Buyer-Optimal Platform Design*

Daniele Condorelli
University of Warwick
d.condorelli@gmail.com

Balazs Szentes
London School of Economics
b.szentes@lse.ac.uk

November 26, 2022

Abstract

A platform matches a unit-mass of sellers, each owning a single product of heterogeneous quality, to a unit-mass of buyers with differing valuations for unit-quality. After matching, sellers make take-it-or-leave-it price-offers to buyers. Initially, valuations of buyers are only known to them and the platform, but sellers make inferences from the matching algorithm. The efficient matching is positive-assortative, but buyer-optimal matchings are, often, stochastically negative-assortative (i.e., compared to lower-quality sellers, high-quality ones are matched to buyers with lower expected valuation). Albeit everyone trades when the platform has full-information, generating rents for the side lacking bargaining power results in inefficient matching.

KEYWORDS: Two-sided markets, matching, asymmetric information, platforms.

JEL CLASSIFICATION: D82, D42, D47, C78

*We thank seminar participants at the 2022 Berkeley / Columbia / Duke / MIT / Northwestern IO Theory Conference, the 2021 Warwick Economic Theory Conference, 2021 Piccolo & friends UniBG SKIO conference, Università di Torino, Universitat Autonoma de Barcelona and at the Center for Competition Policy at the University of East Anglia. We also thank Eduardo Azevedo, Andrea Galeotti, Ludovic Renou and Nima Haghpanah for comments that helped us improving the paper.

1 Introduction

Outcomes in two-sided markets are defined by (i) who transacts with whom and (ii) how the surplus from each transaction is divided. In digital marketplaces, a *matchmaker* normally controls (i), but not necessarily (ii). In many cases, especially if transaction costs for agreeing the terms of the exchange are substantial, matched parties bargain bilaterally, with no, or limited, access to alternatives within the platform. For example, Google exposes consumers to a subset of businesses, which often are chosen among thousands of potential ones.[1] An analogous service is offered by Amazon, eBay and Facebook Marketplace, which recommend products to users who have filled a query. Similarly, for popular destinations, Airbnb does not provide on-screen to lodgers all available rental options but a tailored sample. So do dating services, which restrict large pool of suitors for their users.

In this paper we model the *two-sided matching design* problem of a platform with superior information about its users. With its emphasis on post-match bargaining, our approach raises the following issue. Since platforms can use information that they have on their users to optimize matching, agents are able to make inferences from the matching algorithm that may affect post-match bargaining. For example, Google, who is likely able to infer the wealth of its users with great accuracy, may display more prominently high-quality purchase opportunities to consumers that it may know to have higher willingness-to-pay.[2] As a result, it is conceivable that sellers advertising on Google base their prices not just on their product quality and prior-information about consumers, but also on the sample of buyers that they usually interact with. When this feedback effect is present, platforms can not treat the post-match bargaining outcome of any couple as exogenous, which renders optimal matching design more complex.[3]

In our model, a platform one-to-one matches buyers to quality-differentiated sellers of a single good. A buyer's willingness to pay is known to the platform but not to sellers. The quality of a seller is publicly observed. After being matched, sellers make take-it-or-leave-it price offers to buyers. The matching (indirectly) leaks information, either because it is public or as a result of equilibrium reasoning, which sellers can use to tailor their subsequent price offers and, collectively, price-discriminate against buyers. The main insight we obtain follows from characterising both the welfare optimal matching and the buyer-optimal one. We find that a matching that is meant to favour the side lacking bargaining power, buyers in our case, will often introduce substantial sorting inefficiencies compared to the first-best. However, with a fully informed platform, buyers and sellers trade efficiently post-match.

By assuming that the matching is one-to-one we abstract away from the platform potentially exploiting price-competition.[4] We believe this is often a good simplifying approximation. First, even if consumers are exposed to multiple firms, the top-ranked enjoys substantial market power, as many consumers may be reluctant to search. It is known that consumers are heavily biased toward most prominently located op-

[1] Varian (2006) writes: "First, what does Google do? The answer, I claim is that Google is a yenta — a traditional Yiddish word for matchmaker. [...] From an economics perspective, Google runs a two sided matching mechanism."

[2] Businesses bid for impressions based on information provided by Google and high-quality business likely bid more for high-value consumers. A documented case is that of Orbitz, an online travel agent that used to show to Mac users more expensive hotels than those it showed to Windows users. See Dana Mattioli, "On Orbitz, Mac Users Steered to Pricier Hotels", WSJ (2012).

[3] This differentiates us from most, if not all, of the literature on two-sided matching. Consider the paradigmatic example of the National Residency Matching Program (NRMP) studied in Roth and Peranson (1999). Since salaries are set in advance there is little scope for the platform to reveal information on the preferences of both sides that would affect the outcome post-match.

[4] See Elliott et al. (2021) and Bergemann and Bonatti (2022) for models where consumers have preferences for the products of differentiated firms and a platform decides which firm approach which consumer and what information firms receive about buyers (in the former paper) or buyers about products (in the latter). In contrast with our work, in these papers each seller can in principle supply all buyers and *buyer* optimality does not imply sorting inefficiency.

portunities, such as the Buy-box placement in Amazon or a top place in Google's ranking (see Narayanan and Kalyanam (2015)), placements for which firms are willing to pay higher prices. Second, suppose the platform could minimise such prominence by creating a ranking of all firms and consumers could continue their search beyond their first match at a small cost, as it is sometimes the case. Search might still be limited in equilibrium and the first firm could enjoy market power, as illustrated by Diamond (1971).

