive equilibrium* for the buyer-seller network G is a tuple $(\mathbf{p}, \mathbf{a})$ where $\mathbf{p} = (p_1, \dots, p_m)$ are non-negative item prices, $\mathbf{a} = \{a_{ij}\}_{i \in B j \in S} \in \{0, 1\}^{n \times m}$ is an allocation of the goods to the buyers, and:

- Transactions must respect links: $a_{ij} \leq g_{ij} \quad \forall i \in B\ j \in S$.
- Buyers are allocated at most one good: $\sum_j a_{ij} \leq 1 \quad \forall i \in B$.
- Goods are sold at most once: $\sum_i a_{ij} \leq 1 \quad \forall j \in S$.
- Buyers gets their most preferred outcome: $u_i(\mathbf{p}, \mathbf{a}) \geq v_{ij} - p_j \quad \forall i \in B\ j \in S$. where

$$u_i(\mathbf{p}, \mathbf{a}) = \sum_j a_{ij}(v_{ij} - p_j), \quad (2)$$

  is buyer i's utility for the allocation.
- Buyers have non-negative utility: $u_i(\mathbf{p}, \mathbf{a}) \geq 0 \quad \forall i \in B$.
- Unassigned goods have price 0.

A competitive equilibrium is a canonical model of the steady state in a market, capturing the notion of prices that are set such that supply meets demand. It follows from standard existence results (Kelso and Crawford 1982) that a competitive equilibrium always exists in a unit-demand buyer-seller network (a missing edge can be represented as $v_{ij} = 0$). Moreover, competitive equilibria have a number of desirable properties.

Theorem 1.1 (First Welfare Theorem)**.** *In a competitive equilibrium, the social welfare is maximized with respect to the set of allocations that respect the transaction constraints posed by G.*

Theorem 1.2 (Second Welfare Theorem (Gul and Stacchetti 1999))**.** *Let $(\mathbf{p}, \mathbf{a})$ and $(\mathbf{p}', \mathbf{a}')$ be two competitive equilibria of a buyer-seller network G, then $(\mathbf{p}, \mathbf{a}')$ is also a competitive equilibrium (and so is $(\mathbf{p}', \mathbf{a})$).*

The Second Welfare Theorem implies that prices have the property of either forming a competitive equilibrium with any social-welfare optimal allocation, or forming a competitive equilibrium with none of them. We refer to prices $\mathbf{p}$ that are part of a competitive equilibrium as *competitive prices*. It is also well known that competitive prices have a lattice structure.

Theorem 1.3 (Lattice structure for competitive prices (Gul and Stacchetti 1999))**.** *Let $\mathbf{p}_1$ and $\mathbf{p}_2$ be competitive prices, then $\mathbf{p}_1 \vee \mathbf{p}_2$ and $\mathbf{p}_1 \wedge \mathbf{p}_2$ are also competitive prices, where $\vee$ is the coordinate-wise maximum and $\wedge$ is the coordinate-wise minimum.*

As a result, there are *minimum* and *maximum* competitive prices, denoted $\underline{\mathbf{p}}$ and $\overline{\mathbf{p}}$ respectively, with item-wise minimum and item-wide maximum prices denoted $\underline{p}_j$ and $\overline{p}_j$. For S sellers, B buyers, network G, and buyer values $\mathbf{v}$, we use $W(B, S, \mathbf{v}, G)$ to denote the *optimal welfare* from all feasible transactions between sellers S and buyers B. When clear from context, we omit $\mathbf{v}$ from the notation. We will make use of the following characterization result.

Theorem 1.4 (Characterization of competitive prices (Gul and Stacchetti 1999))**.** *The maximum competitive price for an item j has the following form:*

$$\underline{p}_j = W(B, S \cup \{s_j\}, G) - W(B, S, G) \quad (3)$$

$$\overline{p}_j = W(B, S, G) - W(B, S \setminus \{s_j\}, G). \quad (4)$$

Here, $S \cup \{s_j\}$ denotes adding another copy of seller s_j with all its edges to the market, and $S \setminus \{s_j\}$ is removing seller s_j and all its connected edges from the market. The resulting

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

graph changes correspondingly when adding or removing s_j with all its edges, but for notational convenience we still use G to denote the graph.

