Price Steering in Two-Sided Platforms

Pedro Vaissman Guinsburg*, Gustavo Quinderé Saraiva†

Latest draft: click here. This draft: April 17, 2022.
First draft: April 10th, 2022.

Abstract

We study the incentives from a two-sided platform to segment the market by providing personalized search results. In our environment, a monopolistic platform is in charge of matching sellers to buyers. Upon being matched, each pair of buyer and seller negotiate prices. If they choose to transact, the platform receives a commission fee that is proportional to the value of the transaction. The platform is assumed to have full information over customers' and sellers' outside option. We show that in this environment the platform may have incentives to prioritize finding feasible matches to more expensive products, so as to inflate market prices, and thus, the commissions it receives from transactions. By doing this, the platform maximizes the number of transactions, which can generate excess liquidity.

Keywords: Market Segmentation, Information Design, Two-sided markets

1 Introduction

With the proliferation of online marketplaces, such as Amazon and eBay, many expected that information frictions in these markets would disappear in the long run, thus improving customers' welfare. But theoretical and empirical research have emerged suggesting that platforms may have incentives to obfuscate information or segment the market in a way that could reduce customers' welfare. In this article we conduct a theoretical analysis of a type of market distortion that has been neglected by the literature, which occurs when the platform steers customers towards certain products based on their willingness to pay. We show that, in an environment in which the platform can discriminate search results, but cannot choose the prices of the goods sold by third-party sellers through its platform, the platform may have incentives to match those with higher willingness to pay to more expensive products in a way that may generate more transactions than the socially optimum. This occurs if part of the commission fee charged by the platform is proportional to the price of the product being sold, as it is common in may online market places, such as in Amazon.com and eBay.com.[1]

This research is motivated by growing concerns that e-commerce websites may be discriminating search results based on users' data. As an example, a study in 2012 has found that the travel website Orbitz offered consistently more expensive prices for hotels to Mac users, as opposed to Windows users (Mattioli

*Vaissman Guinsburg: São Paulo University (USP), Avenida Professor Gualberto 908, Prédio FEA 2, São Paulo, 05508-010; guinsburg@usp.br.

†Business School, Pontificia Universidad Católica de Chile,Vicuña Mackenna 4860, Macul, Santiago, Chile.

[1] For Amazon's fee policy, see https://sellercentral.amazon.com/gp/help/external/200336920. For eBay, see https://pages.ebay.com/seller-center/seller-updates/2022-winter/fees-update.html.

(2012)). Other studies, such as Hannak et al. (2014), have also found evidence that many e-commerce websites practice personalized search results as a function of customers' cookies and browsing history. In 2000, it became public that Amazon.com was conducting A/B testing experiments on the prices of certain DVD's. Such policy was then quickly removed after many customers started complaining about the practice (Weiss (2000)).

Motivated by this empirical evidence, we build a theoretical framework in which a platform that intermediates transactions dictates one-to-one matches between buyers and sellers, in order to charge a commission fee per sale, part of which is proportional to the final price chosen by the seller. We show that the platform has incentives to maximize the number of transactions, which can be greater than the socially optimum, Pareto Efficient, competitive equilibrium allocation. The way the platform would induce this excess liquidity is by allocating sellers with high production cost to customers with high willingness to pay, while allocating sellers with low production cost to customers with low willingness to pay, so as to increase market prices, and therefore, the commissions paid to the platform. For such a policy to be feasible, the platform must have information on customers' willingness to pay, as well as sellers' cost structure or willingness to sell.

The intuition as to why a platform would want to implement such a policy can be best described with the help of a simple example. Consider an economy with 5 sellers and 5 buyers. Each seller is willing to sell only one unit of their product, and each buyer has unitary demand. Each seller $s_j \in \{s_1, s_2, \cdots, s_5\}$ sells a product with quality $q_j \geq 0$, with a production costs of $c_j \geq 0$, where c_i could also be interpreted as the seller's outside option, i.e., the price at which they can sell their product outside of the platform. Each buyer $b_i \in \{b_1, b_2, \cdots, b_5\}$ has an outside option $u_i \geq 0$. In figure 1 we order sellers' net value, $q_j - c_j$, in descending order, and buyers' outside option, u_i, in ascending order.

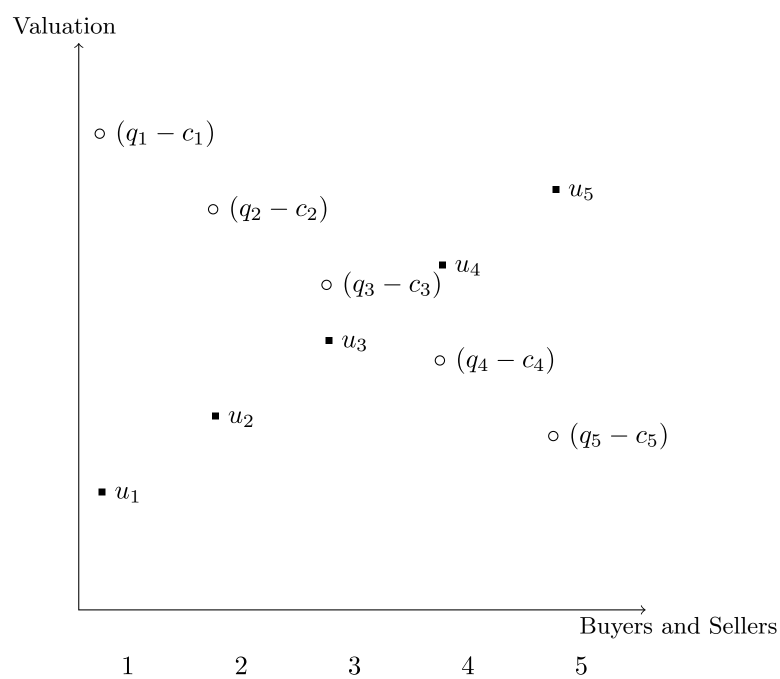

Figure 1: A plot depicting the net value by each seller, and the outside option from each buyer.

From figure 2, we can see that the total surplus from this economy is maximized when buyers b_1, b_2, b_3 (i.e., those with highest willingness to pay) transact with sellers s_1, s_2, s_3 (those who generate most net value), and the remaining buyers and sellers do not transact, so that only 3 transactions take place in the economy. In this case, total surplus could be approximated by the blue region from the graph.

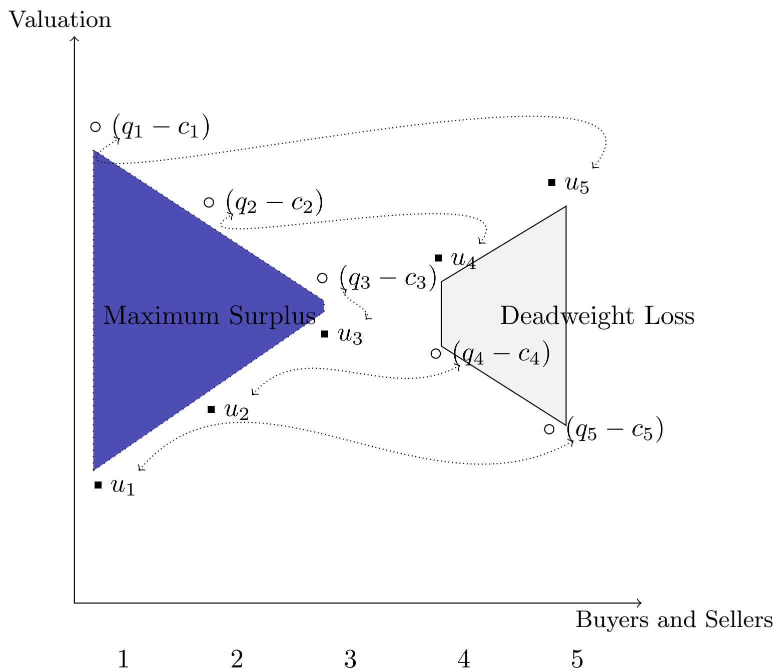

Figure 2: The plot depicts an approximation of the total surplus obtained when agents are matched so as to maximize the total number of transactions.

But the platform could, in principle, induce 5 transactions by matching each buyer $b_i \in \{b_1, b_2, \cdots, b_5\}$ with seller s_{6-i}. This matching is displayed in figure 2 by the curved arrows. Because in each match the net value created by the seller surpasses the customer's outside option, each match should yield a transaction, thus generating a total surplus given by the blue area minus the deadweight loss given by the grey area in the figure.