Let's now elaborate briefly on our main result. The welfare optimal matching is positive-assortative in value and quality. To see this, note two things. First, any deterministic matching induces complete information and results in efficient trade for all buyer-seller couples, albeit sellers extract all the surplus from the transactions. Second, a deterministic positive-assortative matching maximises efficiency of sorting, as the total surplus generated by each match reflects a complementarity between value and quality. Instead, the *buyer-optimal* matching is distorted and, sometimes, (stochastically) negative-assortative. That is, the expected valuation of buyers matched to lower quality sellers is higher than that of buyers matched to higher quality ones. Intuitively, by concealing information about valuations and thereby manipulating sellers'beliefs, such a random matching keeps prices lower and generates rents for buyers. A matching platform wishing to maximise buyer-surplus therefore faces a trade-off between the sorting efficiency of the matching and information rents for the side without bargaining power. This trade-off often resolves in favour of heavily distorting the matching away from positive assortative. Remarkably, matched agents still trade efficiently, but not under complete information.

Our focus on buyer-optimal matching has, primarily, a normative motivation.[5] The view that big-tech (gate-keeping) platforms, such as Google and Amazon, should be regulated is currently popular. In particular, concerns have been raised about what we may call *match-discrimination*, that is the practice of exploiting information on users to determine who they will be able to interact with.[6] We contribute to the debate by offering a framework where some of the welfare trade-offs of designing a "matching algorithm" are exposed. In particular, our results can be read as expressing caution toward tampering with the algorithm in an attempt to increase the surplus of the side with less bargaining power. Generating information rent is expensive. Indeed, we show that in some cases the buyer-optimal matching generates less welfare than a fully random matching, which is chosen by platform that does not match-discriminate.

An implication of our observations is that regulators interested in raising surplus of buyers in buyer-seller platforms with post-match bargaining should explore interventions that increase their bargaining power post-match. Letting buyers make proposals to sellers is, indeed, an existing business model in two-sided platform markets. Such a policy has been popularized by Priceline, an online travel agent, and since then has been widely adopted, quite possibly in an attempt to increase buyers' surplus (e.g., EBay now allows buyers to make offers to consenting sellers).

We now proceed with presenting the model and our results. Formal proofs are in Appendix A. Then, we outline the related literature in Section 4. The model assumes that the platform is fully informed, sellers are only vertically differentiated and buyers have only two possible valuations. The concluding section and the complementary Appendix B discuss how relaxing these assumptions would not affect our insight.

[5]There are several plausible scenarios where maximization of buyer-surplus is a good proxy for the incentives of a two-sided platform. For example, two platforms might be in direct competition to attract buyers, because sellers face switching costs or have already sunk investments that lock them in one or the other. Also, a monopolist platform may be constrained regarding its ability to charge one side, which may lead the platform to maximize the surplus of the side that can be charged; or a platform might earn from advertising to one side only and therefore attempt to maximize participation on that side.

[6]The current chair of the US FTC has expressed this opinion in her landmark paper on Amazon, Khan (2016). In Europe, the Digital Markets Act imposes specific rules of behavior on systemically important platforms, deemed "gatekeepers". Among such rules is one that forbids platforms from distorting their algorithms in favor of own products, a practice called "self-preferencing".

2 The Model

There is a unit mass of buyers. Each buyer has either low valuation, $l\,(>0)$, or high valuation, $h\,(>l)$.The fraction of buyers with high valuation is μ. There is also a unit mass of sellers. Each seller has a single good to sell. The quality of a seller is q and distributed according to the atomless CDF F with support in $[\underline{q},\overline{q}]$. If a buyer with valuation $v\,(\in\{l,h\})$ purchases the good from a seller with quality q at price t, the buyer's payoff is $vq-t$ and the seller's payoff is t. The buyers and sellers are matched by a platform. We assume that the platform and the buyer know the buyer's valuation but the seller does not. We assume that sellers'qualities are publicly observed. Once a buyer and a seller are matched, the seller makes a take-it-or-leave-it price-offer to the buyer. If the buyer accepts the seller's offer, they trade at the price set by the seller. Otherwise, each of them gets their reservation payoff of zero.

Matching.— We describe a *matching* by the probabilities that each seller q is matched with a buyer with valuations h or l. That is, a matching is given by measurable $p=(p_h,p_l)$, $p_h,p_l:\left[\underline{q},\overline{q}\right]\rightarrow[0,1]$, where $p_h\,(q)$ and $p_l\,(q)$ denotes the probabilities that a q-seller is matched with a buyer with valuations h and l, respectively. A *feasible* matching p must satisfy the following constraints:

$$\begin{aligned} p_h\,(q)+p_l\,(q) &\leq 1, \\ \int_{\underline{q}}^{\overline{q}} p_h\,(q)\,dF\,(q) &\leq \mu, \\ \int_{\underline{q}}^{\overline{q}} p_l\,(q)\,dF\,(q) &\leq 1-\mu. \end{aligned}$$

The first constraint guarantees that the probability that a seller with type q is matched with a buyer is weakly less than one. We do not require that each buyer and seller is matched with probability one. The second and third constraints guarantee the measure of high-value (low-value) buyers who are matched with sellers does not exceed the total measure of high-value (low-value) buyers.

We say that a matching p is *positive-assortative* if $p_h(q)=1$ for $q\geq F^{-1}(1-\mu)$, $p_h(q)=0$ elsewhere and $p_l(q)=1-p_h(q)$, where F^{-1} stands for the inverse of F. We say that the matching is *stochastically negative-assortative* whenever p_h is monotonically non-increasing in quality. Conversely, the matching is *stochastically positive-assortative* when p_h is monotonically non-decreasing. A matching is *fully-random* if and only if $p_h(q)=\mu$ and $p_l(q)=1-\mu$ for all $q\in\left[\underline{q},\overline{q}\right]$.

