

Buyer-Optimal Platform Design*

Daniele Condorelli

University of Warwick

d.condorelli@gmail.com

Balazs Szentes

London School of Economics

b.szentes@lse.ac.uk

February 22, 2025

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 stochastically negative assortative when there are few low-value buyers (i.e., compared to lower-quality sellers, high-quality ones are matched to buyers with lower expected valuation). Although everyone trades, 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 University of Tokyo (Center for Market Design), University of Keio, Hitotsubashi University, 2023 Bonn IO Day, 2022 Berkeley / Columbia / Duke / MIT / Northwestern IO Theory Conference, Tilburg University, the 2021 Warwick Economic Theory Conference, 2021 Piccolo & friends UniBG SKIO conference, Università di Torino, Universitat Autònoma de Barcelona and at the Center for Competition Policy at the University of East Anglia. We also thank Eduardo Azevedo, Andrea Galeotti, Ludovic Renou, Nima Haghpanah, Jacopo Perego and Kai Hao Yang for comments that helped us to improve 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* typically controls (i), but not always (ii). Often, especially if transaction costs for agreeing the exchange are significant, matched parties bargain bilaterally, with limited or no access to alternatives within the platform. For example, Google search advertising exposes consumers to a subset of service providers chosen from a vast pool.¹ Similarly, Airbnb displays a curated subset of rental options for popular destinations, while dating services limit the pool of potential matches for their users.

In this paper, we model the *two-sided matching design* problem of a platform that has superior information about its users but does *not* control their bargaining. Our approach raises the following issue: Since platforms can use user information to optimize matching, agents are able to make inferences from the matching algorithm that may affect post-match bargaining. For instance, Google, which can likely infer users wealth with great accuracy, may prominently display higher-quality opportunities to consumers with greater 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 cannot treat the post-match bargaining outcome of any pair as exogenous, which makes optimal matching design more complex.³

In our model, a platform matches buyers one-to-one with sellers, each owning a single good of differentiated quality. 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, enabling sellers to tailor offers and price-discriminate. The main insight we obtain follows from characterising both the welfare-optimal matching and the buyer-optimal one. We find that a platform can use its control of the matching to garble the information of sellers in a way that offsets their bargaining power and generates information rents for buyers. However, this comes at the cost of creating *sorting* inefficiencies compared to the first-best outcome. Sometimes these sorting inefficiencies are so large that even a fully-random matching, which a platform implements if it does not use any information, generates higher overall welfare despite resulting in missed transactions.

¹Varian (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."

²Businesses 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 showed 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).

³This differentiates us from most, if not all, of the literature on two-sided matching. Consider the paradigmatic example of the National Residency Matching Program studied in Roth and Peranson (1999). Since salaries are set in advance any information revealed by the platform on the preferences of both sides does not affect the outcome post-match.

Before we discuss our key assumptions, let us elaborate 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 pairs, 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. In contrast, 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 sellers. 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 control of prices. This trade-off often resolves in favour of heavily distorting the matching away from positive assortative. While sorting inefficiencies are unavoidable, all matched agents trade despite post-match information is asymmetric.⁴

By assuming that the matching is one-to-one we abstract away from the platform potentially exploiting price-competition. We believe this is a good approximation, especially when firms cannot commit to prices before the match and search costs are present. 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, *ceteris paribus*, heavily biased toward most prominently located opportunities, 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).⁵

Our focus on buyer-optimal matching has, primarily, a normative motivation. We present a framework where some of the welfare trade-offs of regulating a "matching algorithm" to benefit the side lacking bargaining power are exposed and can be evaluated. Nonetheless, 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 who join the platform *ex-ante*, 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 to maximisation of the *ex-ante* surplus of the side that can be charged; or a platform might earn from advertising to one side only and

⁴As we illustrate in Section 7, the absence of trading inefficiencies results from the platform being fully informed. Trading may not be always efficient if the platform cannot perfectly tell high value buyers from low-value ones.

⁵The one-to-one assumption would have a strong bite in a frictionless environment. In the limit case of Bertrand competition, having just two firms competing post-match would bring the price down to cost, thus substantially changing the problem of a buyer surplus-maximising platform.

therefore attempt to maximize participation on that side.

The view that gatekeeping platforms, such as Google and Amazon, should be regulated is popular. 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 Then, 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. Instead, regulators should explore alternative interventions that increase the bargaining power of buyers post-match. Letting buyers make proposals to sellers is, indeed, an existing business model in platform markets. Such 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).

The key motivating assumption in our work is *post-match* bargaining. A direct implication is that the platform is unable to condition the matching on prices. We claim this is a realistic assumption in many cases of interest (e.g., Google search advertising). Nonetheless, in other cases, such as for Airbnb or Amazon, matching of products to consumers takes place based on the prices posted by sellers. We therefore consider a variant of the model where the platform chooses a matching *after* sellers have posted their prices. Our main conclusion is that sorting distortions are a direct consequence of the attempt of the platform to mitigate post-match bargaining power.⁷ In fact, when the platform can condition the matching on prices, then there is no trade-off between buyer surplus and sorting inefficiencies. Independently of which side's surplus the platform is maximizing, the equilibrium outcome is positive assortative matching of buyers to sellers. Moreover, trade is efficient for every matched pair. Nonetheless, whether buyers benefit from firms committing to prices beforehand remains ambiguous. On the one hand, sellers lose their market power against matched buyers compared to our benchmark model. On the other hand, the platform loses the ability to directly persuade sellers to post lower prices.

We now proceed with presenting the model (Section 2) and our results (Section 3). Formal proofs for Section 3 are in the Appendix. In Section 4 we relax the assumption of post-match bargaining. In Section 5 we extend some of our observation to the case where the buyers can have more than two values. We outline the related literature in Section 6. The model also assumes that the platform is fully informed, sellers are only vertically differentiated and qualities are known. The concluding section argues that our main insight is robust to relaxing these assumptions.

⁶The previous 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".

⁷Post-match bargaining also renders cross-subsidization ineffective. Suppose the platform charges one strong side and subsidise the weak. Post-match, the subsidised side will be held-up by the side with bargaining power.