2 Characterizations of the Optimal Platform Strategy

Recall from the extended abstract the optimal platform strategy is the set of platform edges that maximizes the sum of item prices sold via platform edges in the resulting competitive equilibrium. Here, we provide two basic characterizations of the optimal platform strategy, which will be useful in the proof of Lemmas 4.7 and 4.6. Clearly, when $E_w = \emptyset$, revenue is entirely aligned with welfare, so the platform computes the maximal welfare matching, which can be done in polynomial time. Thus, we restrict our attention to cases where E_w is nonempty.

Lemma 2.1. *Platform's revenue-optimal strategy adds at most one edge to each buyer and seller.*

Proof. Let E_t denote the set of edges where transactions take place in competitive equilibrium. Assume the platform adds two or more edges to a buyer or seller in E_p, by unit-supply unit-demand there exists edge $e_{ij} = (b_i, s_j) \in E_p$ such that $e_{ij} \notin E_t$. Removing e_{ij} does not change seller s_j's price, because

$$\begin{aligned} W(B,S,G \setminus \{e_{ij}\}) &= W(B,S,G) \\ W(B,S \setminus \{s_j\},G \setminus \{e_{ij}\}) &= W(B,S \setminus \{s_j\},G) \end{aligned}$$

Removing e_{ij} weakly increases other sellers s'_j price, because

$$W(B,S \setminus \{s'_j\},G \setminus \{e_{ij}\}) \leq W(B,S \setminus \{s'_j\},G)$$

Thus the platform weakly prefers to drop e_{ij} that are not transacting. □

Lemma 2.2. *For platform's revenue-optimal strategies, either all sellers sell or all buyers buy. If $n = m$, all buyers and sellers transact.*

Proof. We consider the case when there are weakly more buyers than sellers, though the opposite case follows similarly. Take any equilibrium where a seller s_j does not transact. As there are at least as many buyers as sellers, there must be a buyer b_i who also does not transact (and thus does not have an edge connecting them to s_j by the First Welfare Theorem). Consider connecting the two via a platform edge $e_{ij} = (b_i, s_j) \in E_p$.

Clearly, we obtain an increase in revenue from this new transaction. Now, consider any other seller s'_j. Let G be the graph before adding edge e_{ij}, the change in s'_j's price is expressed as $p_{s'_j}(G) = W(B,S,G) - W(B,S \setminus \{s'_j\},G)$, $p_{s'_j}(G \cup \{e_{ij}\}) = W(B,S,G) + v_{ij} - W(B,S \setminus \{s'_j\},G \cup \{e_{ij}\})$. We will show that $p_{s'_j}(G \cup \{e_{ij}\}) \geq p_{s'_j}(G)$. It suffices to prove

$$v_{ij} \geq W(B,S \setminus \{s'_j\},G \cup \{e_{ij}\}) - W(B,S \setminus \{s'_j\},G)$$

If max weight matching in $(B, S \setminus \{s'_j\}, G \cup \{e_{ij}\})$ does not use the new edge e_{ij}, then the RHS is equal to zero and we are done.

Suppose e_{ij} is in the max weight matching, then

$$W(B,S\setminus\{s'_j\},G \cup \{e_{ij}\}) = W(B \setminus \{b_i\},S \setminus \{s_j,s'_j\},G) + v_{ij} \leq W(B,S \setminus \{s'_j\},G) + v_{ij}$$

This is precisely the inequality we want to prove. Thus all other sellers s'_j's price weakly increases. So by matching s_j to b_j, the platform's revenue strictly increases. □

The analysis for Lemma 2.2 and 2.1 does not require homogeneous good assumption, and can cater for general valuation markets.

3 Proof of NP-Hardness

In this section, we prove the following theorem, establishing hardness for the decision version of the platform problem.