Optimal Prices.— Because the platform is not a player in our game, we treat the matching as an exogenous and known object. Then, if the matching is given by $p=(p_h,p_l)$, the posterior probability of seller with quality q that she is matched with a high-value buyer is $\mu^p\,(q)=p_h\,(q)\,/\,(p_h\,(q)+p_l\,(q))$. Note that $\mu^p=p_h$ if all sellers are always matched. So, the q-quality seller is willing to set price qh if, and only, if $\mu^p\,(q)\geq l/h$. Such seller is willing to set price ql if, and only if, $\mu^p\,(q)\leq l/h$. It is without loss of generality to assume that, when a seller is indifferent between prices, it charges the lowest.

Buyer Surplus.— Let $\chi^p(q)\in[0,1]$ denote the probability that a seller with value q charges price lq following matching p. That is

$$\chi^p\,(q)=\begin{cases} 1 & \text{if } \mu^p\,(q)\leq l/h, \\ 0 & \text{if } \mu^p\,(q)>l/h. \end{cases}$$

Then, for given χ^p, the buyers' surplus can be expressed as

$$\int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, q\, (h-l)\, dF(q)\,.$$

Discussion of assumptions.— We motivated in the introduction the hypothesis that matching is one-to-one. We will comment on the binary-value assumption and on the possibility of having horizontally differentiated sellers in the conclusions and in Appendix B. In the conclusive section we also illustrate the implications of the platform not knowing buyers'values but receiving binary signals about them. Further along the main text we elaborate on the case where quality is not known to buyers. In the rest the section we discuss three other assumptions.

One, we assume that, gross of the price, the utility of a buyer with value v or purchasing a product of quality q is qv. Our results do not qualitatively rely on this specific functional form. The characterization of the welfare-optimal and buyer-optimal matching and the welfare analysis can be extended with little modification to a general $u(q,v)$ assuming u is increasing in both arguments and log-supermodular. We have retained a simpler functional form to avoid burdening the reader with further notation. Two, likewise, assuming that seller had an opportunity cost cq of selling a product of quality q would not affect results.

Three, we are *not* assuming that the platform is able to ask sellers to commit to prices beforehand and then condition the matching on prices. Beyond post-match bargaining being a valid assumption in many cases of interest, for example for Google search, such an assumption would trivialize the problem. A powerful and informed platform would be able to implement any pricing by threatening firms with the prospect of remaining unmatched. A substantial change in modeling would be required to make the optimal matching problem interesting. For instance, see Gomes and Pavan (2016) discussed in Section 4.

3 Optimal Matching

The end goal of this section is to characterize the matching which maximizes buyers'surplus and to study its welfare properties. The primary benchmark against which the buyer-optimal matching will be evaluated is the matching that maximizes welfare, intended as the sum of expected buyers'surplus and sellers'profit. We therefore start with the following result, which should not be surprising in light of Becker (1975).

Proposition 1. *A matching maximizes total welfare if and only if it is a positive-assortative matching (PAM) almost-everywhere. In the PAM buyers obtain zero surplus.*

A formal proof of Proposition 1 relies on two observations. First, for any matching p that induces *complete-information*, that is $\mu^p(q) \in \{0,1\}$, trade will take place with probability one and sellers will obtain all surplus. Second, PAM induces complete information and the total welfare of a match between a seller with quality q and a buyer with value v is given by the supermodular function qv.

Remarkably, the welfare optimal matching not only maximizes expected welfare, but induces a first-best outcome in the classic sense. Since buyers obtain zero surplus, it is immediate to observe that PAM is also the matching that maximizes the profits of sellers.

Having noted that a buyer-surplus maximizing matching must generate either zero surplus or inefficiencies, we now fully characterize the buyer-optimal matching and pricing and show that such inefficiencies can be sizable. We focus on (ex-ante) *Pareto-efficient buyer-optimal* matchings, that is, there is no other matching that gives higher sellers' surplus without reducing buyer-surplus. This has one main implication: no buyer or seller remains unmatched, even if additional matches do not increase buyer surplus.

Theorem 1. *Let $p^* = (p_h^*, p_l^*)$ be a matching defined as follows:*

$$p_h^*(q) = \begin{cases} l/h & \text{if } q \geq q^* \\ 0 & \text{if } q \leq q^* \text{and } \mu \leq l/h \\ 1 & \text{if } q \leq q^* \text{and } \mu \geq l/h \end{cases},$$

$$p_l^*(q) = 1 - p_h^*(q),$$

$$\text{where } q^* = \begin{cases} F^{-1}\left(\frac{l-\mu h}{l}\right) & \text{if } \mu \leq l/h, \\ F^{-1}\left(\frac{\mu h - l}{h-l}\right) & \text{if } \mu \geq l/h. \end{cases}$$

Any Pareto-efficient matching that maximizes buyer-surplus is equal to p^* almost everywhere.

The efficient buyer-optimal matching (henceforth also BOM) of high-value buyers, p_h^*, is exemplified in the two panels below. Observe that, when $\mu \geq l/h$ the matching is stochastically negative assortative. Instead, when $\mu < l/h$ then the matching is stochastically positive assortative.