Assume now that the platform collects a fixed fee per transaction that does not depend on the final price negotiated between buyer and seller, and assume that all products have the same quality q_i. In this case, the platform would have incentives to maximize the total number of transactions. Zhao et al. (2011) have shown that the number of transactions can be maximized by iteratively matching sellers with low costs with buyers with low willingness to pay, until a new match generates no positive surplus. In other words, they propose a matching criterion that prioritizes finding potential buyers to sellers with low production costs (i.e., they prioritize finding matches to more efficient sellers). So after implementing their algorithm, we have that if a seller with cost c_j manages to sell its product, a seller with cost $c'_j < c_j$ will also be able to sell its product. It can be shown that the final allocation obtained after implementing this algorithm is not necessarily Pareto efficient. So this result already illustrates how the platform may have incentives to sacrifice efficiency in exchange of higher revenues.

But we show that, to maximize the total number of transactions, the platform could implement an alternative and even more inefficient algorithm: it could prioritize securing feasible matches to sellers with *high* costs, by iteratively matching them with buyers with low willingness to pay, until a new match yields no positive revenue. So this could lead to situations in which a seller with a high cost manages to sell its product, but a seller with low cost (or low outside option) does not. While the platform has no strict incentives to adopt this more inefficient strategy in this scenario where commissions do not depend on the value of the transaction, things change when the platform also charges a commission fee that is proportional to the final price paid by customers, as it is common in many e-commerce websites, such as Amazon.com and eBay.com. We show that, in this case, the platform will have strict incentives to secure good matches to high cost sellers so as to inflate market prices, which in turn increases the commissions received by the platform per transaction. So in our example, if the commission fees are sufficiently small, it may be beneficial for the platform to match each buyer $b_i \in \{b_1, b_2, \cdots, b_5\}$ with seller s_{6-i}, so as to inflate market prices, and thus, the commissions paid by sellers.

Another generalization from Zhao et al. (2011) is that we can also derive the optimal matching from the platform when products are vertically differentiated, i.e., when the quality q_i from sellers are allowed to differ. We also allow the fees chosen by the platform to be jointly maximized with the matching assignment, and show how to find an optimal solution to this problem.

Notice that in our environment price negotiations between buyers and sellers happen after the match is determined by the platform. While this would be more consistent with market structures in which transactions are made through auctions or bargaining, such as in eBay.com or Upwork.com, respectively, our results also apply to situations in which sellers first commit to a price, and then the platform determines the match, and sellers are not sophisticated enough to predict the effect of their prices on the match that they get from the platform; or to situations in which, though sellers can predict the effect of prices on the matching assignment, they are bound to charge a price (net of commission fees) less than or equal the price they practice outside of the platform lest they suffer public criticism for practicing price discrimination, or violate some price parity clause imposed by the platform. In the latter case, seller s_i would charge a price equal to his outside option c_i, i.e., equal to the price it can sell its product outside of the platform, plus the commissions charged by the platform.

Notice that the source of inefficiency from the model comes from the fact that the platform does not have control over the final prices charged by the third-party sellers. In order to regain some control over prices, the platform practices a market segmentation that prioritizes the sale of more expensive products. Because the higher costs required to produce those more expensive products are not internalized by the platform, the resulting allocation will often be Pareto inefficient.

While the results derived in this paper have been framed in terms of a platform that intermediates transactions, it also applies to other two-sided markets, such as the one intermediated by real estate agencies.

2 Related Literature

The theoretical literature on incentives of e-commerce platforms is filled with examples that show how platforms may be willing to create information frictions that can potentially hurt customers, such as the seminal work of Diamond (1971) and Anderson and Renault (1999). In their theoretical model, customers sequentially search for a product through indirected search, so that at each period customers are randomly matched with a product and decide whether to buy it at the advertised price or keep searching for a better match. Like in our environment, they assume that prices of the final good are set by the third-party sellers, so that the platform cannot directly control the final prices of the products in the market. They show that, if there are infinitely many sellers operating in the platform selling the same homogeneous product, then, for any given search cost, no matter how small, there is an equilibrium in which all sellers choose the monopoly price. This, result known as the Diamond Paradox, is arguably very puzzling, as one would expect that competition in a market with many providers selling the same homogeneous product, would drive prices down to the perfectly competitive level.

An implication of this result is that, if part of the platform's commissions are proportional to the price of the product sold (as it is common in many two-sided platforms, such as Amazon, eBay and AirBnB), then the platform may have incentives to foster search frictions in order to increase the final price charged by sellers. This observation has led others to present specific mechanisms through which platforms can increase search costs, such as the work of Eliaz and Spiegler (2011) and Casner (2020), who show that a platform may have incentives to allow some low quality sellers to enter the market in order to obfuscate search. These results illustrate how two-side platforms are sometimes willing to add inefficiencies into their market in exchange of regaining some control over the prices chosen by third-party sellers. These studies assume, however, that the platform is not in control over who gets matched with whom in the search process, whereas in our environment the platform is the one determining the matches, and uses this ability as a tool to regain some control over market prices.

This paper is also related to the literature on double auctions (i.e., where an auctioneer intermediates transactions between buyers and sellers), such as the work of Deshmukh et al. (2002), Ausubel et al. (2017) and Satterthwaite et al. (2022). But because the auctioneer is usually interested in implementing efficient and/or incentive compatible mechanisms, this literature has usually proposed and studied auction formats in which the sellers with lower opportunity costs are the ones who end up transacting (assuming that the products sold are homogeneous) as it is the case in the uniform and pay-as-bid auctions.[2]

More in line with our approach, Zhao et al. (2011) assume that the auctioneer already knows the valuations from buyers and sellers, so that they can abstract from incentive compatibility issues. They propose the most efficient algorithm to match buyers and sellers among the ones that maximize the total number of transactions in the economy. This algorithm is attractive if the auctioneer collects a fixed revenue per transaction that does not depend on the final price at which the product is sold. This algorithm can generate too much liquidity, as the equilibrium number of transactions can be above the socially optimum. In our specification, on the other hand, we assume that the platform (i.e., the auctioneer) also collects a fee per transaction that is proportional to the final price at which the product is sold. So we predict that the matching mechanism used by the platform will be even more inefficient than the one proposed by Zhao et al. (2011), as the platform will not only try to maximize the total number of transactions, but it will also prioritize finding eligible buyers to sellers with high outside option, so as to inflate the final prices paid by customers, and therefore, the fees received by the platform. Moreover, we also generalize the result from Zhao et al. (2011) by finding an optimal algorithm to the platform when products are vertically differentiated.

This paper is also related to the literature on market segmentation, such as the work of Bergemann et al. (2015), Haghpanah and Siegel (2019) and Hao Yang (2020). This literature, however, usually assumes that the seller segmenting the market is in control over the prices of each one of its products. In such an environment, market segmentation usually has the effect of improving total welfare. As an example, for the extreme case in which the seller practices perfect price discrimination by charging each customer their willingness to pay, total surplus is maximized, though in this case customers earn zero surplus. In our environment, on the other hand, the platform intermediating transactions between buyers and sellers cannot dictate the prices charged by sellers. As a result, attempts by the platform to segment the market in its

[2]Both the uniform and pay-as-bid auctions are Pareto efficient when participants submit their true valuations and the auctioneer does not collect distortionary commission fees. While neither mechanism is strategy-proof (e.g., see AUSUBEL et al. (2014) or Binmore and Swierzbinski (2000)), the uniform price auction is known to be "*almost*" strategy-proof when the number of participants in the market is large and each participant has a small market share on their side of the market (Azevedo and Budish (2019)).

favor will generate inefficiencies not predicted by the previous literature.