Theorem 3.1. *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.*

Similar to that in (Chen and Deng 2014) for revenue-maximizing envy-free pricing, Theorem 3.1 is a result of the following reduction and Lemma 3.1.

Consider a modified version of 3SAT, where each variable x_i appears positively (x_i) and negatively ($\bar{x}_i$) an equal number of times. Given a 3CNF, we add clauses $d_i = (x_i \vee \bar{x}_i)$ for each variable x_i. Additionally, we pad the 3CNF with these clauses such that each variable x_i appears exactly $2t$ times for some $t > 0$.

Denote our modified 3CNF by $(c_1 \wedge \cdots \wedge c_k)$. Let q be the number of unique variables, and let $N = q \cdot k \cdot t$.

Given this 3CNF, we construct an instance of our problem as follows. There are $k+2(t-1)\cdot q+(2t\cdot q-k) = (4t-2)\cdot q$ buyers in total and an equal number of items. We describe the construction below:

- For each variable x_i, we construct $2t$ items $\alpha_{i,j}, \beta_{i,j}$, $j \in \{0, \ldots, t-1\}$. The $\alpha_{i,j}$ are meant to represent the positive instances of x_i while the $\beta_{i,j}$ represent the negative instances. WLOG, let $\alpha_{i,0}, \beta_{i,0}$ be the items that correspond to clauses $d_i = (x_i \vee \bar{x}_i)$.
- For each clause c_i in our 3CNF, we add a buyer U_i that has value N for each item that represents the corresponding literals in clause c_i.
- For each variable x_i, we add $2(t-1)$ buyers $A_{i,j}, B_{i,j}$, $j \in \{1, \ldots, t-1\}$. We also add $2(t-1)$ items $\gamma_{i,j}, \delta_{i,j}$, $j \in \{1, \ldots, t-1\}$. Each buyer $A_{i,j}$ has value $N+1$ for $\alpha_{i,0}, \alpha_{i,j}, \gamma_{i,j}$, and each buyer $B_{i,j}$ has value $N+1$ for $\beta_{i,0}, \beta_{i,j}, \delta_{i,j}$. For each of these buyers, we add world edges to the $\alpha_{i,j}, \beta_{i,j}$ they have positive value for.
- Finally, we add $2t \cdot q - k$ dummy buyers who have value $M \geq N+1$ for all $\alpha_{i,j}, \beta_{i,j}$.

All buyers have value 0 for all items not mentioned above. Finally, we conclude with the proof that this is indeed a valid reduction.

Lemma 3.1. *There is a valid assignment to the original CNF if and only if the optimal revenue is at least* $D := kN + q(t-1)(2N+1) + M \cdot (2t \cdot q - k)$.

Proof. Suppose we have a valid assignment to our original CNF. Then there exists a matching between our buyers U_i and items $\alpha_{i,j}, \beta_{i,j}$ such that no buyer U_i is matched to an item $\alpha_{i,j}$ where x_i is false or an item $\beta_{i,j}$ when x_i is true. We draw these as platform edges. For each buyer $A_{i,j}, B_{i,j}$, we add platform edges to items $\gamma_{i,j}$ and $\delta_{i,j}$ respectively. Finally, we add platform edges between each dummy item and a single unsold item so that all items are sold on platform.

The maximum weight matching then assigns buyers U_i their corresponding items via the platform edges. Each $A_{i,j}$ gets item $\gamma_{i,j}$ and each $B_{i,j}$ gets item $\delta_{i,j}$. Finally, the dummy buyers are matched to the remaining $(2t \cdot q - k)$ items that have not already been sold. As all items are sold and the minimum valuation for any buyer for any item they receive is N, it follows that the price paid for any item is at least N.