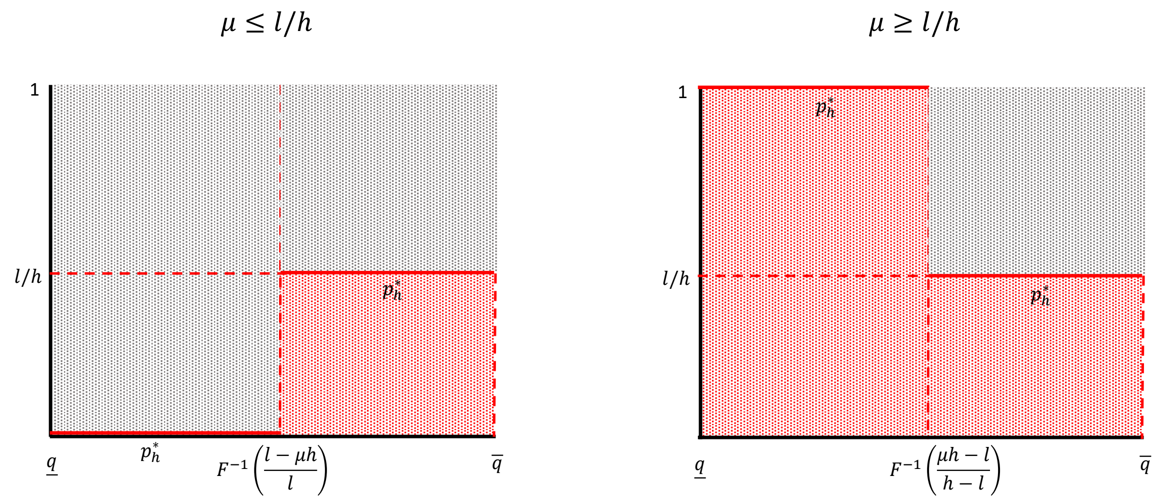

Figure 1: Sketch of Efficient Buyer-Optimal matching, p_h^*.
The red-dotted area has measure μ and the gray-dotted area has measure $1-\mu$.

Let us explain the arguments leading to this proposition. Recall that a q-seller sets price qh if the probability of being matched with a high-value buyer exceeds l/h and sets price ql otherwise. So, a buyer's payoff is positive only if his valuation is high and the seller sets the lower of these prices. The surplus of a high-value buyer who purchases a q-quality good at price ql is $q(h-l)$. Therefore, in order for a seller to generate positive consumer surplus, she must be matched with a mixture of high- and low-value buyers. Furthermore, consumer surplus is maximized at a given seller if she is matched with as many high-value buyers as possible as long as she is willing to set the lower price. That is, she is indifferent between the two prices, so the fraction of high-value buyers among her matches is exactly l/h. Since the high-value buyers' surplus, $q(h-l)$, is increasing in q, the BOM generates consumer surplus at high quality sellers. Above a quality threshold, q^*, q-sellers are matched so that they are indifferent between the two prices, so $p_h^*(q) = 1 - p_l^*(q) = l/h$. Of course, to maximize consumer surplus, the threshold q^* should be as low as possible and it is determined by the initial distribution of the buyers, μ. If low-value buyers are abundant, $\mu \leq l/h$, then q^* is defined so that the fraction of high-value buyers who are matched with sellers with quality above q^* is exactly μ. In this case, sellers with quality below q^* are matched with the remaining low-value buyers. If there are few low-value buyers, $\mu \geq l/h$, then q^* is defined so that the fraction of low-value buyers who are matched with sellers with quality above q^* is exactly $1-\mu$. In this case, sellers below q^* are matched with the remaining high-quality buyers. The matching is stochastically negative assortative.

Before proceeding we make a few observations on the nature of the optimal matching. First, buyers and sellers of the same type may be treated differently *ex-post*. For instance, if $\mu > l/h$, only some of the high-value buyers obtain a surplus, others don't. Hence, there *might* be motives for buyers and sellers to refuse to negotiate and reject their match upon learning the type of the agent they are matched with, if such rejection was allowed by our platform *and* followed by rematching. Second, the profit obtained by sellers in the BOM is not necessarily monotone in their quality. Low-quality sellers may end up with more profit than high-quality ones when BOM is (stochastically) negative assortative. Third, even if the BOM is stochastically negative assortative in the traits that contribute more to produce (first-best) joint surplus, the average surplus of buyers matched to high quality sellers is larger than that of buyers matched to lower-quality ones. Fourth, when low-value buyers are rare, that is $\mu > l/h$, the platform could also benefit buyers by re-balancing participation of high and low values in a somewhat counterintuitive way, that is by raising the share of low-value buyers.[7]

Welfare analysis. We conclude this section by comparing BOM's payoffs with those arising under two benchmarks, the PAM and the fully random matching (henceforth FRM). PAM corresponds to the matching that maximizes welfare and profit, as Proposition 1 states. FRM corresponds to the focal case in which the platform does not condition the matching on any information about buyers.

It is obvious that BOM generates strictly larger buyer surplus than PAM and FRM. It is also immediate to see that buyer surplus is strictly larger in the FRM than in PAM, as long as $\mu \leq l/h$ (i.e., all sellers set the lower price when they have no additional information on buyers). Regarding the remaining welfare variables, we have the following proposition. The main insight from this welfare comparison is that, in some cases, BOM generates lower welfare than if the platform ignored information altogether.

Proposition 2. *(i) PAM generates higher welfare and higher profit than BOM; (ii) If $\mu \leq l/h$, then FRM generates the same profit as BOM and, hence, lower welfare. If $\mu > l/h$ then FRM generates higher-profit than BOM but the welfare ranking between the two matchings is, in general, ambiguous; when μ is sufficiently large, that is $\mu \geq \frac{l}{h} + \frac{h-l}{h}\frac{l}{h}$, then the welfare of BOM is larger than that of FMR.*

Part (i) is an immediate corollary of Proposition 1. It is also immediate to see that when $\mu \leq l/h$ profit of BOM and FRM are the same, as in both cases each seller sets price ql and trades takes place. Hence, because BOM generates higher buyer-surplus by Theorem 1, then the first part of (ii) follows. The proof in the Appendix shows that, when $\mu > l/h$, profits are lower under BOM than under FRM. Because setting the high price maximizes profit for all sellers in both the FRM (since $\mu^p = p_h = \mu > l/h$) and BOM (due

[7]The phenomenon that the marginal value of having additional types of buyers is not only equal to the surplus they generate from trading, which is zero for low-value buyers, is studied in depth in Galperti et al. (2022).

to the indifference condition for $q > q^*$), this inequality in profit is the result of the negative assortative nature of BOM, which creates more sorting distortions than a FRM. Then the proof shows, by means of two non knife-edge examples, that the welfare comparison between BOM and FRM is ambiguous.