This paper is also related to the literature on *online algorithms* used in the market for AdWords, such as the work of Karp et al. (1990) and Mehta et al. (2007). In this market, advertisers bid for search queries (e.g., “Hotel in Miami”, “restaurants near me”, “gas station”, etc.) and provide the platform with a budget constraint, specifying the maximum number of ads that they are willing to pay. Upon receiving those bids, the platform is tasked with matching customers with advertises. In a pay-per-click business model, the advertiser only pays his bid if matched with a customer who actually clicks on the ad. In this case the platform has incentives to only match customers with the ads that they are likely to inspect. While the platform knows customers’ tastes based on their browsing history and other information that the platform can collect (e.g., users’ operating system, location, etc.), it does not know the sequence of queries that customers will make. One way the platform could try to form the matching is by implementing a *greedy algorithm* that prioritizes finding good matches for the advertisers with highest outstanding bids among the ones who have not yet exhausted their budget constraint. One problem with this approach, however, is that it may generate very few matches, which, in this environment, would be inefficient. Indeed, consider the case in which there are two advertisers, A and B. Customers who search for “golf” are willing to be matched with either A or B, but those who search for “sports” only like to be matched with A. Now suppose that the budget constraint from both sellers are such that they can each only pay for one ad inspection, and that seller A pays more per click than seller B (say, because seller A placed a higher bid for both queries in the AdWords auction). Then, if the first query made is “sports”, it is not necessarily on the platform’s best interest to match that customer with advertiser A, as the next query could be “sports” again, in which case, no new matches would be made. If, on the other hand, the platform had matched the first “sports” query to advertiser B, then, in the next period it would still be able to collect revenue from advertiser A by matching him the next “sports” query. Based on this observation, this literature has proposed algorithms that put positive weight not only on how much each agent has bid on a query and how likely someone is to click on that ad (i.e., its click through rate), but also on how easy it is to find good matches for those bidders: if the algorithm notices that it is having a hard time exhausting the budget constraint for a certain seller, it will start prioritizing forming feasible matches for that seller. But while in such an environment, maximizing the number of transactions usually improves the overall level of efficiency in the economy, we show that in our case it can actually reduce the total surplus.

3 Model

A market is comprised of a set of n buyers, $B \equiv \{b_1, b_2, \cdots, b_n\}$, and a set of m sellers, $S \equiv \{s_1, s_2, \cdots, s_m\}$. Buyers have a unitary demand for the products sold by the sellers, and each seller can only sell one unit of their product to a single customer. We assume that all customers have the same cardinal preferences over the products sold by the sellers. We denote $q_j \in \mathbb{R}$ as the quality of the product sold by seller $s_j \in S$, and $c_j \in \mathbb{R}_+$ as the cost of producing such a product or the outside option for selling this product, i.e., c_j can be interpreted as the price at which seller s_j can sell the product outside the platform. Each buyer $b_i \in B$ has an outside option given by $u_i \in \mathbb{R}_+$.

To simplify the exposition, we order buyers in ascending order of their outside options, and sellers in descending order of their net quality, defined by their quality minus their cost. More precisely, unless specified otherwise, we assume that

$$u_1 \leq u_2 \leq \cdots \leq u_n \tag{1}$$

and

$$q_1 - c_1 \geq q_2 - c_2 \geq \cdots \geq q_m - c_m. \tag{2}$$

A monopolistic platform is responsible for intermediating transactions between buyers and sellers. More precisely, the platform will match customers to sellers through a match function μ. In this environment we will only allow for one-to-one matches, i.e., each seller can be matched with at most one customer and vice versa, though the intuition of our main results may be extended to many-to-one matching environments. The matching function μ will map agents on one side of the market to agents on the other side of the market, with the possible exception of an agent being matched with himself to account for the possibility that the agent remains unmatched.

Definition 1. *A matching μ is a function mapping $B \cup S$ into itself, such that:*

i) $\forall b \in B$, $\mu(b) \in S \cup \{b\}$.

ii) $\forall s \in S$, $\mu(s) \in B \cup \{s\}$.

iii) $\forall b \in B$ *and* $s \in S$, $s \in \mu(b)$ *if and only if* $\mu(s) = b$.

Once the matching is determined, matched agents can transact. If customer b_i is matched with seller s_j, they will generate a net surplus equal to $\max\{q_j - c_j - u_i, 0\}$. A matching μ is said to be Pareto Efficient if it maximizes total net surplus.

It can be shown that an efficient allocation should prioritize the matching of customers with the lowest outside option and sellers with the highest net quality (i.e., the highest quality minus cost). One way to achieve this is by implementing *positive assortative matching* (PAM), i.e., by having the customer with the lowest outside option being matched with the seller with highest net quality, the customer with second lowest outside option being matched with the seller with second highest net quality, and so on. Agents who remain unmatched will generate zero surplus, as well as those who end up in a match in which the outside option of the buyer surpasses the product's net quality. But there are other matches that are also efficient. In fact, if we initially apply PAM and then apply any permutation to the matching assignment of those who transact under this match, then the final matching assignment will be Pareto efficient.

Proposition 1. *A matching μ is Pareto Efficient if and only if each customer b_i such that $i \leq \min\{n, m\}$ and $u_i \leq q_i - c_i$ (i.e., if the buyer is willing to transact with the seller with same index) is matched to a seller s_j such that $u_j \leq q_j - c_j$ (i.e., with a seller who, if matched with a buyer with same index, would transact).*

Proof: Proof in the Appendix. ∎

Corollary 1. *(PAM is Pareto Efficient) Market efficiency is achieved whenever $\mu(s_i) = b_i$ for all $i \leq \min\{n, m\}$ (i.e., each seller is matched with the buyer with the same index, if there is any).*

It follows directly from proposition 1 that a Pareto efficient match μ maximizes the number of transactions subject to the constraint that every seller who transacts under μ must also be willing to transact with any other buyer who transacts under μ.

Corollary 2. *For a given match μ, let $FM_\mu \equiv \{(s_j, b_i) \in S \times B; \mu(s_j) = b_i$ and $q_j - b_i \geq 0\}$ (i.e., FM_μ represents all the matches associated with μ, excluding those that do not entail transactions). If a matching*

μ is Pareto efficient, then

$$\begin{aligned} &\mu \in \arg\max_{\mu'} |FM_{\mu'}| \\ &\quad s.t. \quad s_j - b_{i'} \geq 0 \quad and \quad s_{j'} - b_i \geq 0 \quad \forall (s_j, b_i), (s_{j'}, b_{i'}) \in FM_{\mu'} \end{aligned} \tag{3}$$

Now suppose that the platform receives a fixed share $\psi \in [0,1]$ of the price paid by the customer for each realized purchase, plus a fixed commission fee $\tau \geq 0$ per transaction. Prices are determined ex-post, after the matches are formed, and after agents learn about the commission fees ψ and τ charged by the platform. In practice, prices could be determined after the matches takes place if the products are being sold through an auction, or if buyers and sellers are allowed to bargain after they are matched with one another, as in Upwork.com. While in principle one would expect an auction platform to match many potential buyers to a single seller in order to maximize the auctioneer's expected revenue, under the assumption of private independent values, the final price paid by the winning bidder is usually only affected by a few bids. Under the second price auction mechanism, for instance, agents have a weakly dominant strategy of bidding their true valuation, and the price paid only depends on the second highest bid. In this case, if the platform knew buyer's valuation, it would be wasteful to match more than two bidders to a given seller, as the bids below the second highest price would not affect the final price being paid, and those bidders could actually have helped increase the final price from other auctions. So our one-to-one matching environment can still give us some insight as to what could happen in platforms that intermediate some transactions through auctions, such as eBay.com.

We assume that prices are determined through Nash Bargaining, where each buyer has bargaining power $\alpha \in [0,1]$ and each seller has bargaining power $1-\alpha$. More precisely, if $\psi > 0$ and seller s_j is matched with a customer b_i such that $(1-\psi)(q_j - u_i) \geq c_j + \tau$, then a transaction will occur, and the seller will set its price at

$$p_j^*(b_i, \psi, \tau) = (1-\alpha)(q_j - u_i) + \alpha \frac{c_j + \tau}{1-\psi}, \tag{4}$$

generating a revenue of

$$\psi p_j^*(\psi, \tau, b_i) + \tau = \psi \left[(1-\alpha)(q_j - u_i) + \alpha \frac{c_j + \tau}{1-\psi} \right] + \tau$$

to the platform. If $(1-\psi)(q_j - u_i) < c_j + \tau$, then no transaction transpires from the match between seller s_j and buyer b_i, so that the platform earns zero revenues from this match. If $\psi = 1$ and $c_j = \tau = 0$, the seller is indifferent between not selling and selling its product at any price. To guarantee that the profit function is upper semicontinuous, we will assume that in this case the seller will charge $p_j = \max\{q_j - u_i, 0\}$.[3]

Notice that, when $\alpha = 1$, i.e., when buyers have all bargaining power, seller s_j elicits the price $(c_j + \tau)/(1-\psi)$ regardless of who he is matched with. Because in this case the matching algorithm chosen by the platform does not affect sellers' pricing strategy, we could invert the timing of the model by assuming that sellers first elicit a price, and then, based on those prices, the platform decides who gets matched with whom, which would be more consistent with prominent e-commerce websites, such as Amazon.com. In practice, this could happen if, for instance, the platform imposes a price parity clause that prohibits each seller s_j

[3]The hypothesis of upper semicontinuity of the objective function with respect to (ψ, τ) guarantees the existence of an optimal (ψ, τ) for any given match μ. Because there is a finite number of matches, this implies that there exists a tuple (μ, ψ, τ) that maximizes the profits of the platform (see proposition 3).

from charging a price, net of commission fees, higher that the price c_j at which he sells its products outside of the platform (i.e., it prohibits each seller s_j from charging a price above $(c_j+\tau)/(1-\psi)$).