2 The Model

There is a unit mass of buyers. Each buyer has either a low value, $l (> 0)$, or a high value, $h (> l)$. The fraction of buyers with a high value is denoted by μ . We initially assume there is also a unit mass of sellers as this has no impact on the results of the next section. Each seller has a single good to sell. The quality of a seller's product is q and distributed according to the atomless CDF F with support equal to $[\underline{q}, \bar{q}]$, with $\bar{q} > \underline{q} > 0$. If a buyer with value $v (\in \{l, h\})$ purchases the good from a seller with quality q at price t , the buyer's payoff is $qv - t$ and the seller's payoff is t .⁸ The buyers and sellers are matched one-to-one 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, both get their reservation payoff of zero.

Matching.— We describe a *matching* using the probabilities that each seller of quality q is matched with a buyer of valuations h or l . That is, a matching is given by a measurable mapping $p = (p_h, p_l)$ such that $p_h, p_l : [\underline{q}, \bar{q}] \rightarrow [0, 1]$, where $p_h(q)$ and $p_l(q)$ denote 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:

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

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 . 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 $q \in [\underline{q}, \bar{q}]$.

Optimal Prices.— The matching is observed by sellers. Then, if the matching is given by $p = (p_h, p_l)$, the posterior probability of a seller with quality q being matched to 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 matched with

⁸The 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.

probability one. Clearly, the q -quality seller either sets price ql , which gives profit ql , or price qh , which gives an expected profit of $\mu^p(q)qh$. It follows that she will set price ql if, and only if, $\mu^p(q) \leq l/h$ and price qh otherwise. In looking for an optimal matching, it is without loss to focus on equilibria where sellers charge the lowest price when indifferent.

Buyer Surplus.— Let $\chi^p(q) \in [0, 1]$ denote the probability that a seller with quality 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 a given matching p , the buyers' surplus is

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

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, defined as the sum of expected buyer surplus and seller-profit. We therefore start with the following result.

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.*

The 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 . In other words, because every matched pair trades and sorting of buyers to sellers is the best possible, PAM achieves the maximum feasible level of surplus. It is an immediate consequence of Proposition 1 that PAM also maximizes profits of sellers.

Having established that a matching results in either zero surplus for buyers or inefficiencies, we now fully characterize the buyer-optimal matching and equilibrium pricing, showing that such inefficiencies can be sizable. We focus on *buyer-optimal* matchings in the *Pareto frontier*, that is, such that there is no other matching that, in equilibrium, gives higher seller-profit without reducing buyer surplus below its maximum level. This has one main implication: no buyer or seller remains unmatched, even if additional matches do not increase buyer surplus.




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

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 buyer-optimal matching in the Pareto frontier is equal to p^* almost everywhere.

The buyer-optimal matching (henceforth also BOM) of high-value buyers, p_h^* , is exemplified in the two panels of Figure 1. 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.

Let us explain the arguments leading to this theorem. 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 for any 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 buyer 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 mass 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 mass 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.⁹

Let's now define *sorting distortions* in any given matching as the difference between the total welfare generated by PAM and the welfare generated by that specific matching, assuming all transactions occur. In other words, sorting distortions measure the inefficiencies that cannot be directly attributed to informational frictions. Theorem 1 implies that the sorting distortions resulting from maximizing buyer surplus can be greater than those arising from a fully random matching (henceforth FRM), which would be implemented by a platform that does not use any information on buyers and sellers. Visually, when $\mu > l/h$, BOM generates a decreasing p_h , while PAM induces an increasing one and FRM a constant one.

We make a few additional observations on the nature of the optimal matching. First, 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. Second, even if the BOM is stochastically negative assortative in value and quality, the average surplus of buyers matched to high quality sellers is larger than that of buyers matched to lower-quality ones. Third, 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.¹⁰

Welfare analysis.— We conclude this section by comparing BOM's payoffs with those arising under two benchmarks, the PAM and the fully random matching.

⁹An alternative buyer-optimal matching *not* in the Pareto frontier would leave high-valuation buyers unmatched. In this case, matching would not be negative assortative. Hence, readers may wonder whether negative assortativeness is an essential property of buyer-optimal matchings. In Section 5, Example 1, we consider the case of buyers with three possible valuations and show that any buyer-optimal matching is *necessarily* negative assortative.

¹⁰The 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. (2024).

It follows from Proposition 1 that PAM generates higher welfare and higher profit than BOM. It is a consequence of Theorem 1 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), but is equal to zero as in PAM when $\mu > l/h$. Regarding the comparison between BOM and FRM, we have the following proposition.

Proposition 2. *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. In particular, if $\frac{l}{h} < \mu < \frac{l}{h} + \frac{h-l}{h} \frac{l}{h}$, then the welfare of BOM is larger than that of FRM. However, for any given $\mu > 0$, if $\frac{l}{h}$ is sufficiently small, the welfare of FRM exceeds that of BOM.*

A significant consequence of this proposition is that a stochastically negative assortative BOM generates, in some cases, not only larger sorting distortions, but also lower welfare than if the platform ignored information altogether. This observation is not obvious because there is an efficiency trade-off between lost sales arising from too high prices in FRM against the sorting distortions introduced by PAM. Moreover, 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. As we discuss further in the concluding section, more information always reduces bargaining disagreement but may bring about a larger sorting distortion.

Let's now conclude this section by discussing the intuition behind the proof Proposition 2. When $\mu \leq l/h$, the profit of BOM and FRM are the same, as in both cases each seller sets price ql and trade takes place. Hence, the first part of the proposition follows because BOM generates higher buyer surplus by Theorem 1. For the case where $\mu > l/h$, Proposition 2 formalises the following two observations about overall welfare. First, when μ is close to l/h , the loss from not selling to low-value buyers remains positive while the distortion loss arising from BOM goes to zero since $\lim_{\mu \rightarrow l/h} p_h^*(\underline{q}) = \lim_{\mu \rightarrow l/h} p_h^*(\bar{q}) = l/h$. Second, if $l/h \rightarrow 0$, the loss from not selling to low-value buyers in FRM becomes negligible while BOM tends to a fully negative assortative matching given that $\lim_{l/h \rightarrow 0} q^* = F^{-1}(\mu)$. Turning to profit, observe that when $\mu > l/h$ charging h is optimal for all sellers both in FRM and in BOM because of the indifference condition. Hence, higher profit is attained in FRM rather than in BOM, because the value of lost sales to low-value buyers is larger under BOM with h being charged than in FRM, given the negative assortative nature of BOM.

