

Who benefits from improved search in platform markets? *

Gregory Lewis

Harvard University and NBER

Albert Wang

Harvard University

April 8, 2013

Abstract

Online platforms invest large sums in their search technology. Motivated by this, we investigate how lowering search costs affects the welfare of market participants, in a model where buyers with horizontally differentiated tastes search and compete for goods in an auction. We identify a “matching effect”, whereby lower search costs lead to better matches; and a “segmentation effect” whereby lower search costs endogenously shift market participation in favor of some goods and against others. We prove that there is a unique equilibrium, and demonstrate that the decentralized market achieves the social planner’s solution. Decreasing search costs thus improves joint welfare; and yet surprisingly joint seller revenue may fall.

*. We are grateful to Susan Athey, Stephanie Hurder, Sonia Jaffe, David Laibson, Zhenyu Lai and Eric Maskin for useful comments that helped us to improve the paper. Greg has benefited from the support of the National Science Foundation.

1. Introduction

New technologies often bring new markets and market structures. Just as advances in shipping technology opened up countries to international trade in goods that were previously too heavy to ship, advances in communication and computation have created new marketplaces in goods, media, and advertising that could not have existed earlier. Many of these new markets are organized around platforms, which bring together large numbers of consumers and producers, offering each side easy access to the other. In some ways, these platforms resemble marketplaces that have existed for centuries: platforms act as intermediaries to facilitate transactions between agents that would otherwise not occur. In other respects, however, these new platforms are *sui generis*. The platforms themselves often represent substantial investments, and they embody many proprietary technologies that facilitate trade. In many cases, the platforms themselves are far more visible than any of the transactions that take place on them: eBay and Amazon, for instance, have become household names.

Platforms help to overcome several key frictions: (1) they provide market thickness; (2) they reduce transaction costs; and (3) they reduce search costs (Hagiwara 2009). The first two functions are relatively well-understood, and they represent the two foremost challenges that a platform must solve in terms of its development strategy. The consequences of search within platforms, however, are less well understood, in spite of the growing importance that search facilitation seems to be playing on many platforms. For instance, media platforms, such as Netflix and Amazon, invest heavily in their recommendation systems, which help to direct viewers and readers to movies and books that they are likely to enjoy. In fact, Netflix, for several years, offered annual prizes to developers who could improve their user rating prediction algorithms. Some platforms, such as dating sites, live or die by the efficacy of their search systems: what matters for the paying members of a dating site is not the number of *potential* matches that might exist on the platform, but rather the *quality* of the few matches that actually do take place. The lesson seems to be that creating market thickness on the other side is not always enough. In situations where match *quality* is a central part of platform outcomes, a platform also has a role in making desirable matches easier to achieve.

The implicit intuition behind such efforts to facilitate search within platforms is that lower search costs represent a “win-win” situation for the agents on the platform and lead to Pareto-improving outcomes. In the case of Amazon, for instance, low search costs make it more likely that readers find books that they will want – they are better off since they

end up getting books that they like more. Sellers are also better off because the readers most likely to purchase their books get exposure to them, and so their revenues are likely to increase.

This “win-win” logic implicitly requires that the agents on the one side are not competing for the goods on the other (e.g. book buyers do not interact). Yet in many markets, such competition is the norm. In a dating market in which men search for women, a relatively unattractive man has a better chance of being matched to an attractive woman if the search technology is poor; since when the technology is good, she will be identified and receive offers from equally attractive men. Similarly, in an auction market, bargains are more likely to be had when search is difficult. Indeed the history of eBay is littered with anecdotes about people who managed to earn a living arbitraging spelling mistakes.

To the best of our knowledge, this is the first paper to examine the implications of search and search costs on outcomes and welfare, in a platform setting where *agents on one side compete for goods on the other side*. We consider a market where buyers with horizontal “Hotelling-style” preferences over two goods decide whether or not to search for their preferred good, before participating in an auction for that good. A finite version of this model — say with N buyers and exactly 2 goods/sellers — is not analytically tractable, and an important contribution of the paper is to set up a continuum model with an infinite number of buyers and sellers. This allows us to reduce the large platform game to a continuum of Poisson games, much as Myerson (2000) did in voting games, enabling analytic results.