In this environment we can easily find examples in which the platform has no incentives to implement a Pareto efficient match.

Example. *For a given commission fee, $\psi \in [0,1]$ and $\tau \geq 0$, the optimal match μ^* chosen by the platform is not necessarily Pareto efficient.*

Indeed, consider a market where the set of sellers is given by $\{s_1, s_2, s_3, s_4\}$, and the set of buyers is given by $\{b_1, b_2, b_3, b_4\}$. Sellers' quality and costs are given by

$$(q_1, q_2, q_3, q_4) = (20, 20, 20, 20) \quad \textit{and} \quad (c_1, c_2, c_3, c_4) = (2, 8, 14, 16),$$

respectively. Customers' outside option are given by

$$(u_1, u_2, u_3, u_4) = (0, 3, 5, 7).$$

In what follows, we use the notation $s_j : b_i$ to indicate that seller s_j is matched to buyer b_i. In other words, $s_j : b_i$ is equivalent to writing $\mu(s_j) = b_i$ or $\mu(b_i) = s_j$.

In a Pareto Efficient allocation, buyers b_1, b_2 and b_3 should be each matched with one of the first three sellers with highest net quality (i.e., with sellers s_1, s_2 and s_3). But if the platform charges a commission fee of $\psi = 0.5$ and $\tau = 0$, then all of the Pareto efficient matches would generate a total profit of 9.75 *to the platform, whereas if it implemented the match*

$$s_1 : b_4, \quad s_2 : b_3, \quad s_3 : b_2, \quad s_4 : b_1$$

its profits would equal 12.25.

Intuitively, the reason why the platform may not have incentives to implement a Pareto efficient allocation can be explained as follows. We can write the platform's revenue for a match between a seller s_j and a customer b_i such that $(1-\psi)(q_j - u_i) \geq c_j + \tau$, as

$$\psi(1-\alpha)(q_j - c_j - \tau - u_i) + \frac{\psi(1-\psi+\alpha\psi)}{1-\psi}(c_j+\tau) + \tau. \tag{5}$$

While the first term from expression 5 is increasing in the total surplus created from the transaction between seller s_j and buyer b_i, $(q_j - c_j - u_i)$, the last terms from this expression are increasing in the opportunity cost from selling, c_j, and on τ. So the platform is not exclusively interested in maximizing the creation of value, but also on maximizing the total number of transactions so as to collect more revenues through τ, and also on facilitating the transaction of more expensive products, as that increases the final price charged by the seller, and therefore the commissions collected through ψ.

There is an instance, however, in which the platform will have incentives to implement a Pareto efficient match: if all sellers have zero production cost (i.e., if $c_j = 0$ for all $j \in \{1, 2, \cdots, m\}$) and $\tau = 0$. Indeed, in this case expression 5 reduces to

$$\psi(1-\alpha)(q_j - u_i), \tag{6}$$

so that the platform's revenue is proportional to the sum of value created form transactions.

Proposition 2. *Suppose $c_j = 0$ for all $j \in \{1, 2, \cdots, m\}$. Then, for any given $\psi \in [0, 1]$ and for $\tau = 0$, a Pareto efficient match μ maximizes the profits of the platform.*

Proof: If $c_j = 0$ for all $j \in \{1, 2, \cdots, m\}$ and $\tau = 0$, then, for any given $\psi \in [0, 1]$ and any matching function μ, a matched agent will transact under μ and $\psi \in [0, 1]$ if and only if he transacts under μ and $\psi = 0$ (i.e., the commission ψ does not affect agents' willingness to transact). Therefore, from expression 6 we have that the platform's revenue for a given match μ is proportional to the total value created from transactions when there are no commissions. So the platform has incentives to implement a match that maximizes total surplus. ∎

Intuitively, proposition 2 holds because, if costs are zero, the fee charged by the platform is non-distortionary (i.e., it does not affect agents' choices), in which case maximizing the creation of value allows the platform to extract more surplus.

The hypothesis that marginal costs are zero is, however, arguably inconsistent with the hypothesis that sellers face capacity constraints. Indeed, in our model we assume that each seller can only sell one unit of their product. So for the markets that we are considering, one would expect production costs to be non negligible, so that the fee charged by the platform affects equilibrium prices, in which case the platform may have incentives to choose an inefficient match.

4 Optimal Match with exogenous commissions

In this section we will consider the problem of finding the optimal match taking the commission fees, ψ and τ, as given, i.e., we will characterize the solution to the following problem:

$$\max_{\mu} \tau |S_\mu| + \psi \sum_{s_j \in S_\mu} p_j^*(\mu(s_j), \psi, \tau), \tag{7}$$

where

$$S_\mu \equiv \{s_j \in S; \mu(s_j) = b_i \in B \wedge (1 - \psi)(q_j - u_i) \geq c_j + \tau\}$$

corresponds to the set of sellers who end up transacting given the fees ψ and τ and matching function μ, and $p_j^*(\mu(s_j), \psi, \tau)$ is the equilibrium price given by equation 4.

In principle, finding the optimal match can be computationally costly. Indeed, for a market with n buyers and n sellers, there would be $n!$ different ways to match each seller to a (different) customer. So if, for instance, $n = m = 20$, the number of possible matches between buyers and sellers would be $20! = 2.43\text{e}{+}18$, in which case finding the optimal match through brute force would be infeasible. This motivates the creation of an algorithm that can solve the platform's problem in polynomial time.

For pedagogic purposes, it is useful analyzing the platform's optimum policy when sellers' products are homogeneous, i.e., when $q_j = \overline{q} \, \forall j \in \{1, 2, \cdots, m\}$. In this case, a simple algorithm can be used to find the platform's optimal match for a given $\psi \in [0, 1]$ and $\tau \geq 0$. The algorithm consists in first matching the seller with highest cost with the buyer with lowest outside option who, if matched with this seller, would end up purchasing the product from this seller. If there is no such buyer, the seller remains unmatched. Then, for the seller with second highest cost, match this seller with the buyer with lowest outside option who would be willing to transact with this seller among the buyers who have not been matched yet. Again, if there is no such buyer, this seller remains unmatched. Continue this process until the match from each seller has been determined or until all buyers have been matched. As it turns out, this algorithm not only maximizes the profits from the platform, but also the number of transactions that can be made, given (ψ, τ).

Algorithm 1. *For each $B' \subseteq B$, each $s_j \in S$ and each $\psi \in [0,1]$ and $\tau \geq 0$ define*

$$F(s_j, B', \psi, \tau) \equiv \{b_i \in B'; (1-\psi)(q_j - u_i) \geq c_j + \tau\},$$

i.e., $F(s_j, B', \psi, \tau)$ is the set of "feasible buyers" in B' for seller s_j, or more precisely, it is the set of buyers in B' who, if matched with seller s_j, would end up purchasing their product.

i) Initialize $j = m$ and define $B_j = B$.

ii) Compute $F(s_j, B_j, \psi, \tau)$. If $F(s_j, B_j, \psi, \tau) \neq \emptyset$, set $\mu(s_j)$ equal to $b_i \in F(s_j, B_j, \psi, \tau)$ such that $u_i \leq u_l$ for all $u_l \in F(s_j, B_j, \psi, \tau)$ (i.e., match this seller with the buyer with lowest outside option who has not been matched yet) and define $B_{j-1} = B_j \backslash \{b_i\}$, else set $\mu(s_j) = s_j$ (i.e., keep seller s_j unmatched) and define $B_{j-1} = B_j$. Then proceed to the next step.

iii) If $j > 1$, redefine $j = j - 1$ and repeat step ii), else, stop the algorithm.

Theorem 1. *Suppose that $q_j = \overline{q} \, \forall j \in \{1, 2, \cdots, m\}$ and $c_1 \leq c_2 \leq \cdots \leq c_m$. Then, given $\psi \in [0,1]$ and $\tau \geq 0$, the matching allocation obtained from algorithm 1:*

i) Maximizes the number of transactions that can be generated in the economy, i.e., it maximizes $|S_\mu|$.

ii) Maximizes the profits of the platform, i.e., it solves the maximization problem 7.