Since a platform without information on at least one of the two sides implements a FRM, these findings suggest that the overall welfare effect of the platform collecting more information about users may depend on the platform's objective and is non-obvious. As we discuss further in the concluding section, more information always reduces bargaining disagreement but may bring about a larger sorting distortion.

Private qualities. The assumption that qualities of sellers are observable is often plausible, but it is worth asking whether our characterization and main insight would survive if quality was not observable by buyers. We now argue that this is the case. However, while analogous sorting inefficiencies remain, we observe that the platform can achieve a higher payoff for buyers in some equilibria.

It is easy to see that, even with unknown quality, the platform can implement the same buyer-surplus as in the BOM. For instance, assume $\mu > l/h$ and that sellers in $(q^*, \overline{q}]$ are matched to a mix of high and low-vale buyers in such a way that each such seller is indifferent and each buyer's posterior is such that the quality of sellers they are matched with is distributed in $(q^*, \overline{q}]$ according to the truncated prior. Also assume that the remaining high-value buyers are matched *one-to-one* to sellers in $[\underline{q}, q^*]$. Then, there is an equilibrium in which all sellers with q in $[\underline{q}, q^*]$ charge qh, since buyers know the sellers they are matched with, while all sellers with q in $(q^*, \overline{q}]$ charge $\mathbb{E}[q \mid q > q^*]\, l$. In this equilibrium, sellers in $(q^*, \overline{q}]$ deviating to a different price are believed to be of quality q^*. It is immediate to see that buyer surplus in this equilibrium is equal to that achieved by BOM with observable quality. An analogous construction can be performed for $\mu < l/h$.

However, there are other equilibria of the matching above in which buyers do better. In particular, there is an equilibrium where all sellers in $[q^*, \overline{q}]$ post price $q^* l$ sustained by the same out-of-equilibrium beliefs that a deviating seller is of type q^*. The fact that the set of equilibrium prices will depend on the beliefs of buyers, suggests that the buyer-optimal matching might might induce posterior beliefs for *sellers* that are different than those induced by our BOM with known quality. Remarkably, this is not the case. Roughly speaking, by placing sellers in a neighborhood of $\underline{q}$ in the support of the two segments of buyers we built for our Theorem 1, we can construct a matching and an equilibrium such that all sellers in $(\underline{q}, q^*]$ charge $\underline{q}h$ and sellers in $(q^*, \overline{q}]$ charge $\underline{q}l$. Note that the share of sellers charging a price that results in a purchase by high-value buyers at a price acceptable also by the low-value buyers is maximized by posteriors induced by the BOM. Hence, this is the buyer-optimal equilibrium. We omit the tedious details involved in formalizing such a construction.

4 Literature Review

Starting with Shapley and Shubik (1971), an important strand of the literature on two-sided markets has imposed joint restrictions on who matches with whom and at what prices by requiring that no coalition of agents benefits from a different matching and sharing of output that they can implement.[8] We depart from tradition by considering an environment where, due to the informational spillovers and bargaining

[8]A complementary literature has explored conditions for decentralizing such efficient outcomes by means of dynamic matching and bargaining games, where buyers and sellers meet repeatedly according to some *exogenous* random matching process. See Lauermann (2013) for a recent theoretical synthesis of this approach.

under asymmetric information, the ex-post surplus-sharing of a matched couple is not exogenous, but depends on the matching. A notable exception, within the smaller literature that considers matching under asymmetric information, is Liu (2020). There, in order to evaluate potential deviations, agents form interim expectations on the value of their match which, as in our paper, depend on the putative mapping from states of the world into matchings, which is publicly known.

A number of other papers have looked at the incentives of platforms who charge per-click to distort matching in order to boost costly search (i.e., clicking). In Eliaz and Spiegler (2011) a consumer searches from a pool of firms whose boundary is determined by the platform. Inefficiently too many low-quality sellers may be allowed in the pool, to induce consumers to search more. In De Corniere (2016) a platform can match consumers to their preferred segment of firms in a more or less noisy way. Consumers search within the set of firms they are matched with. Because perfect matching traps consumers into monopoly prices, a platform may want to bias the algorithm to foster consumer participation. In a similar vein, in Hagiu and Jullien (2011) an intermediary may match consumers with stores that are worse for them in order to persuade them to search more. They also show that the intermediary gains from marginally biasing the matching away from the perfect one, if it reduce firms' prices and convinces more consumers to visit at least one store. While focusing strictly on platform's incentives, these papers suggest reasons why perfect matching may not be consumer-optimal, with De Corniere (2016) and Hagiu and Jullien (2011) also identifying a feedback effect of matching on prices. In our work, we abstract from horizontal differentiation and competition. The additional simplicity allows us to fully characterise the buyer-optimal matching, thus highlighting a trade-off between matching efficiency and consumers'information rent.

The design of an optimal matching, but by a revenue maximising platform, is studied in some relatively recent papers, including Damiano and Li (2007), Johnson (2013), Gomes and Pavan (2016) and Gomes and Pavan (2022). We share with these works the presence of a monopolistic platform that matches two-sides of the market and the emphasis placed on the distortion introduced by a matching that maximises an objective other than total welfare. In stark contrast to our model, the platform in these papers is uninformed about valuations and aims at maximizing its own profit by setting prices to agents on both sides. The approach is in the spirit of optimal mechanism design, where the outcome is now given by the matching rather than by the allocation of an an object as in classic Myersonian mechanism design. We believe these papers are complementary to ours, as they study optimal matching in a different environment where the platform also restricts the terms of each trade.