We find that the “win-win” intuition does not hold under competition on the one side. It is true that the total surplus generated by the platform increases with a better search technology, since matches will generally be of higher quality. However, this need not be a Pareto improvement. Sellers of the less-preferred product may end up with few bidders on their goods, harming revenue. And buyers with weak preferences for a universally preferred product may get lucky and win it when search is difficult; but be completely uncompetitive when search is easy, and have to make do with the the other good (again the analogy to dating makes this effect clear). In fact, we demonstrate that sellers can jointly lose with lower search costs, if one product gains so many more buyers that revenues from the other product collapse. We find that a sufficient condition for lower search costs to raise the revenues of all sellers is symmetry: there are an equal number of sellers of each good, and the distributions of valuations for each good are identical. In such cases, the measure of buyers faced by each seller remains constant regardless of search intensity, so that the only channel affecting seller revenues is the creation of better matches, which is revenue increasing.

Literature Review Two papers are closely related. The first is Hagiu and Jullien (2011), who analyze a model of search intermediation. In their model, two types of consumers must access two types of sellers through a search intermediary (i.e. a platform), who may direct them to their preferred good or not. Their main result is that platforms sometimes have an incentive to divert search, in the sense that they may direct buyers to their less preferred sellers and prevent some surplus-enhancing trades from taking place. Our model has some similarities to theirs but differs notably in treating the platform as fixed, rather than a strategic player, and in allowing for competition between buyers for goods through the auction process.

The other recent paper is a study by Tadelis and Zettelmeyer (2011), which provides a direct motivation for our analysis. In a field experiment involving an auction platform for used car transactions, they found that making car quality information available online — rather than requiring potential buyers to inspect the cars in person — improved revenues for all car grades, *including those which were poorly rated*. They attributed to this to improved matching: buyers who valued low quality cars relatively more than average were able to find those cars easily due to the information provision, and similarly those with a high taste for quality were better able to match. Notice that this experimental treatment essentially reduced search costs — albeit only for some subset of cars — and so is very similar to the comparative statics exercise we will investigate below. We extend the simple model presented in their paper to a much more general environment.

Most of the literature on platform markets has focused almost exclusively on network effects (Armstrong 2006; Rochet and Tirole 2006; Weyl 2010; Weyl and White 2012), and the implications of these effects for optimal pricing, either under monopoly or oligopoly. These network effects also have important implications for the number of platforms that can exist in equilibrium, and how they might be structured. In general, platform markets will support only one or two platforms in equilibrium because of “tipping” effects (Ambrus and Argenziano 2009; Ellison and Fudenberg 2003). Recently, Hagiu and Lee (2011) have examined the issue of exclusivity, and shown that a content provider will generally multi-home across multiple platforms when they are able to maintain control over pricing. Halaburda and Yehezkel (2011) have extended existing work on platform competition to the setting where both sides of the market are uncertain about the value of a new technology.

The search mechanics of our model are based on well-established models of sequential search (Diamond 1971; Mortensen 1970). What differentiates our model of search on platforms from existing models of consumer and job search, however, is that there is many-to-one

matching, so that changes in search costs affect a buyer’s utility indirectly through competition in the auctions. We address these complications by using results on Poisson games, building upon the work of Myerson (2000).

Our paper is organized as follows: the next section presents a simple discrete example to illustrate the effects of search within platforms. The third section establishes the formal model and notation; while the fourth presents the analysis and results, before we conclude. All proofs are in the appendix.

2. A Motivating Example

In this section, we present an example similar to that in Tadelis and Zettelmeyer (2011) that illustrates how decreasing search costs can increase average seller revenue. Consider an auction platform market with two goods available for sale, A and B , and four buyers. Buyers have a private valuation v of receiving either good that is independently and uniformly distributed on $[0, 1]$. In addition, two of the buyers are commonly known to prefer good A , and get an additional payoff of 1 if they win it; while the other two are commonly known to prefer good B , and get a corresponding additional payoff of 1 if they win it. We assume quasi-linear utility, so that the payoff of a buyer who prefers and wins good A is $1 + v - p$ for p the price they pay.