We get revenue $N \cdot k$ from the U_i. If x_i is true, $\alpha_{i,0}$ is sold to a buyer U_i, so $A_{i,j}$ has price N as it has an opportunity path to this buyer. However, for $B_{i,j}$, both $\beta_{i,0}$ and $\beta_{i,j}$ are sold to dummy buyers – thus, removing $\delta_{i,j}$ would leave $B_{i,j}$ with no item, reducing welfare by $N+1$. It follows that $B_{i,j}$ must pay $N+1$. Thus, $\forall j, A_{i,j}$ has price N and $B_{i,j}$ has price $N+1$. The converse holds when x_i is False. In conclusion, from each $A_{i,j}$ and $B_{i,j}$, we get revenue $(t-1) \cdot N + (t-1) \cdot (N+1) = (t-1)(2N+1)$. Finally, we get revenue $M \cdot (2t \cdot q - k)$ from the dummy buyers as they are only connected to a single item. Thus, in total we get revenue $kN + q(t-1)(2N+1) + M \cdot (2t \cdot q - k) = D$ as desired.

Now, suppose that there exists a set of platform edges and a matching that generates revenue at least D. Then it must be the case that each buyer U_i gets an item as otherwise we could have maximal revenue $(k-1)N + q \cdot 2(t-1)(N+1) + M \cdot (2t \cdot q - k) < D$. For each U_i corresponding to clause $d_i = (x_i \vee \bar{x}_i)$, they must receive either $\alpha_{i,0}$ or $\beta_{i,0}$. We construct our satisfying assignment as follows:

- If U_i receives $\alpha_{i,0}$, set x_i to True.
- If U_i receives $\beta_{i,0}$, set x_i to False.

It suffices to show that this is a satisfying assignment. To do this, we show that no buyer U_k receives an item $\alpha_{i,j}$ when x_i is False or an item $\beta_{i,j}$ when x_i is True.

If x_i is True, U_i is allocated $\alpha_{i,0}$. Thus, removing the seller that $A_{i,j}$ transacts with (if one exists) can yield at most a reduction of N in total welfare, as we could always connect $A_{i,j}$ with $\alpha_{i,0}$ and drop U_i from the matching. Similarly, if x_i is False, we can get at most revenue N from $B_{i,j}$.

Thus, the platform makes revenue at most N from each buyer $A_{i,\cdot}$ or $B_{i,\cdot}$ that corresponds to the true literal. For the false literal, note that the platform has to get maximal revenue $N+1$ from the corresponding buyer because $D - (k-q)N - M \cdot (2t \cdot q - k) - q(t-1)N = q(t-1)(N+1)$. If U_i receives an item corresponding to the false literal, then the buyer $A_{i,\cdot}$ or $B_{i,\cdot}$ that corresponds to the false literal will only pay price N, contradicting the above.

It follows that our assignment is indeed a satisfying assignment as desired, concluding the proof. ☐

4 Constant Valuation AMOS Setting

Here, we examine the constant valuation AMOS (at most one seller) case, where buyers have *homogeneous valuations* and have at most one incident world edge. We give a quadratic time algorithm to solve the platform's revenue-maximization problem.

Opportunity Paths

We first introduce the concept of opportunity paths, which will be crucial in our analysis of the AMOS setting. This is taken from the authors' unpublished work (Eden, Ma, and Parkes 2023).

For homogeneous good markets, competitive prices are directly related to a buyer's next best forgone trade opportunity, or more specifically its direct and indirect competitors. This notion is characterized by *opportunity paths* in Kranton and Minehart (2000) as follows.

Definition 4.1 (Opportunity Path (Kranton and Minehart 2000))**.** For an allocation $\mathbf{a}$ of goods on a buyer-seller network G, buyer i_1's opportunity path linking to another buyer i_t is a path

$$(i_1, j_1, i_2, j_2, \ldots, j_{t-1}, i_t),$$

where for every $\ell \in \{1, \ldots, t-1\}$,

$$g_{i_\ell, j_\ell} = 1 \text{ and } g_{i_{\ell+1}, j_\ell} = 1,$$

and

$$a_{i_\ell, j_\ell} = 0 \text{ and } a_{i_{\ell+1}, j_\ell} = 1.$$

Kranton and Minehart (2000) show the following connection between buyers' opportunity paths and sellers' maximum competitive prices, which we make use of.