In search for a buyer-optimal matching, a designer resolves a trade-off between a more efficient matching and information rent for buyers. This trade off, between efficiency and rent, recurs in other contexts that also share with us a flexible information-design-like approach. In Condorelli and Szentes (2020) a buyer can choose her distribution of value for the product of a seller with bargaining power. Therefore, it faces a related trade-off between having a higher valuation and larger information rent. In Armstrong and Zhou (2021), perfectly informing consumers about which of two differentiated products is best for them relaxes competition but maximizes welfare, while the consumer-optimal information structure dampens differentiation to some extent. In a search model, Dogan and Hu (2021) show that total welfare would be maximized by giving a buyer as much information as possible to find a good match among several firms, but that would lead to too little competition within firms. In our own papers, information is a byproduct of choices (i.e., distribution of value and matching) that affect both the information structure and feasible surplus, while in Armstrong and Zhou (2021) and in Dogan and Hu (2021) the information structure is a primitive that affects agents'choices, which may end up being inefficient.

Intuitively, one would expect the above trade-off between efficient matching and information rents to be exacerbated when sellers are more heterogenous, as in that case it becomes more costly to move away from PAM. Indeed, if all sellers are the same, sorting efficiency is irrelevant and a platform that maximises buyer surplus can focus on producing information rent. In fact, it is not difficult to see that in this case our problem is isomorphic to that of identifying the segmentation of buyer's demand that maximizes buyer-surplus under a single price-discriminating monopolist. It then follows from Bergemann et al. (2015) (BBM) that there exists an efficient matching where buyers obtain all surplus minus the profit from random matching. As it turns out, in our BOM the buyer market is subdivided into two segments, one composed entirely by either high or low types, and the other by a mix of high and low type such that a monopolist facing such segment would be indifferent between setting the high or low price. Remarkably, these are the same two segments composing the segmentation that maximises consumer surplus in BBM.

5 Concluding Remarks

We close the paper by discussing some results for the case where buyers may have more than two values and further elaborating on the relation of our work with BBM. We also discuss extending the analysis to the case of a partially informed platform and that of horizontally differentiated sellers. The take-away is that, in all cases, sorting inefficiency remains a feature of the buyer-optimal matching.

Multiple values. Our key characterisation result for the multiple-value case is formally stated in Appendix B. Any Pareto efficient matching outcome can be obtained starting from a segmentation of the buyer market that is *extremal*, in the sense of BBM.[9] Assuming the designer maximises some weighted average of buyer and seller-surplus, there will be an extremal segmentation such that an optimal matching can be built by ordering segments in terms of the average weighted surplus they generate for *unit-quality* and then positively assortatively matching them to groups of sellers. Buyers in each segment are matched to sellers so that all sellers matched to a certain segment believe they are facing a random buyer from that segment. Trade always takes place because sellers tailor the price to the lowest value in the support of the segment to which they are matched. Any inefficiency is due to sorting.

We make three further observations regarding optimal matching, which we substantiate in Appendix B. First, PAM will continue to maximise welfare and sellers' profits even when buyers have multiple values. The relevant extremal segmentation will be one where each segment contains all and only buyers with the same value. Second, in the extremal segmentation that maximises buyer-surplus, there is at most a single segment for which buyers have only one possible value. If two such segments existed, buyer surplus could be raised by merging them. This implies that sorting will be inefficient. Third, the buyer optimal matching can still be stochastically negative assortative.

Last but not least, in light of our discussion in this and the previous section, it could be conjectured that the buyer-optimal matching problem may be solved using consumer-surplus maximising extremal segmentations characterised by BBM. Unfortunately, this is not always true and obtaining a full solution remains an open question. As we show in Appendix B with an example, the segmentation of buyers in our buyer-optimal matching with heterogenous quality and more than two values may deliver higher than BBM's monopoly profit under fixed unit-quality. The idea underlying the example is that when only a

[9] Extremal markets are those where the seller is indifferent among any price equal to valuations in the support of the market and an extremal segmentation is simply a segmentation of the initial buyer market made up only of extremal markets.

small fraction of sellers is able to generate value, the platform maximising buyer surplus may want to focus on extremal segmentations that, albeit resulting in higher than monopoly profit and, therefore, lower total consumer surplus, nonetheless produce higher consumer surplus for some small fraction of buyers, which will be those matched with the few high-quality sellers.

Partially informed platform. Suppose the platform does not know buyers'values. Instead, it receives independent binary signals about them, which are informative in the sense that a monopolist seller sets the low price following one realization and the high one following the other. It is easily seen that, on the one hand, the matching that maximizes sellers' revenues is positive assortative in the binary signal and sellers' quality. On the other hand, in the buyer-optimal matching a portion of higher quality sellers are matched with a mixture of low-signal and high-signal buyers, so that they are, as in the full information case, indifferent between charging the low and the high prices. Therefore, if $\mu > l/h$ BOM is also stochastically negative assortative. However, compared to the case of full information, sorting distortions when $\mu > l/h$ ($\mu < l/h$) are mitigated (amplified) by the lower leverage that a partially informed platform has to manipulate sellers' beliefs. A further important difference with the full information case is that PAM is not always welfare optimal. Indeed, sometimes BOM is optimal. This is because, for lower levels of signal informativeness, charging a high price results in a substantial likelihood of even the high-signal buyer refusing the offer. Hence, persuading sellers to charge a low price becomes welfare relevant and an additional trade-off arises, which either resolves in favor of BOM or PAM.