Proof: In the Appendix. ■

Intuitively, part i) from theorem 1 holds because, when all sellers have the same quality and only differ on their production cost, algorithm 1 prioritizes matching sellers for which it is hardest to find a feasible buyer. By iteratively matching those sellers first, the market is iteratively left with sellers for which it is easiest to find a feasible match, thus allowing the platform to form more matches.

Intuitively, part ii) from theorem 1 holds because, when $q_j = \overline{q} \, \forall j \in \{1, 2, \cdots, m\}$, the transactions that generate most revenue to the platform are the ones between sellers with high production cost and customers with low outside option. Because algorithm 1 prioritizes the matching of those sellers and buyers, and because it is the matching that generates most transactions (part i) from the theorem), this then implies that this matching algorithm maximizes the platform's revenues.

So we conclude that, if the products sold are homogeneous, the platform will implement algorithm 1, which does not always yield a Pareto efficient allocation (see example 3). The next corollary presents a sufficient condition under which the platform's optimal policy is not Pareto efficient.

Corollary 3. *(If the Buyer's side is the short side then the platform's optimal matching will be Pareto inefficient) For each seller $s_j \in S$, let*

$$F(s_j, \psi, \tau) \equiv \{b_i \in B; (1-\psi)(q_j - u_i) \geq c_j\},$$

i.e., $F(s_j, \psi, \tau)$ is the set of buyers who, if matched with seller s_j, would end up purchasing their product. Without loss of generality, suppose that $F(s_j, \psi, \tau) \neq \emptyset$ for all $s_j \in S$ (i.e., all sellers have at least one *feasible buyer).*[4] *Suppose in addition that*

$$c_1 < c_2 < \cdots < c_m.$$

and

$$q_1 = q_2 = \cdots = q_n = \overline{q}.$$

Then, if $|S| > |B|$ (i.e., there are more sellers than buyers) and $q_j = \overline{q}$ for all $s_j \in S$ (i.e., the products are homogeneous) then the platform's optimal policy will be Pareto inefficient.

Proof: If the products are homogeneous, then clearly, in a Pareto efficient allocation we must have that, whenever a seller $s_j \in S$ transacts, a seller $s'_j \in S$ such that $j' < j$ must also transact, as $c_{j'} < c_j$. Because algorithm 1 starts by matching seller s_m, the one with lowest willingness to sell, and because $F(s_m, \psi, \tau) \neq \emptyset$, seller s_m will transact under algorithm 1. But because $|S| > |B|$, this implies that there will be at least one seller $s_j \in S$ such that $j < m$, who does not transact, which implies that the final allocation is not Pareto efficient. ∎

When products are vertically differentiated, so that we may have $q_j \neq q_{j'}$ for $j \neq j'$, the platform can implement an algorithm similar to algorithm 1. In this case, instead of ordering sellers from lowest to highest cost, we index sellers such that

$$(1-\psi)q_1 - c_1 \geq (1-\psi)q_2 - c_2 \geq \cdots \geq (1-\psi)q_m - c_m,$$

i.e., such that sellers with lower indexes are the ones who have more feasible matches. Then we implement algorithm 1. If, after implementing this algorithm we have an unmatched seller who could be replaced by a matched seller and generate a higher revenue to the platform, we make the replacement. We keep making those replacements iteratively, until every agent who is unmatched cannot replace a matched agent and generate a higher revenue.

Algorithm 2. *For a given (ψ, τ), index sellers so that*

$$(1-\psi)q_1 - c_1 \geq (1-\psi)q_2 - c_2 \geq \cdots \geq (1-\psi)q_m - c_m,$$

i.e., we order sellers based on how "easy" it is to find a feasible match for that seller, with seller s_1 being the one with most feasible matches, and seller s_m being the one with least feasible matches.

i) Using this indexing for sellers, initialize μ^* as the resulting matching function obtained by performing algorithm 1.

ii) If there is some $s_j \notin S_\mu$ such that

$$(1-\alpha)q_j + \alpha\frac{c_j}{1-\psi} > (1-\alpha)q_{j'} + \alpha\frac{c_{j'}}{1-\psi},$$

[4] This hypothesis is without loss of generality, because the platform would never have incentives to match buyers or sellers who do not have feasible counterpart with whom they would be willing to transact.

for some $j' > j$ such that $s_{j'} \in S_{\mu^*}$, we define

$$s_{\tilde{j}} = \underset{s_{j'} \in S_{\mu^*}}{\arg\min} (1-\psi)q_{j'} - c_{j'}$$
$$s.t. \quad (1-\alpha)q_j + \alpha\frac{c_j}{1-\psi} > (1-\alpha)q_{j'} + \alpha\frac{c_{j'}}{1-\psi},$$

and replace the match $s_{\tilde{j}} : \mu^*(s_{\tilde{j}})$ with $s_j : \mu^*(s_{\tilde{j}})$, and leave seller $s_{\tilde{j}}$ unmatched. Redefine μ^* as the matching function obtained after this exchange and repeat step ii).

Else, stop the algorithm.

It can be shown that the algorithm eventually converges, as an unmatched seller with index $j \in \{1, 2, \cdots, m\}$ can only create a chain of up to $m - j$ replacements. That is because each unmatched seller can only replace a matched seller with a higher index.

Theorem 2. *For a given $\psi \in [0, 1]$ and $\tau \geq 0$, the matching allocation obtained from algorithm 2:*

i) Maximizes the number of transactions that can be generated in the economy, i.e., it maximizes $|S_\mu|$.

ii) Maximizes the profits of the platform, i.e., it solves the maximization problem 7.

5 Optimal commissions and optimal match

In the previous section we assumed the commission fees ψ and τ to be exogenous. But in reality one would expect commissions to be set by the platform along with the matching function μ. In this case, the platform would choose ψ, τ and μ so as to maximize

$$\max_{\mu,\psi,\tau} \tau|S_{\mu,\psi,\tau}| + \sum_{s_j \in S_{\mu,\psi,\tau}} \psi p_j^*(\mu(s_j), \psi, \tau), \tag{8}$$

where

$$S_{\mu,\psi,\tau} \equiv \{s_j \in S; \mu(s_j) = b_i \in B \wedge (1-\psi)(q_j - u_i) \geq c_j + \tau\}$$

corresponds to the set of sellers who end up transacting given the fee ψ and matching function μ, and $p_j^*(\mu(s_j), \psi, \tau)$ is the equilibrium price given by equation 4.

It can be show that a maximum to this problem always exists. Indeed, for a given match μ, we can show that the platform's objective function is upper semicontinuous in $(\psi, \tau) \in [0, 1] \times \mathbb{R}_+$. Moreover, we can limit τ to be less than or equal to a constant $\overline{\tau}$, since, for τ sufficiently high, no transactions are made, so the platform earns zero profits. Because, for a given μ, the objective function is upper semicontinuous, and because we only consider (ψ, τ) defined on the compact set $[0, 1] \times [0, \overline{\tau}]$, it then follows that the problem has a maximum for a given μ, since every upper semicontinuous function defined on a compact set has a maximum (Leininger (1984)). Because there is a finite number of matches, we conclude that a maximum to problem 8 exists.

Proposition 3. *There exists a pair of commission rates (ψ, τ) and a matching function μ that maximizes the profits of the platform.*

Proof: In the Appendix. ■

Let $\mu^*_{\psi,\tau}$ be the matching algorithm obtained by implementing algorithm 2 when the fees are given by (ψ, τ). For a given match μ, let $M_\mu \equiv \{(s_j, b_i) \in S \times B; \mu(s_j) = b_i\}$ (i.e., M_μ represents all the matches associated with μ, excluding agents who are matched with themselves).

Proposition 4. *If $(\mu^*_{\psi,\tau}, \psi, \tau)$ is a solution to the platform's objective function 8, we must have $(1-\psi)(q_j - u_i) = c_j + \tau$ for at least one $(s_j, b_i) \in M_{\mu^*_{\psi,\tau}}$.*

Intuitively, proposition 4 holds because, if all agents who transact were strictly better off transacting vs not transacting, the platform would be able to marginally increase ψ or τ, thus increasing its revenues per transactions, without reducing the number of transactions.

When we assume that the platform's fixed fee per transaction, τ, is set exogenously, we can use theorem 2 and proposition 4 to build an algorithm to find the optimal $\psi \in [0, 1]$ and optimal match $\mu^*_{\psi,\tau}$. The algorithm basically finds all the points in $\psi \in [0, 1]$ such that, after implementing the match $\mu^*_{\psi,\tau}$, at least one of the sellers who transacts under this match is indifferent between transacting and not transacting, as those points will be our only candidates for an optimum. A similar algorithm can be used to find the optimum τ, given an exogenous ψ.