4 Matching on Prices

Because bargaining is *post-match*, the platform is unable to condition the matches on prices. In this section, we consider a variant of the model where firms post (observable) prices *before* the platform implements its matching. Our main conclusion is that, in this case, there is no trade-off between buyer surplus and efficient sorting. The outcome is PAM and trading always takes place. Once the platform is unable to steer pricing, it loses its need to distort the matching.

There are two natural avenues to modify the model, adding the platform as an explicit player. First, we could allow the platform to *commit* to a price-dependent matching before prices are posted. This option trivializes the problem. The platform can implement PAM and any pricing that gives sellers more than their outside option. It does so by threatening sellers with the prospect of remaining unmatched. Second, we could assume that sellers post prices *before* the matching algorithm is chosen. In the remainder of the section we study this variant, which may also appear more realistic in light of the flexibility that platforms have in updating their algorithms.

In this modified game, first, sellers simultaneously post prices. Then, after observing prices, the platform chooses a matching p and, finally, the buyers decide whether to buy or not. We focus only on perfect equilibria which are limits of equilibria in a finite model¹¹. The competition among sellers may intensify if there are more sellers than buyers because sellers are willing to lower their prices in order to guarantee that they are matched. Therefore, we generalise the previous environment by assuming that the measure of sellers is larger than the measure of buyers. We assume that F is a strictly increasing function with $F(\bar{q}) = k \geq 1$. We denote with $\tau : [q, \bar{q}] \rightarrow \mathbb{R}_+$ the pure strategy of sellers. Regarding the platform's objective, we continue to consider three scenarios: it maximizes either buyer surplus, profits, or welfare.

Proposition 3. *The matching is PAM irrespective of the platform's objectives.*

We provide an intuitive argument for the proof, which is relegated to the Appendix. Let $\tau(q)$ denote the posted price of a q -seller. First, we argue that only the worst sellers are unmatched in every equilibrium. If there were two sellers, q_1 and q_2 , $q_1 < q_2$, such that the q_1 -seller is matched but the q_2 -seller is not, the q_2 -seller could deviate and post a slightly higher price than that posted by the q_1 -seller. Since this seller generates more surplus than the q_1 -seller, the platform strictly prefers to match the q_2 -seller, irrespective of its objective, so the deviation is profitable.

It remains to show that higher quality sellers are matched with h -buyers. Suppose, by contradiction, that $q_1 < q_2$ and that the q_1 -seller sells to an h -buyer and the q_2 -seller sells to an l -buyer. We next show $\tau(q_1) > lq_1$. If the platform maximizes profits, the q_1 -seller can always set a price of $hq_1 (> lq_1)$ without changing the platform's desire to match him with an h -buyer. If the platform maximizes welfare or buyer surplus and $\tau(q_1) \leq lq_1$, then it could strictly increase its payoff by matching the q_1 -seller with an l -buyer and the q_2 -seller with an h -buyer.

¹¹This is motivated by the fact that, since there are continuum many sellers and buyers, even when the platform cannot commit to a matching mechanism, it can punish individual deviations at no cost. This gives rise to multiple equilibria, some of which do not capture the platform's commitment problem.

Finally, we demonstrate that the q_2 -seller could profitably deviate by setting a price slightly above $\max\{\tau(q_1), lq_2\}$. Since an l -buyer would not trade at this price, such a deviation essentially commits the seller to not generate surplus unless she is matched with an h -buyer. However, since an h -buyer was willing to trade with the q_1 seller at $\tau(q_1)$, he would also be willing to trade with the q_2 -seller at the new price. In fact, since $q_2 > q_1$ and $\tau(q_1) > lq_1$, the h -buyer's payoff generated by the deviating q_2 -seller is strictly larger than that generated by the q_1 -seller. In addition, the profit of the deviating q_2 -seller would also be larger than that of the q_1 -seller because the deviating price exceeds $\tau(q_1)$. Consequently, irrespective of the objective, the platform strictly prefers matching an h -buyer with the q_2 -seller to matching him with the q_1 -seller irrespective of its objectives. Consequently, any equilibrium matching must be PAM.¹²

The seller-optimal matching, when the platform can condition on prices, aligns with the seller-optimal outcome in our benchmark model. Specifically, a q -seller matched with an l -buyer (or an h -buyer) sells at a price of ql (or qh , respectively). However, the buyer-optimal matching differs between the two versions of the model. The price set by a seller matched with an l -buyer is determined by the requirement that she must generate the same buyer surplus as the highest-quality unmatched seller would if they set a price of zero and were matched with an l -buyer. Similarly, the equilibrium price paid by an h -buyer is established to ensure his surplus matches what he would receive if he were matched with the highest-quality seller who is paired with an l -buyer. As competition among sellers intensifies—i.e., as k increases—prices decline.

Despite the sorting distortions of BOM in the benchmark model, whether buyers prefer one platform type to the other is not a priori obvious. On the one hand, if sellers post prices before being matched, they lose bargaining power and their prices are lower. On the other hand, if the platform moves after prices are posted, it loses ability to persuade sellers and control their prices. In fact, the next example shows that the comparison of the two models from the buyers' perspective is, in general, ambiguous.

Example 1. Let's assume there are as many sellers as buyers and quality is uniformly distributed in $[\underline{q}, \bar{q}]$. That is $F(q) = \frac{q - \underline{q}}{\bar{q} - \underline{q}}$. We can then compute buyers surplus in the PAM where the platform conditions on prices as

$$\mu(h-l)F^{-1}(1-\mu) = \mu(h-l) \left((1-\mu)(\bar{q} - \underline{q}) + \underline{q} \right) \quad (1)$$

¹²This argument is incomplete because one must also show that after the deviation of the q_2 -seller, the platform cannot find an optimal matching in which neither the q_1 -seller nor the q_2 -seller is matched. We provide a proof in the appendix where we construct the unique equilibrium for each of the objectives of the platform.

and buyer surplus in the BOM of the benchmark model as

$$\begin{cases} \mu(h-l)\mathbb{E}[q \mid q \geq F^{-1}(1 - \mu \frac{h}{l})] = \mu(h-l) \left(\frac{q+\bar{q}}{2} + \frac{\bar{q}-q}{2} (1 - \mu \frac{h}{l}) \right) & \text{if } \mu < l/h \\ (1-\mu)l\mathbb{E}[q \mid q \geq F^{-1}(1 - \frac{(1-\mu)h}{h-l})] = l(1-\mu) \left(\frac{q+\bar{q}}{2} + \frac{\bar{q}-q}{2} \left(1 - \frac{(1-\mu)h}{h-l}\right) \right) & \text{if } \mu > l/h. \end{cases} \quad (2)$$