Vertically differentiated sellers. We conclude this section with an example suggesting vertical differentiation is not key to observing inefficient sorting in buyer-optimal matchings. Consider a model with two buyers, 1 and 2, and two sellers, A and B. Assume that, with some probability, buyer 1 has value h for the product of A and value l for the product of B, while buyer 2 has value h for the product of B and l for that A. With the remaining probability, preferences are reversed. Given perfect correlation in values, an informed platform can always match each product to the buyer that values it the most. However, such a matching fully informs both sellers that they are facing a value h buyer, thus leaving no surplus to buyers. Instead, by mismatching sufficiently often, that is matching sellers to buyers that like them least until both sellers become indifferent between asking h or l, the platform can make sure that both sellers charge price l, thus raising buyer surplus.

Appendix A: Proofs

Proof of Proposition 1. Denote with G the distribution of buyers'values. Since vq is supermodular, a classic result implies that $\sup_{\pi\in\mathscr{M}(F,G)}\mathbb{E}_\pi[vq]$, where $\mathscr{M}(F,G)$ is a *coupling* of probabilities F and G, has a unique *comonotone* solution which is given by PAM. We have already argued the second statement follows from the fact that PAM induces complete information. $\square$

Proof of Theorem 1. Focusing on the case $\mu \geq l/h$, we show that p^* generates strictly larger consumer surplus than p unless $p = p^*$ almost everywhere. The case $\mu \leq l/h$ is analogous and we omit the proof.

For each p, let us defined the CDF G^p as follows:

$$G^p(x) = \frac{\int_{\underline{q}}^{x} \chi^p(q)\, p_h(q)\, dF(q)}{\int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, dF(q)}.$$

Also, define q^p by

$$\int_{q^p}^{\overline{q}} \frac{l}{h} dF(q) = \int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, dF(q). \tag{1}$$

Finally, define the CDF H^p by $H^p(x) = 0$ if $x \leq q^p$ and by

$$H^p(x) = \frac{\int_{q^p}^{x} \frac{l}{h} dF(q)}{\int_{q^p}^{\overline{q}} \frac{l}{h} dF(q)}$$

if $x > q^p$.

We now show that H^p first-order stochastically dominates G^p. To see this, first note that if $x \leq q^p$ then $H^p(x) = 0 \leq G^p(x)$. Moreover, for all $x > q^p$,

$$1 - H^p(x) = \frac{\int_{x}^{\overline{q}} \frac{l}{h} dF(q)}{\int_{q^p}^{\overline{q}} \frac{l}{h} dF(q)} = \frac{\int_{x}^{\overline{q}} \frac{l}{h} dF(q)}{\int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, dF(q)} \geq \frac{\int_{x}^{\overline{q}} \chi^p(q)\, p_h(q)\, dF(q)}{\int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, dF(q)} = 1 - G^p(x),$$

where the first and last equalities are the definitions of H^p and G^p, respectively, the second equality follows from (1) and the inequality follows from $\chi^p(q) \leq 1$ and $p_h(q) \leq l/h$ whenever $\chi^p(q) = 1$. Then previous inequality chain implies that $H^p(x) \leq G^p(x)$ even when $x > q^p$.

Therefore,

$$\begin{aligned}
&\int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, q\,(h-l)\, dF(q) = \left[\int_{\underline{q}}^{\overline{q}} \chi^p(q)\, p_h(q)\, dF(q)\right] \left[\int_{\underline{q}}^{\overline{q}} q\,(h-l)\, dG^p(q)\right] \\
= &\left[\int_{q^p}^{\overline{q}} \frac{l}{h} dF(q)\right] \left[\int_{\underline{q}}^{\overline{q}} q\,(h-l)\, dG^p(q)\right] \leq \left[\int_{q^p}^{\overline{q}} \frac{l}{h} dF(q)\right] \left[\int_{\underline{q}}^{\overline{q}} q\,(h-l)\, dH^p(q)\right] \\
= &\int_{q^p}^{\overline{q}} \frac{l}{h} q\,(h-l)\, dF(q),
\end{aligned}$$

where the first and last equalities follows from the definitions of G^p and H^p, respectively. The second equality is implied by (1) and the inequality follows from the fact that H^p first-order stochastically dominates G^p.

It remains to show that

$$\int_{q^p}^{\overline{q}} \frac{l}{h} q\,(h-l)\,dF\,(q) \leq \int_{q^*}^{\overline{q}} \frac{l}{h} q\,(h-l)\,dF\,(q)\,.$$

In order to do so, it is enough to argue that $q^p \geq q^*$. Observe that

$$\begin{aligned} \int_{q^*}^{\overline{q}} \left(1-\frac{l}{h}\right) dF\,(q) = 1-\mu \geq \int_{\underline{q}}^{\overline{q}} \chi^p\,(q)\,p_l\,(q)\,dF\,(q) \geq \frac{h-l}{l} \int_{\underline{q}}^{\overline{q}} \chi^p\,(q)\,p_h\,(q)\,dF\,(q) \\ = \quad \frac{h-l}{l} \int_{q^p}^{\overline{q}} \frac{l}{h} dF\,(q) = \int_{q^p}^{\overline{q}} \left(1-\frac{l}{h}\right) dF\,(q)\,, \end{aligned}$$

where the first equality is the explicit definition of q^* and the first inequality is a feasibility constraint for the matching p. The second inequality follows from the fact that if $\chi^p\,(q) = 1$ then $p_h\,(q)\,h \leq (p_l\,(q) + p_h\,(q))\,l$, that is, $p_l\,(q) \geq p_h\,(q)\,[(h-l)\,/l]$. The second equality is again implied by (1). □