Theorem 4.1 (Opportunity Path Theorem (Kranton and Minehart 2000))**.** *Consider a competitive equilibrium with maximum competitive prices* $(\bar{\mathbf{p}}, \mathbf{a})$ *where* $a_{ij} = 1$. *Seller* j*'s price* $\bar{p}_j$ *is equal to the lowest valuation of any buyer linked by an opportunity path from buyer* b_i. $\bar{p}_j = 0$ *if and only if buyer* b_i *has an opportunity path linking to a seller who does not sell.*

Polynomial-Time Algorithm

Here, we provide a proof of the polynomial time algorithm to optimally solve for the homogeneous buyer AMOS case. We first assume that $|B| = |S|$ and will later relax this assumption. By Lemma 2.2, all sellers and buyers transact. To guarantee maximum revenue, the platform can take the buyer perspective and make sure every buyer pays the highest possible fees. As an intuition, this can be done by raising the lowest buyer's valuation on each buyer's opportunity path. To facilitate easier analysis, we identify the following structure in the bipartite graph.

Definition 4.2 (Seller Subgraphs)**.** Denote a *Seller Subgraph* S_j as a collection of nodes and edges $S_j := \{s_j\} \cup \{b_i | (b_i, s_j) \in E_w\} \cup \{(b_i, s_j) \in E_w\}$ to seller s_j who is known by at least one buyer off platform. Let $v(S_j) := \max_{(b_i, s_j) \in E_w} v_i$ be the value of the largest buyer connected to s_j.

A seller subgraph is a fundamental building block for our analysis. The world graph $G_w = (B, S, E_w)$ is composed of seller subgraphs, *dangling sellers* (sellers with no incident world edges), and *dangling buyers* (buyeres with no incident world edges). Buyers in the same seller subgraph S_j who do not directly buy from s_j have a shared opportunity path by Theorem 4.1. We further sort and index all seller subgraph S_j by their value: $v(S_1) \geq v(S_2) \geq ... \geq v(S_{m'})$ where $m' \leq m$.

Definition 4.3 (Chains and Cycles)**.** Consider a set of transactions on a graph $G_p = (B, S, E_w \cup E_p)$ such that all buyers and sellers transact. E_p further connects seller subgraphs, forming chains and cycles. If there are no dangling sellers, then it is clear that all transactions must be formed of cycles, as defined below.

Otherwise, starting from a seller s_k with no world edges, consider the path defined as follows:

- For each seller, the next vertex in the path is the buyer that it transacts with. If this seller has already appeared in the path, terminate the path.
- For each buyer, the next vertex in the path is the seller whose subgraph it belongs to. If this is a dangling buyer, terminate the path.

As there are a finite number of buyers and sellers, this path must terminate. Additionally, since all buyers and sellers transact, this path must terminate either at a dangling buyer or at a seller who has already appeared previously in the path. Given such a path, we define *0-chains, chains, and cycles* as shown below:

$$\underbrace{s_k \rightarrow b_j}_{\text{0-Chain}} \rightarrow \underbrace{s_i \ldots}_{\text{Chain}} \rightarrow \underbrace{s_j \rightarrow \cdots \rightarrow s_j}_{\text{Cycle}}$$

Note that this path always begins with a dangling seller transacting with a buyer – we term this part of the path a 0-chain. The chain (if one exists) connects the next seller, all the way up to, but not including, the seller who begins the cycle. If the path terminates at a dangling buyer b, we consider this to be a chain connected to a 0-cycle, where the 0-cycle is simply the dangling buyer b. Additionally, we define the length of a cycle to be the number of unique sellers that it includes.

To clean up the proof, from here forward, we will suppress any mention of 0-chains and deal only with chains and cycles, leaving implicit the fact that any extra buyers who do not belong to a chain or cycle transact with a dangling seller. Note that the set of chains and cycles uniquely defines the set of transactions (up to deciding which dangling sellers a buyer transacts with, which does not affect platform revenue).

We start with the following lemma, which places a constraint of the length of any cycles in the optimal solution.