Algorithm 3. *Let $\mu^*_{\psi,\tau}$ be the matching algorithm obtained by implementing algorithm 2 when the fees are given by (ψ, τ). For a given $\tau \in \mathbb{R}_+$, perform the following algorithm.*

i) Initialize $k = 0$ and $\psi = 0$.

ii) Compute $\mu^*_{\psi,\tau}$. If $M_{\mu^*_{\psi,\tau}} \neq \emptyset$, proceed to the next step, else, stop the algorithm.

iii) Set $k = k + 1$. If, for all $(s_j, b_i) \in M_{\mu^*_{\psi,\tau}}$, we have that $(1 - \psi)(\overline{q} - u_i) > c_j + \tau$, then set

$$\psi_k = \min_{(s_j,b_i)\in M_{\mu^*_{\psi,\tau}}} 1 - \frac{c_j + \tau}{\overline{q} - u_i},$$

else, set $\psi_k = \psi$.

iv) If there is no $(s_j, b_i) \in M_{\mu^*_{\psi,\tau}}$ such that $(1 - \psi)(\overline{q} - u_i) > c_j + \tau$, stop the algorithm, else, redefine

$$\psi = \min_{(s_j,b_i)\in M_{\mu^*_{\psi,\tau}}} 1 - \frac{c_j + \tau}{\overline{q} - u_i}$$
$$s.t. \quad (1 - \psi)(\overline{q} - u_i) > c_j + \tau$$

and go back to step iii).

If $k = 0$, then any $\psi \in [0, 1]$ and any match μ generates zero profits to the platform. If $k > 0$, choose the commission

$$\psi^* = \underset{\psi\in\{\psi_1,\psi_2,\cdots,\psi_k\}}{\arg\max} \pi(\mu^*_{\psi,\tau}, \psi, \tau).$$

and the match $\mu^*_{\psi^*,\tau}$.

Theorem 3. *For a given $\tau \in \mathbb{R}$, let ψ^* and $\mu^*_{\psi^*,\tau}$ be the fee and matching function, respectively, obtained after implementing algorithm 3. Then,*

$$(\mu^*_{\psi^*,\tau}, \psi^*) = \underset{\mu,\psi}{\arg\max}\ \tau|S_{\mu,\psi,\tau}| + \sum_{s_j\in S_{\mu,\psi,\tau}} \psi p^*_j(\mu(s_j), \psi, \tau).$$

While algorithm 3 does not simultaneously find the optimal ψ and optimal τ, one could create a grid for $\tau \in [0, \overline{\tau}]$, and then perform algorithm 3 for each τ in the grid, and then select the τ from the grid that generates the maximum revenue to the platform.

References

Anderson, S. P. and R. Renault (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics 30*, 719–735.

Ausubel, L. M., C. Aperjis, and O. Baranov (2017). Market design and the fcc incentive auction.

AUSUBEL, L. M., P. CRAMTON, M. PYCIA, M. ROSTEK, and M. WERETKA (2014). Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies 81*(4 (289)), 1366–1400.

Azevedo, E. M. and E. Budish (2019). Strategy-proofness in the large. *The Review of Economic Studies 86*, 81–116.

Bergemann, D., B. Brooks, and S. Morris (2015, March). The limits of price discrimination. *American Economic Review 105*(3), 921–57.

Binmore, K. and J. Swierzbinski (2000). Treasury auctions: Uniform or discriminatory? *Review of Economic Design 5*, 387–410.

Casner, B. (2020). Seller curation in platforms. *International Journal of Industrial Organization 72*.

Deshmukh, K., A. V. Goldberg, J. D. Hartline, and A. R. Karlin (2002). Truthful and competitive double auctions. In R. Möhring and R. Raman (Eds.), *Algorithms — ESA 2002*, Berlin, Heidelberg, pp. 361–373. Springer Berlin Heidelberg.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory 3*, 156–168.

Eliaz, K. and R. Spiegler (2011). A simple model of search engine pricing. *The Economic Journal 121*, F329–F339.

Haghpanah, N. and R. Siegel (2019, February). Consumer-optimal market segmentation. Working Paper 3333940, Social Science Research Network (SSRN).

Hannak, A., G. Soeller, D. Lazer, A. Mislove, and C. Wilson (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, New York, NY, USA, pp. 305–318. Association for Computing Machinery.

Hao Yang, K. (2020, December). Selling consumer data for profit: Optimal market-segmentation design and its consequences. Working paper, Yale.

Karp, R. M., U. V. Vazirani, and V. V. Vazirani (1990). An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pp. 352–358.

Leininger, W. (1984). A generalization of the 'maximum theorem'. *Economics Letters 15*(3), 309–313.

Mattioli, D. (2012). On orbitz, mac users steered to pricier hotels. *The Wall Street Journal*.

Mehta, A., A. Saberi, U. Vazirani, and V. Vazirani (2007, oct). Adwords and generalized online matching. *J. ACM 54*(5), 22–es.

Satterthwaite, M. A., S. R. Williams, and K. E. Zachariadis (2022). Price discovery using a double auction. *Games and Economic Behavior 131*, 57–83.

Weiss, T. R. (2000). Amazon apologizes for price-testing program that angered customers. *Computerworld*.

Zhao, D., D. Zhang, M. Khan, and L. Perrussel (2011). Maximal matching for double auction. In J. Li (Ed.), *AI 2010: Advances in Artificial Intelligence*, Berlin, Heidelberg, pp. 516–525. Springer Berlin Heidelberg.

A Appendix

Proof of proposition 1:

Necessity: Consider a customer b_i such that $i \leq \min\{n, m\}$ and $u_i \leq q_i - c_i$.

Case 1) Suppose by way of contradiction that customer b_i is matched with a seller s_j such that $u_j > q_j - c_j$. Then there must be a seller s_l such that $l \leq i$, that is either unmatched or is matched with a buyer b_k such $u_k > q_k - c_k$. If seller s_l is unmatched, the allocation is clearly inefficient, as total surplus would be increased if buyer b_i was matched with the idle seller s_l instead of being matched with seller s_j, as the quality of the product from seller s_l is higher (indeed, the value created from the match between seller s_l and customer b_i is $q_l - c_l - u_i$, which is greater than $q_j - c_j - u_i$, the value created from the original match). If, on the other hand, seller s_l is matched with a buyer b_k such that $q_k - c_k < u_k$, the sum of the value created from the matches between customer b_i and seller s_j and customer b_k and seller s_l is given by

$$(q_j - c_j - u_i) + (q_l - c_l - u_k) = (q_l - c_l - u_i) + (q_j - c_j - u_k). \tag{9}$$

If $k > j$, then $q_k > q_j$, which implies that $(q_j - c_j - u_k) < (q_j - c_j - u_j)$. From the assumption that $u_j > q_j - c_j$, this implies that $(q_j - c_j - u_k) < 0$. If $k \leq j$, then $(q_j - c_j - u_k) \leq (q_k - c_k - u_k) < 0$. Either way, we must have $(q_j - c_j - u_k) < 0$, which implies that the value created from these two matches (expression (9)) is less than or equal to

$$q_l - c_l - u_i,$$

the value created if seller s_l is matched to customer b_i, and seller s_j and customer b_k remain unmatched (or matched with one another generating a surplus of zero). Because a new allocation can be found that increases the total surplus, we conclude that the original allocation is Pareto inefficient.

Cbse 2) Suppose by way of contradiction that customer b_i is unmatched. Then there must be a seller s_l with $l \leq i$ that is either unmatched or is matched to a customer b_k such that $q_k - c_k < u_k$. If seller s_l is unmatched, the allocation is clearly inefficient, as total surplus would be increased if customer b_i was matched with the idle seller s_l instead of being unmatched, increasing total surplus from 0 to $q_l - c_l - u_i$ ($q_l - c_l - u_i > 0$ stems from the fact that, since $i \geq l$, we have that $q_l - c_l - u_i > q_i - c_i - u_i > 0$).

Analogously, if seller s_l is matched with a buyer b_k such that $q_k - c_k < u_k$, then, the value created from this match is given by

$$\max\{q_l - c_l - u_k, 0\},$$

which is less than $q_l - c_l - u_i$, the value created if seller s_l was matched with customer b_i instead, and customer b_j remained unmatched.