Proof of Proposition 2. Profit under the FRM matching is $\mu h \mathbb{E}[q]$. If $\mu > l/h$, then the profit under BOM is smaller than that under FRM if

$$F(q^*)\mathbb{E}\,[q \mid q \leq q^*]\,h + (1-F(q^*))\frac{l}{h}\mathbb{E}\,[q \mid q > q^*]\,h < \mu h \mathbb{E}[q]$$

where we have used the fact that albeit sellers with $q \geq q^*$ charge l, they are indifferent between charging l and h. Dividing both sides by μ, simplifying h away and rewriting $\mathbb{E}[q]$ we get

$$\frac{F(q^*)}{\mu}\mathbb{E}\,[q \mid q \leq q^*] + \frac{1-F(q^*)}{\mu}\frac{l}{h}\mathbb{E}\,[q \mid q > q^*] < \mathbb{E}[q] = F(q^*)\mathbb{E}\,[q \mid q \leq q^*] + (1-F(q^*))\mathbb{E}\,[q \mid q > q^*]\,.$$

Now focus on the inequality between the left and right side of the above. Since $q^* = F^{-1}\left(\frac{\mu h - l}{h-l}\right)$, we have

$$\frac{F(q^*)}{\mu} + \frac{1-F(q^*)}{\mu}\frac{l}{h} = 1.$$

Hence, both sides are weighted sums of the same conditional expectations. Then, to conclude the proof of the statement that profit under BOM is below profit with FRM observe that

$$\mathbb{E}\,[q \mid q > q^*] > \mathbb{E}\,[q \mid q \leq q^*]$$

and, because $\mu \geq l/h$,

$$\frac{F(q^*)}{\mu} > F(q^*) \text{ and } \frac{1-F(q^*)}{\mu}\frac{l}{h} < 1-F(q^*).$$

Again, continue to assume $\mu > l/h$. The FRM welfare is as profit, that is $\mu h \mathbb{E}[q]$. Rewrite it as

$$\mu h \mathbb{E}[q \mid q \leq q^*] F(q^*) + \mu h \mathbb{E}[q \mid q \geq q^*](1-F(q^*))$$

The BOM welfare is

$$h\mathbb{E}[q \mid q \leq q^*]F(q^*) + \left(\frac{l}{h}h + \frac{h-l}{h}l\right)\mathbb{E}[q \mid q > q^*](1-F(q^*)).$$

Subtracting FRM from BOM we get

$$h(1-\mu)\mathbb{E}[q \mid q \leq q^*]F(q^*) - \left[\left(\mu - \frac{l}{h}\right)h - \frac{h-l}{h}l\right]\mathbb{E}[q \mid q > q^*](1-F(q^*)).$$

The first term is clearly positive. We now focus on when the second term of the difference. It is positive if

$$\left[\left(\mu - \frac{l}{h}\right)h - \frac{h-l}{h}l\right] > 0,$$

or

$$\mu > \frac{l}{h} + \frac{h-l}{h}\frac{l}{h}$$

It is now easy to see that when $\mu \leq \frac{l}{h} + \frac{h-l}{h}\frac{l}{h}$ then the welfare of BOM is higher than the welfare of FRM. To construct an example where the welfare of BOM is lower then FRM one assumes $\mathbb{E}[q \mid q \leq q^*]$ is sufficiently small and $\mu > \frac{l}{h} + \frac{h-l}{h}\frac{l}{h}$. □

References

ARMSTRONG, M. AND J. ZHOU (2021): "Consumer information and the limits to competition," *American Economic Review*.

BERGEMANN, D. AND A. BONATTI (2022): "Data, Competition, and Digital Platforms," *Working Paper*.

BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): "The limits of price discrimination," *American Economic Review*, 105, 921–57.

CONDORELLI, D. AND B. SZENTES (2020): "Information design in the holdup problem," *Journal of Political Economy*, 128, 681–709.

DAMIANO, E. AND H. LI (2007): "Price Discrimination and Efficient Matching," 30, 243–263.

DE CORNIERE, A. (2016): "Search advertising," *American Economic Journal: Microeconomics*, 8, 156–88.

DIAMOND, P. A. (1971): "A model of price adjustment," *Journal of economic theory*, 3, 156–168.

DOGAN, M. AND J. HU (2021): "Consumer search and optimal information," *The RAND Journal of Economics*.

ELIAZ, K. AND R. SPIEGLER (2011): "A simple model of search engine pricing," *The Economic Journal*, 121, F329–F339.

ELLIOTT, M., A. GALEOTTI, A. KOH, AND L. WENHAO (2021): "Marketplace access and information design," *Working Paper*.

GALPERTI, S., A. LEVKUN, AND J. PEREGO (2022): "The Value of Data Records," *Working Paper*.

GOMES, P. AND A. PAVAN (2022): "Price Customization and Targeting in Matching Markets," *Working Paper*.

GOMES, R. AND A. PAVAN (2016): "Many-to-many matching and price discrimination," *Theoretical Economics*, 11, 1005–1052.

HAGIU, A. AND B. JULLIEN (2011): "Why Do Intermediaries Divert Search?" *The RAND Journal of Economics*, 42, 337–362.

JOHNSON, T. (2013): "Matching through position auctions," *Journal of Economic Theory*, 148, 1700–1713.

KHAN, L. M. (2016): "Amazon's antitrust paradox," *Yale lJ*, 126, 710.

LAUERMANN, S. (2013): "Dynamic Matching and Bargaining Games: A General Approach," *American Economic Review*, 103, 663–89.

LIU, Q. (2020): "Stability and Bayesian Consistency in Two-Sided Markets," *American Economic Review*, 110, 2625–66.

NARAYANAN, S. AND K. KALYANAM (2015): "Position effects in search advertising and their moderators: A regression discontinuity approach," *Marketing Science*, 34, 388–407.

ROTH, A. E. AND E. PERANSON (1999): "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *The American Economic Review*, 89, 748–780.

SHAPLEY, L. S. AND M. SHUBIK (1971): "The assignment game I: The core," *International Journal of game theory*, 1, 111–130.

VARIAN, H. (2006): "The Economics of Internet Search," *Rivista di Politica Economica*, 96, 9–23.