Lemma 4.1. *There exists an optimal set of transactions that can be decomposed into chains and/or cycles. Furthermore, there exists an optimal solution where all such cycles are of length at most* 3*.*

Proof. By Lemma 2.2, there exists a platform optimal strategy where all buyers and sellers transact. By Definition 4.3, this implies that the set of transactions can be decomposed into cycles and chains.

Let us now bound the length of these cycles. Suppose that we had a cycle of length at least 4. Then we could always split this cycle into smaller cycles of length at most 3 (note that we can represent every integer larger than 3 as a sum of multiples of 2 and 3). Furthermore, these cycles can only result in weakly higher revenue among sellers included in the original cycle as we are decreasing the number of opportunity paths. Finally, any chain connected to the original cycle can now be attached to the same subgraph in this new split set of cycles – again, there are fewer opportunity paths, so each seller in the chain also obtains weakly higher revenue.

It follows that there always exists an optimal solution where all cycles have length at most 3. □

We make the following simple observation about the optimal configuration of such cycles and chains.

Observation 4.1. The cycle on k subgraphs that yields optimal revenue connects the highest buyer in each subgraph together. The optimal-revenue chain connects subgraphs from smallest to largest in terms of their max-value bidder. The chain and cycle that yields optimal revenue is constructed by attaching the optimal chain to the largest second-highest bidder in any subgraph in the optimal cycle.

In fact, we can also show that there exists an optimal solution with only a single chain.

Lemma 4.2. *There exists an optimal solution with only a single chain. Additionally, this optimal solution can still restrict the cycle length to at most* 3*.*

Proof. Suppose we had two chains, each attached to a different cycle, or potentially attached to the same cycle. We could always consolidate these two chains into a single chain and attach it to the cycle that yields the largest minimum opportunity path – all sellers in the chain now receive weakly more revenue than they were before. This is doable because each seller subgraph has at least one seller and one buyer to be connected into a chain. Note that this doesn't affect the restriction on cycle length. □

To conclude the discussion of the structure of chains in the optimal solution, we show that these chains must be contiguous and moreover, chains "fill out" the rest of the subgraphs once started.

Lemma 4.3. *If S_i is part of a chain in the optimal solution, then S_j is also part of this chain, for $j > i$, provided that S_j is not part of the cycle that the chain attaches to.*

Proof. Suppose otherwise. Consider the optimal solution with at most one chain. Since S_j is not part of a chain, it must be part of a cycle. As this solution is optimal, it must be the case that adding this whole cycle to the chain decreases the overall revenue. That is, the minimum opportunity path in S_j's cycle is larger than the minimum opportunity path in the cycle that S_i's chain is attached to.

However, this implies that we could add S_i to S_j's cycle, increasing the revenue generated from the seller in S_i without affecting the revenue from any of the sellers in S_j (as S_i

has a larger max-value buyer than S_j), which is a contradiction. □

Lemma 4.4. *Suppose that there is a chain in the optimal solution that connects to a cycle $\mathcal{C}$, and let $S_{min,\mathcal{C}}$ be the smallest subgraph belonging to $\mathcal{C}$. Then there exists an optimal solution where no subgraph larger than $S_{min,\mathcal{C}}$ is part of the chain.*

Proof. Suppose otherwise. If there are two or more such subgraphs, then we could always connect them in a cycle, and their minimum opportunity path would be weakly larger than that obtained by placing them in the chain. If there is a single such subgraph, we could add it to $\mathcal{C}$ – note that this does not affect the minimum opportunity path of the chain or of the sellers in $\mathcal{C}$ and weakly increases the revenue obtained from this subgraph.

If at any point during this process, we obtain a cycle with more than three subgraphs, we can simply split it up into smaller cycles, preserving our desired property. □

Observation 4.2. Note that this implies that once a subgraph is part of a chain, all smaller subgraphs (in terms of max bidder value) must also be part of this chain.

We make one final observation regarding the structure of the optimal solution – namely, all cycles in the optimal solution can be made to be contiguous.

Lemma 4.5. *There exists an optimal solution where all cycles are contiguous.*

Proof. Suppose we have an optimal solution satisfying all the above properties where at least one cycle is non-contiguous. Let S_m be the smallest subgraph that belongs to a non-contiguous cycle (clearly $S_m \neq S_1$).