In either of these contingencies, an alternative allocation can be found that improves the total surplus, which implies that an allocation in which customer b_i remains unmatched cannot be Pareto efficient.

Sufficiency: Let $M \equiv \{i \in \mathbb{N}; i \leq \min\{n, m\} \wedge q_i - c_i > u_i\}$. If each customer $b_i \in M$ is matched with a seller s_j such that $b_j \leq s_j$, then the maximum total surplus that can be achieved from such a match is given by[5]

$$\sum_{i \in M} (q_i - c_i - u_i),$$

which is greater than or equal to the maximum total surplus that can be obtained from any other match. Indeed, for any subset of sellers $\tilde{S} \subseteq S$ and any subset of buyers $B' \subseteq B$, if we have each seller from $\tilde{S}$ **purchasing** from a different buyer from set B', then the total surplus from these transactions would be given by

$$\sum_{j \in \tilde{S}} (q_j - c_j) - \sum_{i \in B'} u_i,$$

which is less than or equal to $\sum_{i \in M} (q_i - c_i - u_i)$.

■

The following proofs will make use of the following definition:

Definition 2. *For a given pair of ψ and τ, we define*

$$\widehat{B_\mu} \equiv \{i \in \{1, 2, \cdots, n\}; \mu(b_i) = s_j \in S \wedge (1 - \psi)(q_j - u_i) \geq c_j\},$$
$$\widehat{S_\mu} \equiv \{j \in \{1, 2, \cdots, m\}; \mu(s_j) = b_i \in B \wedge (1 - \psi)(q_j - u_i) \geq c_j\},$$
$$B_\mu \equiv \{b_i\}_{i \in \widehat{B_\mu}},$$
$$S_\mu \equiv \{s_i\}_{i \in \widehat{S_\mu}},$$

In words, B_μ and S_μ correspond to the set of buyers and sellers, respectively, who end up transacting after the match μ is formed, given that the platform's transaction fee is ψ; whereas $\widehat{B_\mu}$ and $\widehat{S_\mu}$ correspond to the indexes of those buyers and sellers, respectively. Notice that, because the matching is one-to-one, we must have $|B_\mu| = |S_\mu|$ for any matching function μ. Also notice that these sets depend on ψ. But because ψ will be treated as given in the following proofs, we did not index these sets by ψ to avoid clutter notation.

Proof of theorem 1:

Part i)

Let μ^* be the matching obtained by implementing algorithm 1, and let μ be any other matching. Clearly, if $B_{\mu^*} = \emptyset$, then $B_\mu = \emptyset$. So suppose that $|B_{\mu^*}| > 0$.

Suppose by contradiction that $|B_\mu| > |B_{\mu^*}|$ (which happens iff $|S_\mu| > |S_{\mu^*}|$). Then there must be an $i \in \widehat{B_\mu}$ such that $i \notin \widehat{B_{\mu^*}}$. Because μ^* sequentially matches buyers with lowest outside options first (i.e., those with lower indexes), this implies that $i > |B_{\mu^*}|$. If $\mu(b_i) = s_j \notin S_{\mu^*}$, we get a contradiction. Indeed, since $\mu(b_i) = s_j \in S_\mu$, we have that buyer b_i is a *feasible* match to seller s_j. But because $i > |B_{\mu^*}|$, by the time seller s_j was selected in algorithm 1 (i.e., the algorithm used to form the match μ^*) to form a match with one of the remaining customers, customer b_i was still available to that seller, so that $b_i \in F(s_j, B_j, \psi, \tau)$, a contradiction with $\mu^*(s_j) = s_j$.

[5] Customers not in this set can either remain unmatched, or matched a seller $j \notin M$; either way, no transaction would take place for those agents, as the buyer would prefer to evert to his outside option.

Therefore, for each $b_i \in B_\mu$ such that $b_i \notin B_{\mu^*}$, we must have that $\mu(b_i) \in S_{\mu^*}$, i.e., if a buyer transacts under μ but not under μ^*, then this buyer must transact with one of the sellers who transacts under μ^*. Therefore, if there is a seller $s_k \in S_\mu$ such that $s_k \notin S_{\mu^*}$, there exists $b_l \in B_{\mu^*}$ such that $\mu(b_l) = s_k$. Moreover, because $|B_\mu| > |B_{\mu^*}|$, there must be a $b_i \in B_\mu \backslash B_{\mu^*}$ such that and $\mu(b_i) = s_j \in S_{\mu^*}$. But if $\mu(b_i) = s_j$ and $b_i \notin B_{\mu^*}$, then every seller s_k with $k < j$ (i.e., with $c_k \leq c_j$) must be matched under μ^*, as b_i is a feasible buyer to seller s_j (i.e., $b_i \in F(s_j, B, \psi, \tau)$), and thus, feasible to all sellers with a lower cost, and is available to be matched by the time seller s_k with $k < j$ is selected in algorithm 1.

Therefore, if $s_k \notin S_{\mu^*}$, we must have $k > j$ (i.e., $c_k \geq c_j$). Moreover, if $\mu(b_l) = s_k$, then b_l is a feasible match to seller s_k (i.e., $b_l \in F(s_k, B, \psi, \tau)$). But since $s_k \notin S_{\mu^*}$, then by the time s_k is selected to form a match in algorithm 1, buyer b_l was already matched with some other seller $s_{k'}$ such that $k' > k$ (i.e., with cost $c_{k'} \geq c_k$). So there must be some $k' > k$ with $k' \notin S_{\mu^*}$ such that $s_{k'} \notin S_\mu$ and $s_{k'} \in S_{\mu^*}$. So we conclude that for every $s_k \in S_\mu$ such that $s_k \notin S_{\mu^*}$ there is a different $s_{k'} \in S_{\mu^*}$ such that $s_{k'} \notin S_\mu$, so that $|S_{\mu^*}| \geq |S_\mu|$. Because the matching is one-to-one, this implies that $|B_{\mu^*}| \geq |B_\mu|$, a contradiction with $|B_\mu| > |B_{\mu^*}|$.

Part ii)

For pedagogic purposes, we will prove the proposition for the case in which all inequalities in 1 and 2 are strict. The result can be easily extended to cases in which those inequalities are not strict (see footnotes 6 and 7).

Clearly, the profit from the platform as a function of μ is given by

$$\begin{aligned}\pi(\mu) &= \psi \sum_{j \in S_\mu} \left[(1-\alpha)(\overline{q} - u_{\mu(s_j)}) + \alpha \frac{c_j + \tau}{1-\psi} \right] \\ &= |S_\mu| \psi (1-\alpha) \overline{q} - \psi(1-\alpha) \sum_{i \in \widehat{B_\mu}} u_i + \frac{\psi\alpha}{(1-\psi)} \sum_{j \in \widehat{S_\mu}} (c_j + \tau).\end{aligned}$$

So the total profit only depends on the sets B_μ and S_μ.

Now let μ^* be the matching allocation obtained by implementing algorithm 1, and let μ be any other matching allocation.

Clearly, if $B_{\mu^*} = \emptyset$, then we must have $B_\mu = \emptyset$, in which case the proof is trivial. So let us assume that $B_{\mu^*} \neq \emptyset$.

If $S_\mu \subseteq S_{\mu^*}$ and $B_\mu \subseteq B_{\mu^*}$, then μ^* clearly generates a weakly better match to the platform than μ. This happens because the profit that the platform gets by implementing μ only depends on the sets S_μ and B_μ, and because every feasible match that is formed adds a positive profit to the platform. Because $S_\mu \subseteq S_{\mu^*}$ and $B_\mu \subseteq B_{\mu^*}$, this then implies that the platform obtains a higher profit by implementing μ^* as opposed to μ.

So consider the following alternative cases:

Case 1) Suppose there is some $b_i \in B_\mu$ such that $b_i \notin B_{\mu^*}$. Because algorithm 1 iteratively matches buyers with lowest outside options first, it must be the case that there are $|B_{\mu^*}|$ buyers with outside option lower than u_i. Because from theorem 1, $|B_{\mu^*}| \geq |B_\mu|$, there must be at least one $b_j \in B_{\mu^*} \backslash B_\mu$ such that $u_j < u_i$.[6] Then the platform's profits obtained by implementing μ can be improved by matching seller $\mu(b_i)$ with b_j as opposed to matching $\mu(b_i)$ with b_i.

[6] If the inequalities in 1 and 2 were not necessarily strict, then we would suppose that there is some $b_i \in B_\mu$ such that $\#\{j \in B_\mu; u_j = u_i\} > \#\{j \in B_{\mu^*}; u_j = u_i\}$, and conclude that this would imply that there is at least one $j \in B_{\mu^*} \backslash B_\mu$ such that $u_j < u_i$.