To establish our conclusion that a comparison of the two models in terms of buyer surplus is ambiguous it is sufficient to compare (1) and (2) when $\mu < l/h$. After factoring out $\mu(h-l)$ from both formulas, we can see that buyer surplus from BOM in the benchmark model is lower than buyer surplus in the PAM arising from the variant model if

$$\frac{q+\bar{q}}{2} + \frac{\bar{q}-q}{2} \left(1 - \mu \frac{h}{l}\right) < (1-\mu)(\bar{q}-q) + q.$$

That is, if

$$(\bar{q}-q) \left(1 - \mu \frac{h}{2l}\right) < (1-\mu)(\bar{q}-q) \quad \text{or} \quad (\mu \leq) l/h < 1/2.$$

5 Multiple Buyer Values

In this section we present a partial characterisation of Pareto efficient matchings under the assumption that there are a finite number of possible buyers values drawn from the naturally ordered set $V = \{v_1, v_2, \dots, v_K\}$. We continue to assume that quality is distributed according to F .

Our key finding is the following. Any Pareto efficient matching outcome can be constructed using a segmentation of the buyer market that is *extremal*, in the sense of Bergemann et al. (2015) (BBM). More precisely, assuming the designer maximises some weighted average of buyer and seller surplus, an optimal matching is built by ordering segments of an extremal segmentation 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. As in the two-value case, any inefficiency is due to sorting.

We make three additional observations regarding optimal matching that confirm the insights obtained in the two-value case. First, PAM maximises welfare and sellers' profits. 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, in general, inefficient. Third, the buyer optimal matching can still be stochastically negative assortative.

As a first step we need to introduce some notation, which we borrow from Bergemann et al. (2015) (BBM). We define a market on the buyer-side as a distribution on V and we denote the

set of all markets as $X = \Delta(V)$. For each $x \in X$ let $x(v)$ be the probability of $v \in V$. We denote the market defined by the prior distribution with full support on V as $x^* \in X$. A *segmentation* σ of x^* is a subdivision of buyers in x^* into various submarkets, called segments of σ . Formally, σ is a distribution on X such that the aggregate market, that is the mixture distribution, is equal to x^* . If σ has finite support, we let $\sigma(x)$ be the probability of market $x \in \text{supp } \sigma$. Finally, we say that a segmentation σ is *extremal* if, for each $x \in \text{supp } \sigma$, a monopolist selling to it is indifferent between setting any price equal to the valuations in the support of the segment. Denote an extremal market with support $E \subseteq V$ as x^E .

We now introduce some additional pieces of notation that are needed for our purposes. First, we define a matching as the assignment of a market $x \in X$ to each quality level q . Denoting x_q the market assigned to any seller with quality q , a Pareto efficient matching must satisfy the feasibility condition

$$\int x_q(v) dF(q) = x^*(v) \text{ for all } v \in V. \quad (3)$$

Next, we introduce the platform's objective. We will be interested in matchings that maximize a weighted average of buyer surplus and seller-profit. That is, for $\lambda \in [0, 1]$ our objective function is

$$\lambda \int q cs(x_q) dF(q) + (1 - \lambda) \int q \pi(x_q) dF(q) = \int q [\lambda cs(x_q) + (1 - \lambda) \pi(x_q)] dF(q),$$

where $cs(x)$ and $\pi(x)$ denote the consumer surplus and the monopoly profit in a standard monopoly market $x \in X$ with *unit quality*, respectively. Writing $u_\lambda(x) = \lambda cs(x) + (1 - \lambda) \pi(x)$, the platform problem we aim to solve is

$$\max_{x_q \text{ satisfies (3)}} \int u_\lambda(x_q) q dF(q). \quad (\text{PP-}\lambda)$$

A matching will be Pareto-efficient if and only if it solves (PP- λ) for some $\lambda \in [0, 1]$.

We can think of a matching as a segmentation of x^* made up of a continuum of segments. Conversely, to solve (PP- λ) it will be useful to construct a matching from a generic segmentation of the buyer market with finite support. To this end, we will say that a segment $x \in \text{supp } \sigma$ is matched *uniformly at random* with a certain mass $\sigma(x)$ of sellers if for each of those sellers $x_q = x$. Hence, each of those sellers has posterior equal to $x(v)$ for any $v \in \text{supp } x$.

Which segments are matched with which seller is crucial for our purposes. For any segmentation σ with finite support, let's order the segments in its support as $\{x_1^\sigma, x_2^\sigma, \dots, x_n^\sigma\}$, with the property that $u_\lambda(x_1^\sigma) \geq u_\lambda(x_2^\sigma) \geq \dots \geq u_\lambda(x_n^\sigma)$. We then introduce the following key definition. A matching is a *u_λ -assortative matching based on a segmentation σ* (with finite support) if it is built by pairing sellers in $(F^{-1}(1 - \sigma(x_1^\sigma)), \bar{q}]$ to consumers in x_1^σ uniformly at random, then pairing sellers in $(F^{-1}(1 - \sigma(x_1^\sigma) - \sigma(x_2^\sigma)), F^{-1}(1 - \sigma(x_1^\sigma))]$ to consumers in x_2^σ , and so on until all buyers and sellers are exhausted.

Our main observation in this environment is the following.

Proposition 4. *There is solution to $(PP-\lambda)$ which is an u_λ -assortative matching based on an extremal segmentation of x^* .*

The proposition above indicates that the problem of finding an optimal matching can be simplified by restricting attention to the finite set of all extremal segmentations of the buyers market, further characterized in BBM. The proof follows from two observations. First, for any market $x \in X$ there exists an extremal segmentation that has the same distribution over valuations and where the price charged for all segments of such extremal segmentation is the same as that charged in x (Proposition 2 in BBM). Second, in maximizing $(PP - \lambda)$ we must have that $q' > q''$ implies $u_\lambda(x_{q'}) \geq u_\lambda(x_{q''})$. If not, then we could swap segments without affecting the feasibility condition and without lowering the objective function.

A consequence of the proposition is that, as for the case of two values, in a (Pareto-efficient) optimal matching (in the sense made clear in the statement of the proposition) every matched couple will trade. There will be no inefficiency due to private information, despite information not being symmetric at the bargaining stage. This is because sellers set the lowest possible price, which is equal to quality times the lowest value in the support of the (extremal) segment they are matched with.