Firstly, if S_m is a part of a three-cycle $S_a - S_b - S_m$. If $a < b-1$, then we can connect $S_b - S_m$ in a two cycle. By Lemma 4.5, S_{b-1} cannot be in the chain otherwise S_b and S_m would be in the chain. Thus, S_{b-1} is in a cycle. We attach S_a to the cycle that S_{b-1} is in. This weakly increases S_a's revenue because S_{b-1} only connects to larger seller subgraphs.

Now we can focus on non-contiguous two-cycles $S_a - S_m$ and three-cycles $S_a - S_{a+1} - S_m$. Again, by Lemma 4.5, S_{m-1} cannot be a part of the chain otherwise S_m would have been in the chain as well. There are two cases. Suppose that the second highest bidder in S_{m-1} has a weakly higher valuation than the highest bidder in S_m. Consider the following alternative configuration. Connect S_m to S_{m-1} as a chain. (If a chain is already connected to S_{m-1}, add S_m to this chain.) Combine the two cycles that S_{m-1}, S_m belonged to (save for S_m, works also if S_{m-1} is in a 1-cycle.) Additionally, keep any existing chains that were connected to any of the subgraphs involved in the two cycles.

Note that all subgraphs that were previously connected to S_m via a cycle now have weakly increased revenue – the minimum opportunity path in the new cycle is now S_{m-1} rather than S_m. All subgraphs that were previously connected to S_{m-1} have the same revenue; all such subgraphs still have S_{m-1} as their smallest opportunity path within the cycle. It is easy to verify that all chains also have weakly higher revenue.

Now, consider the second case where the second highest buyer in S_{m-1} has a smaller valuation than the highest bidder in S_m, or S_{m-1} does not have a second world buyer. There are now a couple of possible cases:

- Two or more subgraphs were previously connected to S_m, S_{m-1}. Connect S_m, S_{m-1} in a two-cycle. In this case, connect all the subgraphs that were previously connected to S_m, S_{m-1} together into a single cycle. Note that we lose $v(S_{m-1}) - v(S_m)$ in revenue from connecting S_m and S_{m-1} together, but we gain at least $v(S_{m-2}) - v(S_m)$ from the other subgraphs.
- S_m belonged to a 2-cycle and S_{m-1} belonged to a 1-cycle. If the other subgraph in the S_m cycle was S_{m-2}, then connect S_{m-2}, S_{m-1} together and leave S_m as a 1-cycle. Note that we get weakly more revenue from S_{m-2} in this case, and a 1-cycle with S_m is less costly than a 1-cycle with S_{m-1}.
- S_m belonged to a 2-cycle and S_{m-1} belonged to a 1-cycle. Connect S_m, S_{m-1} in a two-cycle. If the other subgraph in the S_m cycle was not S_{m-2}, connect the other subgraph in S_m's cycle to the cycle that S_{m-2} is currently a part of. Note that we gain $v(S_m)$ from connecting S_m and S_{m-1} together, and at least $v(S_{m-2}) - v(S_m)$ from the other subgraphs.

In all the above cases, we keep all existing chains as they are – they are weakly better off under this new configuration because the second highest buyer in S_{m-1} is smaller than $v(S_m)$. We continue this process until there are no non-contiguous cycles, at each step weakly increasing our revenue. If at step t, the smallest subgraph belonging to a non-contiguous cycle is S_m, then at step $t+1$, the process ensures that the smallest possible subgraph belonging to a non-contiguous cycle is S_{m-1}, meaning that this process must terminate. □

Corollary 4.1. There exists an optimal solution where all cycles are contiguous, there is (possibly) a single contiguous chain, and all cycles are limited to length 3 or less.

Finally, we have placed sufficient structure on the optimal solution to be able to develop a poly-time algorithm. In essence, we restrict our search to configurations such as those described in the above corollary.