Cbse 2) Suppose there is some $s_i \in S_\mu$ such that $s_i \notin S_{\mu^*}$. Because $F(s_j, B, \psi, \tau) \leq F(s_i, B, \psi, \tau)$ for all $j > i$, there must be at least one $j > i$ such that $s_j \in S_{\mu^*} \backslash S_\mu$.[7] Then the platform's profits obtained by implementing μ can be improved by matching buyer $\mu(s_i)$ with s_j as opposed to matching $\mu(s_i)$ with s_i.

■

Proof of proposition 3:

For a given match μ, we first show that the platform's objective function is upper semicontinuous (u.s.c.) in (ψ, τ). Indeed, for each $s_j \in S$ define

$$\pi_{s_j,b_i}(\psi,\tau) \equiv \begin{cases} \tau + \psi\left[(1-\alpha)(q_j - u_i) + \alpha\frac{c_j+\tau}{1-\psi}\right], & \text{if } (1-\psi)(q_j-u_i) \geq c_j + \tau \text{ and } \psi < 1 \\ \max\{q_j - u_i, 0\}, & \text{if } \psi = 1 \text{ and } c_i = \tau = 0 \\ 0, & \text{else} \end{cases}.$$

In this case, defining $M_\mu \equiv \{(s_j, b_i) \in S \times B; \mu(s_j) = b_i\}$ (i.e., M_μ represents all the matches associated with μ, excluding agents who are matched with themselves), we have that the profit from the platform can be written as

$$\pi(\mu, \psi, \tau) = \sum_{(s_j,b_i)\in M_\mu} \pi_{s_j,b_i}(\psi,\tau).$$

Because the sum of u.s.c. functions is u.s.c., it suffices to show that $\pi_{s_j,b_i} : [0,1] \times \mathbb{R}_+ \to \mathbb{R}$ is u.s.c. for every possible $(s_j, b_i) \in S \times B$.

Clearly, for values of (ψ, τ) such that $(1-\psi)(q_j - u_i) > c_j + \tau$ or $(1-\psi)(q_j - u_i) < c_j + \tau$, the function $\pi_{s_j,b_i}(\cdot,\cdot)$ is continuous, and therefore, u.s.c.

So suppose that (ψ, τ) is such that $(1-\psi)(q_j - u_i) = c_j + \tau$. If $q_j - u_i < 0$, the platform earns zero profits regardless of the commissions it charges, in which case $\pi_{s_j,b_i}(\cdot,\cdot)$ is continuous and therefore u.s.c. So throughout, let us assume that $q_j - u_i \geq 0$. Then, we must analyze the two possible scenarios:

1. Suppose that $\psi < 1$. Then there is a $\delta > 0$ such that $\psi + \delta < 1$.

   So suppose that
   $$||(\psi', \tau') - (\psi, \tau)|| = \sqrt{(\psi' - \psi)^2 + (\tau' - \tau)^2} < \delta.$$

   Then $\psi' - \psi < \delta \Rightarrow \psi' < \psi + \delta < 1$.

   So, for any $(\psi', \tau') \in B_\delta((\psi, \tau))$, we have that $\psi' < 1$, which also implies that

$$\begin{aligned}\pi_{s_j,b_i}(\psi',\tau') - \pi_{s_j,b_i}(\psi,\tau) \leq& \tau' + \psi'\left[(1-\alpha)(q_j-u_i) + \alpha\frac{c_j+\tau'}{1-\psi'}\right] - \tau - \psi\left[(1-\alpha)(q_j-u_i) + \alpha\frac{c_j+\tau}{1-\psi}\right] \\ \leq& \underbrace{(\tau'-\tau)}_{<\delta} + \underbrace{(\psi'-\psi)}_{<\delta}(1-\alpha)(q_j-u_i) + \alpha\left[\frac{(1-\psi)\psi'(c_j+\tau') - (1-\psi')\psi(c_j+\tau)}{(1-\psi')(1-\psi)}\right] \\ <& \delta + \delta(1-\alpha)(q_i-u_i) + \alpha\left[\frac{(1-\psi)(\psi+\delta)(c_j+\tau+\delta) - (1-\psi-\delta)\psi(c_j+\tau)}{(1-\psi+\delta)(1-\psi)}\right].\end{aligned}$$

[7] If the inequalities in 1 and 2 were not necessarily strict, then we would suppose that there is some $s_i \in S_\mu$ such that $\#\{j \in S_\mu; c_j = c_i\} > \#\{j \in S_{\mu^*}; c_j = c_i\}$, and conclude that this would imply that there is at least one $j \in S_{\mu^*} \backslash S_\mu$ such that $c_j > c_i$.

The term

$$\delta + \delta(1-\alpha)(q_i - u_i) + \alpha \left[\frac{(1-\psi)(\psi+\delta)(c_j+\tau+\delta) - (1-\psi-\delta)\psi(c_j+\tau)}{(1-\psi+\delta)(1-\psi)} \right]$$

clearly converges to zero as $\delta \to 0$. This implies that, for any $\varepsilon > 0$, there is a δ sufficiently small such that, for any $(\psi', \tau') \in B_\delta((\psi', \tau'))$, $\pi_{s_j,b_i}(\psi', \tau') - \pi_{s_j,b_i}(\psi, \tau) < \varepsilon$, which implies that $\pi_{s_j,b_i}(\cdot,\cdot)$ is u.s.c. in (ψ, τ).

2. Suppose that $\psi = 1$ and $c_j = \tau = 0$. Then, for any $\varepsilon > 0$ and any $(\psi', \tau') \neq (\psi, \tau) = (1, 0)$,

$$\pi_{s_j,b_i}(\psi', \tau') = \begin{cases} \psi'(1-\alpha)(q_j - u_i) + \alpha \frac{\tau'}{1-\psi'}, & \text{if } (1-\psi')(q_j - u_i) \geq c_j + \tau' \text{ and } \psi' < 1 \\ 0, & \text{else} \end{cases}$$

Notice that if $q_j - u_i < 0$, then $\pi_{s_j,b_i}(\psi', \tau') = 0$ for all $(\psi', \tau') \in [0,1] \times \mathbb{R}_+$, in which case $\psi(\cdot,\cdot)$ is continuous, and therefore u.s.c. So suppose that $q_j - u_i \geq 0$. Then, for any $(\psi', \tau') \neq (\psi, \tau) = (1, 0)$ such that $\psi' < 1$, we have that,

(a) If $(1-\psi')(q_j - u_i) \geq \tau'$ and $\psi' < 1$, then

$$\begin{aligned} \pi_{s_j,b_i}(\psi', \tau') - \pi_{s_j,b_i}(\psi, \tau) &= \tau' + \psi' \left[(1-\alpha)(q_j - u_i) + \alpha \frac{\tau'}{1-\psi'} \right] - (q_j - u_i) \\ &\leq -\psi'(q_j - u_i) + \psi' \left[(1-\alpha)(q_j - u_i) + \alpha \frac{(1-\psi')(q_j - u_i)}{1-\psi'} \right] = 0. \end{aligned}$$

(b) If $(1-\psi')(q_j - u_i) < \tau'$ or $\psi' = 1$, then

$$\pi_{s_j,b_i}(\psi', \tau') - \pi_{s_j,b_i}(\psi, \tau) = 0 - (q_j - u_i) < 0.$$

So we conclude that, for any $(\psi', \tau') \neq (\psi, \tau) = (1, 0)$ and any $\varepsilon \geq 0$ we have that $\pi_{s_j,b_i}(\psi', \tau') - \pi_{s_j,b_i}(\psi, \tau) < \varepsilon$, which implies that $\pi_{s_j,b_i}(\cdot,\cdot)$ is u.s.c. in $(\psi, \tau) = (1, 0)$.

Now clearly, if τ is sufficiently large, no transactions are made, which results in zero profits to the platform. So, without loss of generality, we can assume that the platform will only consider choosing $\tau \in [0, \overline{\tau}]$ for some $\overline{\tau} > 0$. In this case, for a given μ, the feasible set from the platform's problem is given by $[0,1] \times [0, \overline{\tau}]$, which is compact. Because any u.s.c. function defined on a compact set has a maximum (Leininger (1984)), we have that for a given μ there exists a (ψ, τ) that maximizes the platform's profits. Because the set of possible matches is finite, this implies that there exists a tuple (μ, ψ, τ) that maximizes the profits of the platform. ■