It is a simple corollary of the previous proposition that PAM, defined now as u_0 -assortative matching based on the perfect price discrimination segmentation of x^* , maximizes producer surplus and achieves a first-best. Formally, a perfect-price discrimination segmentation of x^* is the σ composed by K extremal segments $x^{\{v_i\}}$ with $\sigma(x^{\{v_i\}}) = x^*(v_i)$ for all $i = 1, \dots, K$.

Corollary 1. *$PP - 0$ is solved by PAM.*

Our final result further characterises the matching and the extremal segmentation arising in the maximisation of buyer surplus. It suggests that the BOM will continue to be coarse, hence inefficient compared to PAM, even with multiple valuations. It extends our main qualitative insight from the previous sections to the case where buyers have multiple valuations.

Corollary 2. *$PP-1$ is solved by a u_1 -assortative extremal segmentations of x^* which contains at most one extremal segment that has a single valuation in its support.*

The corollary restates that a buyer-optimal matchings can be constructed from some extremal segmentation (but, as we shall see, not necessarily one that maximises consumer surplus in BBM) by ordering the segments in terms of consumer surplus. It also states that it is not possible that the segmentation contains two segments that both have a single value in their support.

We conclude this Section with two examples. First, we present an example where the consumer-surplus-maximizing-segmentation of BBM can be used to construct a buyer-optimal matching that result in stochastically negative assortative matching. Second, we show that, in some cases, the optimal extremal segmentation of the buyer-market used in the BOM is not an efficient one that maximises consumer surplus in BBM. The example also demonstrates that the

problem of fully characterising a buyer-optimal matching is not trivial. The idea underlying the example is that when only a small fraction of sellers is able to generate value, the platform maximising buyer surplus may want to focus on segmentations that, albeit resulting in higher profit and therefore lower total consumer surplus, produces higher consumer surplus for some small fraction of buyers that are those matched with the few high-quality sellers.

Example 2. Suppose buyers have three equally likely valuations, (l, m, h) , with $l = 1$, $m = 2$ and $h = 3$, as in BBM's leading example. Note that a unit-quality monopoly price is 2.

A consumer-surplus-maximising extremal segmentation from BBM (see Figure 2) is composed by a segment of mass $2/3$ with support $\{l, m, h\}$, that is $x^{\{l, m, h\}}$, a segment of mass $1/6$ with support $\{m, h\}$, that is $x^{\{m, h\}}$, and a segment with support $\{h\}$ of mass $1/6$. As shown later, in the BOM all buyers in $x^{\{l, m, h\}}$ are matched to high quality sellers, buyers in $x^{\{m, h\}}$ are matched to medium quality sellers and buyers in $x^{\{h\}}$ are matched to low-quality sellers.

Segment	$x(1)$	$x(2)$	$x(3)$	$\sigma(x)$
$x^{\{1,2,3\}}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{2}{3}$
$x^{\{2,3\}}$	0	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{6}$
$x^{\{3\}}$	0	0	1	$\frac{1}{6}$
x^*	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	1

Figure 2: Segmentation of the buyer market into the buyer-optimal matching of Example 2.

It is a matter of algebra to verify that the average buyer value in $x^{\{l, m, h\}}$ is $11/6$ while it is $16/6$ in $x^{\{m, h\}}$ and $18/6$ in $x^{\{h\}}$. Hence the buyer-optimal matching above is stochastically negative assortative and it remains so even if agents in x^h are left unmatched since they provide no strictly positive surplus. Notably, the price is only piecewise monotone in quality. It is equal to hq for the lowest $1/6$ of firms, then it jumps down to mq for the next $1/6$ and jumps down again to lq for the remaining highest quality firms.

To verify the above is the BOM we proceed by using the characterisation obtained above and, hence, by computing average consumer surplus for unit quality in the various extremal segments. We have $cs(x^{\{l, m, h\}}) = 5/6$, $cs(x^{\{m, h\}}) = 1/3$, $cs(x^{\{l, m\}}) = 1/2$, $cs(x^{\{l, h\}}) = 2/3$, $cs(x^{\{j\}}) = 0$ for $j = l, m, h$. Then, it is a matter of algebra and brute force to show that the BBM segmentation described above first-order stochastically dominate any other extremal segmentation in terms of consumer surplus generated. Hence, it is optimal for any distribution of quality.

Example 3. Suppose the three possible valuations of buyers, (l, m, h) , with $l = 1/8$, $m = 1/2$ and $h = 1$. The prior x^* is such that the medium and high value have both probability $1/5$. Note that a unit-quality monopoly profit is $8/40$. If everyone buys in x^* , the unit-quality total welfare is equal to $15/40$. We assume the distribution of quality is such that the quality is either high, $q = 1$, with probability $2/5$, or otherwise low, where we set $q = 0$ for simplicity.

The consumer-surplus-maximising extremal segmentation according to BBM is composed by a segment of mass $4/5$ with support $\{l, m, h\}$, that is $x^{\{l, m, h\}}$, and a segment of mass $1/5$ with support $\{m, h\}$, that is $x^{\{m, h\}}$. Because it is needed for computations, we note that $x^{\{l, m, h\}}(h) = x^{\{l, m, h\}}(m) = 1/8$, while $x^{\{m, h\}}(m) = x^{\{m, h\}}(h) = 1/2$. By construction, while the unit-quality monopolist is indifferent among all prices equal to valuations in the support of each segment, price l is charged in the first segment and price m in the second. Hence, consumer surplus with a unit-quality monopolist is $7/40$. Now suppose we were to build a u_1 -assortative matching based on it. Observe that $cs(x^{\{m, h\}}) = 1/4 > 5/32 = cs(x^{\{l, m, h\}})$. Hence, we would match segment $x^{\{m, h\}}$ first with $1/5$ of the high quality sellers, while the remaining sellers will be matched to buyers in segment $x^{\{l, m, h\}}$. Of these, another $1/5$ will be of high quality. We conclude that the buyer surplus generated by the matching built using the BBM segmentation, is $13/160$.