Theorem 4.2. *There exists a poly-time algorithm for the AMOS case, when $|B| = |S|$.*

Proof. Fix an arbitrary contiguous cycle for the chain to attach to. Note that we are limited to $0, 1, 2$, or 3-cycles, for which there are $O(n)$ possibilities. Fixing this cycle, consider the induced graph G' obtained by removing all the subgraphs in the cycle from the original graph.

Let $DP[i]$ denote the maximum obtainable revenue from connecting the first i subgraphs in G' via cycles. Filling in the base cases for $DP[1], DP[2], DP[3]$, we have the following recurrence:

$$DP[i] = \max(DP[i-3] + \text{3-cycle revenue}, \\ DP[i-2] + \text{2-cycle revenue}, DP[i-1] + \text{1-cycle revenue})$$

Let $R[i]$ denote the revenue obtained 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 $\max_i DP[i] + R[i]$.

Repeating this process for all $O(n)$ possible fixed cycles, we can take the maximum revenue from any of these configurations. This process 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, where cycles are of length at most 3, there is at most one chain, and this chain is contiguous. By Corollary 4.1, there exists an optimal solution that satisfies these properties and thus we find it, concluding the proof. □

Lemma 4.6. *When $n = |B| > |S| = m$, sort buyers by their valuation, and denote the set of the m largest buyers by $B_m = \{b_1, b_2, ..., b_m\}$. Then perform the revenue-optimal matching on S, B_m while ignoring $B \setminus B_m$. This is revenue optimal.*

Proof. By Lemma 2.2, there exists a revenue-optimal transaction where all sellers trade. Otherwise, all sellers connecting to this non-trading seller through an opportunity path have a price of zero. By adding an edge between this seller and a non-trading buyer (guaranteed to exist as $|B| > |S|$), other sellers' prices weakly increase.

Further, no buyer in $B \setminus B_m$ trades. Suppose otherwise and take any configuration where a buyer in $B \setminus B_m$ trades. Let b_{min} be the smallest trading buyer, and let b_{max} be the largest non-trading buyer, noting that $v(b_{min}) < v(b_{max})$. Suppose that b_{min} currently transacts with seller s.

As the market clears according to the maximum weight matching, it must be that s is not connected to b_{max} – otherwise, matching them would yield a strictly higher welfare matching. Thus, we can add a platform edge from s to b_{max}, ensuring that they transact in the new matching. Note that b_{min} no longer transacts as it was the minimum value transacting buyer, and thus the maximum welfare matching can no longer include it.

We claim that all sellers' prices weakly increase. The only modified opportunity paths that previously existed are those which previously went through b_{min} – as previously mentioned, there are no opportunity paths that terminate at a non-trading seller, so this was the minimum possible opportunity path. Now, no opportunity paths terminate at b_{min} (as they no longer trade), meaning that the minimum opportunity path weakly increased (in the event of a tie with b_{min}).

We can continue this process until no buyers in $B \setminus B_m$ trade. Buyers not trading will not be in any opportunity path hence not affect the characterization results nor optimality of the procedure in Theorem 4.2. □

Lemma 4.7. *When $n = |B| < |S| = m$, discard $m - n$ of the sellers who do not have world edges, and perform the revenue-optimal matching on the induced graph. This is revenue optimal.*

Proof. By Lemma 2.2, note that there exists a revenue optimal matching in which all buyers transact. Suppose that there is a seller s who has at least one adjacent world edge who does not transact. Consider any buyer b who is in seller s's subgraph.

Clearly, buyer b must transact with some seller s' via a platform edge (as they do not transact with the seller in their subgraph). Consider eliminating this platform edge (b, s'). No revenue is lost from b as they had an opportunity path to s, who did not transact, so they paid price 0. Secondly, any sellers who previously had opportunity paths to b also received price 0, so removing this edge does not decrease their price either.

It follows that there exists a revenue optimal matching where all sellers with an adjacent platform edge transact, which implies that we can eliminate $m - n$ sellers with no adjacent world edges, as they would not transact in this revenue optimal matching regardless. □

References

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

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

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

Kelso, A. S.; and Crawford, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, 1483–1504.

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