Next, we show that there is an extremal segmentation of the buyer-market that generates higher unit-quality monopoly profit than $8/40$ but allows a matching with higher buyer surplus. The idea is to maximise the size of the segment producing the highest consumer surplus among all extremal segments, in this case $x^{\{m, h\}}$. Indeed, the buyer-optimal matching in this example requires the creation of an extremal segmentation with a segment with support $\{h, m\}$ of mass $2/5$ and a segment with support $\{l\}$ of mass $3/5$ (see Figure 3). When the segment $x^{\{m, h\}}$ is matched uniformly at random to the high-quality sellers the buyer surplus is $1/10$, which is larger than $13/160$. Crucially, note that this segmentation generates a profit to a unit-quality monopolist equal to $11/40$, which is higher than $1/5$. Hence, it also generates lower consumer surplus overall.

Segment	$x(\frac{1}{8})$	$x(\frac{1}{2})$	$x(1)$	$\sigma(x)$
$x^{\{\frac{1}{2}, 1\}}$	0	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{2}{5}$
$x^{\{\frac{1}{8}\}}$	1	0	0	$\frac{3}{5}$
x^*	$\frac{3}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	1

Figure 3: Segmentation of the buyer market into the buyer-optimal matching of Example 3.

6 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. We depart from tradition by considering an environment where, due to the informational spillovers and bargaining under asymmetric information, the ex-post surplus-sharing of a matched couple is not exogenous, but depends on the matching. Notable exceptions, within the smaller literature that considers stability of matching under asymmetric information, are Liu et al. (2014) and 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.

The design of an optimal matching, but by a revenue maximizing platform, is studied in Damiano and Li (2007), Johnson (2013), Gomes and Pavan (2016), Aoyagi and Yoo (2022) and Gomes and Pavan (2024). 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 object as in classic Myersonian mechanism design. These papers are complementary to ours. Aoyagi and Yoo (2022) is the closest in spirit. There, bargaining is post-match and the sorting distortions, as in our paper, result from an attempt to manipulate information rents. All the other works look at platforms that dictate the conditions at which trading takes places among matched pairs.

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 the incentives of the platform, 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 the platform's incentives, horizontal differentiation and competition. The addi-

tional simplicity allows us to fully characterise the buyer-optimal matching, thus highlighting a trade-off between matching efficiency and consumers' information rent.

In our model, if a pair is formed its participants won't be available to form other matches. Hence, the platform encounters a trade-off when forming a match. We refer to Elliott et al. (2023b), Bergemann et al. (2024) and Bergemann and Bonatti (2024) for models where sellers do not have limited supply and the role of the platform is to control which sellers compete for which buyers. In Elliott et al. (2023b) a fully informed platform manipulates the outcome by making firms uncertain about the set of buyers they can sell to and the other firms they are competing with. In contrast to our work, consumer optimal outcomes does not imply sorting distortions.¹³ In Bergemann and Bonatti (2024), the consideration set of each consumer is auctioned off by the platform, which also reveals the information it has about the buyer to the winning firm.

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.¹⁴ 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. Hidir and Vellodi (2020) and Ichihashi (2020) study the incentives of a consumer to reveal information to a multiproduct monopolist who chooses which product to sell. Information improves match quality but exposes the buyer to price discrimination. The information revealed does not lead to an efficient outcome. In our own work, information is a byproduct of choices that affect both the information structure and feasible surplus, while in the papers above 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 this case, it is not difficult to see that, in light of Proposition 4, our problem is isomorphic to that of identifying the segmentation of consumer's demand that maximizes buyer surplus under a single price-discriminating monopolist. It then follows from Bergemann et al. (2015) that there exists an efficient matching where buyers obtain all surplus minus the profit from random matching.

¹³In Elliott et al. (2023a), the platform does not control the consideration set of buyers but the information that sellers receive about the buyers. On the theme of a platform providing information to sellers see also Yang (2022).

¹⁴As we explained in Section 5, we can map to each matching a distribution of posteriors over buyer valuations that satisfies the martingale property. However, there are many matching with different payoff consequences corresponding to any single distribution of posteriors. Hence, the techniques used to solve Bayesian persuasion problems (e.g., see Kamenica and Gentzkow (2011), Bergemann and Morris (2016), and Dworzak and Martini (2019)) are not directly applicable.

Indeed, one could see our model as a generalisation of BBM whereby the informed monopolist is a seller of an inventory of heterogeneous good and is prescribed by a third party which products to allocate to which segments.

7 Concluding Remarks

In this last section, we discuss extending the analysis to the case of a partially informed platform, private qualities and horizontally differentiated sellers. The take-away is that sorting inefficiency remains a feature of the buyer-optimal matching with post-match bargaining.

Private Qualities.— The assumption that qualities of sellers are observable is plausible, but it is worth asking whether our main insights survive if quality is observable by the platform but not 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 $(q^*, \bar{q}]$ are matched to a mix of high and low-value 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 $(q^*, \bar{q}]$ 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 infer sellers they are matched with from the algorithm, while all sellers with q in $(q^*, \bar{q}]$ charge $\mathbb{E}[q \mid q > q^*]l$. In this equilibrium, sellers in $(q^*, \bar{q}]$ deviating to a different price are believed to be of quality q^* . buyer surplus 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 in which buyers do better. In particular, there is one where all sellers in $[q^*, \bar{q}]$ 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 surplus-maximizing matching 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 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 $(q^*, \bar{q}]$ 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.

Partially informed platform.— Suppose the platform does not know buyers' values. Instead, it gets 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.

Horizontally differentiated sellers.— We conclude this section with an example suggesting vertical differentiation is not key to observing inefficient sorting in buyer-optimal matchings, when bargaining is post-match. 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: Proofs

Proof of Proposition 1. Denote with G the distribution of buyers' values. Since vq is supermodular, a classic result, the Fan-Lorentz Theorem, implies that $\sup_{\pi \in \mathcal{M}(F,G)} \mathbb{E}_{\pi}[vq]$, where $\mathcal{M}(F,G)$ is a *coupling* of probabilities F and G , has a unique *comonotone* solution which is given by PAM. Let w^* be the maximised value of the problem above and note it is an upper bound to the achievable welfare in any matching. To complete the proof of the Proposition we observe that PAM induces complete information. Therefore, each seller q sets price vq where v is the value of the buyer they are matched with. We conclude that the total welfare achieved by PAM is equal to w^* and accrues entirely to sellers. \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:

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

Also, define q^p by

$$\int_{q^p}^{\bar{q}} \frac{l}{h} dF(q) = \int_{\underline{q}}^{\bar{q}} \chi^p(q) p_h(q) dF(q). \quad (4)$$

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

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

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 $x > q^p$,

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

where the first and last equalities are the definitions of H^p and G^p , respectively, the second equality follows from (4) 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,

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

where the first and last equalities follows from the definitions of G^p and H^p , respectively. The second equality is implied by (4) and the inequality follows from the fact that H^p first-order stochastically dominates G^p .

It remains to show that

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

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

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

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 (4). \square

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

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

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

$$\begin{aligned}
& \frac{F(q^*)}{\mu} \mathbb{E}[q | q \leq q^*] + \frac{1 - F(q^*)}{\mu} \frac{l}{h} \mathbb{E}[q | q > q^*] < \\
& < F(q^*) \mathbb{E}[q | q \leq q^*] + (1 - F(q^*)) \mathbb{E}[q | q > q^*].
\end{aligned}$$

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 $\mu \geq l/h$,

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

We now move to discuss welfare comparisons continuing to assume $\mu > l/h$. The FRM welfare in this case is the same as profit, that is $\mu h \mathbb{E}[q]$. Recalling the decomposition above and noting that $F(q^*) = \frac{\mu h - l}{h - l}$, rewrite it as

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

The BOM welfare is

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

Subtracting FRM from BOM we get

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

This is greater than zero if

$$(\mu h - l) \mathbb{E}[q \mid q \leq q^*] \frac{h(1 - \mu)}{h - l} + \left(l + \frac{h - l}{h} l - \mu h \right) \mathbb{E}[q \mid q > q^*] \frac{h(1 - \mu)}{h - l} \geq 0$$

or, as long as $\mu < 1$,

$$(\mu h - l) \mathbb{E}[q \mid q \leq q^*] + \left(\frac{h - l}{h} l - \mu h + l \right) \mathbb{E}[q \mid q > q^*] \geq 0. \quad (5)$$

Now note that $(\mu h - l) \geq 0$ by assumption. So if $\frac{h - l}{h} l - \mu h + l > 0$ then the welfare of BOM is above that of FRM. Rewriting this term gives the condition

$$\mu < \frac{l}{h} + \frac{h - l}{h} \frac{l}{h}.$$

To conclude the proof, now rewrite (5) as

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

Then, note that as l/h goes to zero the term on the left also goes to zero while the term on the right does not since $q^* = F^{-1}\left(\frac{\mu h - l}{h - l}\right)$. \square

Proof of Proposition 3. In what follows we present the main arguments without explicitly approximating our model with a sequence of finite games. It will be clear that the proof can be made precise by introducing such sequences but the notations associated would be demanding.

Consider first the case of a platform which maximises buyer surplus. We first show that the worst sellers will be unmatched, that is, the set of unmatched sellers is an interval starting from zero. By way of contradiction, suppose there are two qualities, q_1 and q_2 , $q_1 < q_2$, the q_1 -seller is matched but the q_2 -seller is not. Then the q_2 -seller could deviate and post a slightly higher price than the q_1 -seller. Since this seller generates more surplus than the q_1 -seller, the platform strictly prefers to match the q_2 -seller, so the deviation is profitable.¹⁵

Next, we consider the set of sellers who are matched with l -buyers. We explain the price of these sellers are determined by the largest quality of unmatched sellers. To this end, let q_c denote the unmatched seller's quality who would provide the largest surplus to an l -buyer among all the unmatched sellers. That is, $q_c = \arg \max_q \{lq - \tau(q) : q \text{ is unmatched}\}$. We argue that if a q -seller is matched with an l -buyer then

$$ql - \tau(q) = q_c l - \tau(q_c).$$

First, the right-hand-side cannot exceed the left-hand-side for otherwise the platform can increase the buyers' surplus by matching the q_c -seller with an l -buyer instead of the q -seller. Second, if the left-hand-side is strictly larger then the q -seller can deviate and increase $\tau(q)$ slightly so that the inequality is still satisfied. Since the platform's response to such a deviation is still to match the q -seller, the deviation is profitable¹⁶. Next, we explain that $\tau(q_c) = 0$. If not, the q_c -seller can deviate and lower $\tau(q_c)$. By the previous displayed equality, the deviating q_c -seller could generate strictly more surplus than any of the seller who was matched with an l -buyer. Hence, the platform will find it optimal to match her. As long as the decreased price is positive, the deviation is profitable because the q_c -seller was unmatched and received a payoff of zero. Now we argue that q_c is the largest quality for which a q -seller is unmatched. If there is a $q > q_c$ such that the q -seller is unmatched, then this q -seller can set a strictly positive price so that it generates

¹⁵Note that the platform cannot find an optimal matching in which the q_2 -seller is unmatched. The reason is that in such a match the platform's payoff must be the same as in the absence of the deviation. However, if the platform uses the original matching except it matches the deviating q_2 -seller instead of the q_1 -seller, its payoff increases strictly.

¹⁶The precise argument is that, by the previous paragraph, any optimal matching matches the q -seller. Therefore, if the platform does not match the deviator, buyer surplus decreases strictly. This implies that, if the change in $\tau(q)$ is small, the platform still find it optimal to match the deviator.

strictly more surplus to an l -buyer than q_c -seller would. By the previous displayed equality, this seller would generate more surplus to an l -buyer than any other seller who is matched with an l -buyer. Therefore, the platform would find it optimal to match him and the deviation would be profitable. To summarize, we have shown that $q_c = F^{-1}(1 - k)$, the set of unmatched sellers is $[0, q_c]$, $\tau(q_c) = 0$, and the price of a q -seller who is matched with a l -buyer is

$$\tau(q) = (q - q_c)l. \quad (6)$$

Let us turn our attention to those sellers who are matched with h -buyers. As will be argued, the price of such sellers will be determined by the largest quality seller who is matched with an l -buyer. To this end, let q^c denote the seller who is matched with an l -buyer who would provide the largest surplus to an h -buyer among all sellers matched with l -buyers. That is, $q^c = \arg \max_q \{hq - \tau(q) : q \text{ is matched with an } l\text{-buyer}\}$. Note that q^c is indeed the largest quality of a seller who is matched with an l -buyer¹⁷. Then, it must be that

$$hq - \tau(q) = hq^c - \tau(q^c). \quad (7)$$

If the right-hand-side was strictly larger than the left-hand-side, the platform could strictly increase buyers' surplus by matching an h -buyer with the q^c -seller instead of the q -seller, letting the q -seller to remain unmatched, and matching the l -buyer (who was originally matched with the q^c -seller) with the q_c -seller. By equation (6) and $\tau(q_c) = 0$, this l -buyer's payoff is unaffected by such a change. However, the h -buyer's payoff would strictly increase. Suppose now that the left-hand-side is strictly larger. Then the q -seller could slightly increase the price so that the inequality remains strict. We have to argue that after such a deviation, the platform still matched the q -seller. If the platform lets the q -seller be unmatched, the best it can do is to create the same matches except to match the h -buyer who was matched with the q -seller with the q^c -seller and match the l -buyer who was matched with the q^c -seller with the q_c -seller¹⁸. Again, the payoff of the l -buyer would remain the same but now, the payoff of the h -buyer would decrease.

To conclude the equilibrium matching is PAM, it remains to show that the sets of sellers who are matched with l -buyers and h -buyers are also intervals. In other words, we need to show that the quality of each seller matched with an h -buyer is larger than the quality of a seller matched with an l -buyer. Since the largest seller matched with an l -buyer is q^c , it is enough to show that

¹⁷To see this, note that equation (6) implies

$$hq - \tau(q) = hq - (q - q_c)l = (h - l)q - lq_c,$$

which is increasing in q .

¹⁸To see this note that it is without loss to assume that, in the new matching, the h -buyer (who was matched with the deviating seller) is matched with a seller who was previously either unmatched or matched with an l -buyer. The latter cannot be optimal because the deviating q -seller still generates a larger surplus for the buyer than any of the sellers who were matched with an l -buyer. The former also cannot be optimal because, by equation (6), unmatched sellers generate less surplus to h -buyers than those sellers who are matched with l -buyers.

the quality of any seller matched with an h -buyer exceeds q^c . By contradiction, suppose that $q \in (q_c, q^c)$ and the q -seller is matched with an h -buyer. Observe that

$$\tau(q) = h(q - q^c) + l(q^c - q_c) < l(q - q_c),$$

where the equality follows from (6) and (7) and the inequality follows because $q < q^c$ (so $h(q - q^c)$ is negative) and $q_c < q$ (so $l(q^c - q_c) < l(q - q_c)$). Suppose now that this q -seller deviates by increasing his price slightly so that the previous inequality is still satisfied. By the previous displayed inequality chain and equation (6), this q -seller would generate more surplus to an l -buyer than any seller who is matched with an l -buyer. Consequently, the platform would find it optimal to match the seller, and hence, the deviation is profitable. We conclude that q^c solves $F(q^c) = k - 1 + [1 - \mu] = k - \mu$.

Consider now a welfare-maximizing platform. It is not hard to show that the q -seller who is matched with an l -buyer sets price lq and the q -seller who is matched with an h -buyer sets price hq in equilibrium. If the price was lower, the seller could raise it slightly. Since that increase in price does not affect welfare, the platform will still match this seller with the same buyer. Let us now argue that the equilibrium matching is PAM. First, it is clear that the worst sellers will be unmatched because a high quality seller can offer a small but positive price and guarantee that she will be matched. It remains to show that higher quality sellers are matched with h -buyers. Suppose, by contradiction, that $q < q'$ and that the q -seller sells to an h -buyer and the q' -seller sells to an l -buyer. Then the q' -seller could deviate by setting price slightly below hq' instead of setting lq' . Such a deviation essentially commits the seller not to generate surplus unless she is matched with an h -buyer. After such a deviation, the platform strictly prefers matching an h -buyer with the q' -seller to matching him with the q -seller. Finally, note that the arguments we made in this paragraph are also applicable if the platform maximizes the sellers' profits. Hence, the matching will also be PAM in this case. \square

References

- AOYAGI, M. AND S. H. YOO (2022): “Matching strategic agents on a two-sided platform,” *Games and Economic Behavior*, 135, 271–296.
- ARMSTRONG, M. AND J. ZHOU (2021): “Consumer information and the limits to competition,” *American Economic Review*.
- BERGEMANN, D. AND A. BONATTI (2024): “Data, Competition, and Digital Platforms,” *American Economic Review*, 114, 2553–95.
- BERGEMANN, D., A. BONATTI, AND N. WU (2024): “How Do Digital Advertising Auctions Impact Product Prices?” *The Review of Economic Studies*, rdae087.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The limits of price discrimination,” *American Economic Review*, 105, 921–57.
- BERGEMANN, D. AND S. MORRIS (2016): “Bayes correlated equilibrium and the comparison of information structures in games,” *Theoretical Economics*, 11, 487–522.
- 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,” *Economic Theory*, 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.
- DWORCZAK, P. AND G. MARTINI (2019): “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy*, 127, 1993–2048.
- 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 (2023a): “Market segmentation through information,” *Working Paper*.
- (2023b): “Matching and Information Design in Marketplaces,” *Working Paper*.
- GALPERTI, S., A. LEVKUN, AND J. PEREGO (2024): “The Value of Data Records,” *The Review of economic studies*, 91, 1007–1038.

- GOMES, R. AND A. PAVAN (2016): “Many-to-many matching and price discrimination,” *Theoretical Economics*, 11, 1005–1052.
- (2024): “Price customization and targeting in matching markets,” *The Rand journal of economics*, 55, 230–265.
- HAGIU, A. AND B. JULLIEN (2011): “Why Do Intermediaries Divert Search?” *The RAND Journal of Economics*, 42, 337–362.
- HIDIR, S. AND N. VELLODI (2020): “Privacy, Personalization, and Price Discrimination,” *Journal of the European Economic Association*, 19, 1342–1363.
- ICHIHASHI, S. (2020): “Online Privacy and Information Disclosure by Consumers,” *American Economic Review*, 110, 569–95.
- JOHNSON, T. (2013): “Matching through position auctions,” *Journal of Economic Theory*, 148, 1700–1713.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KHAN, L. M. (2016): “Amazon’s antitrust paradox,” *Yale LJ*, 126, 710.
- LIU, Q. (2020): “Stability and Bayesian Consistency in Two-Sided Markets,” *American Economic Review*, 110, 2625–66.
- LIU, Q., G. J. MAILATH, A. POSTLEWAITE, AND L. SAMUELSON (2014): “Stable Matching With Incomplete Information,” *Econometrica*, 82, 541–587.
- 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.
- YANG, K. H. (2022): “Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences,” *American Economic Review*, 112, 1364–93.