The buyers play a two stage game. In the first stage, buyers make a search decision. To keep this example simple, we focus on two cases. Either search is *costless*, and the buyer can find their preferred good; or it is *infinitely costly*, and the buyer gets matched to a random good. We will generalize to a full sequential search model with continuous search costs below. Once buyers and goods are matched, the goods are sold in a sealed-bid second-price auction. Buyers observe the kind of good that they’re bidding on and make a bid.

Now consider the revenues in each of the cases. In the costless case, the two buyers who prefer good A will choose to enter the good A auction; and those who prefer good B will enter the auction for good B . This will result in (symmetric) expected revenues from each auction of $1 + \frac{1}{3} = \frac{4}{3}$, since the expected value of the second order statistic from a sample of n uniform random variables is $\frac{n-1}{n+1}$. In the infinitely costly case, random matching implies that the number of bidders in each auction are $n_A + n_B$ and $4 - n_A - n_B$ respectively, where n_i is the number of buyers who prefer good i . n_A is independent of n_B and both are distributed Binomial $(2, \frac{1}{2})$. This leads to nine possible outcomes — one for each pair of draws of n_A and n_B . Doing the relevant algebra, we get expected revenues of $\frac{27}{48}$.

This confirms the intuition that improving search should raise seller revenues. There are two reasons for this. First, with lower search costs, matching improves: the agents who prefer a good end up bidding on that good. Second, there is less chance of coordination failure, whereby all the agents end up in one auction and none in the other auction. To see the impact of this second effect, notice that expected revenues under random matching *conditional on there being exactly two agents in the auction* are $\frac{33}{48}$, higher than in the case of pure random matching.

This example is special in a number of ways. The environment was entirely symmetric, so that costless search led to the market for each good being equally thick. This need not be true in general: we will later show that lower search costs can dramatically thin the market for one good, harming seller revenues. We also chose to tackle only the extreme cases of costless and infinitely costly search, essentially making the search decision exogenous. Once we endogenize it, each individual buyer's search decision will depend in equilibrium on who else is searching, making the problem more complex – the more buyers searching for a particular good, the less attractive entering into an auction for that good becomes. Finally, this example was discrete, which led to a lot of different cases. As the number of buyers and sellers on the platform grows, this model very quickly becomes unmanageable. For this reason we will move to a continuum framework so that auction participation will vary continuously rather than discretely.

3. Model

There are a measure μ of buyers, and a measure 1 of sellers on the platform. Sellers are one of two types, A and B , which denotes the type of good that they have. A fraction α of sellers are type A . Goods are sold off in second-price auctions. Buyers have quasi-linear utilities, and their types are independently and identically distributed on $\mathbf{X} = [\underline{x}, \bar{x}]$, where a buyer's type x also denotes his valuation for the A good. A buyer of type x has a valuation for a B good given by $y(x)$, where y is continuous and strictly decreasing in x , so buyers with relatively high valuations for the A good have relatively low valuations for the B good.¹ Let \mathbf{Y} denote the range of y .

Buyer types are independently and identically distributed according to F , where we assume F is atomless with probability density function f . We will also let $\mu(X)$, for measurable

1. Notice that this does not impose a restriction on the absolute valuations: it could be that every buyer values the A good more than the B good; rather, this is a restriction on relative valuations.

$X \subseteq \mathbf{X}$, denote the measure of buyers with types in X , and $\mu(x)$ denote the density of buyers at x . Also, let $F(\cdot)$ be the cdf of the value distribution of buyers, i.e., $F(x) = (1/\mu)\mu(\{x' : x' \leq x\})$. It will also be convenient to define $H(y) = 1 - F(y^{-1}(y))$ to be the cumulative distribution function of valuations for the B good.

The goods are allocated according to a two-stage game. In the first stage, buyers are matched to sellers via sequential search; in the second stage, goods are sold off in second-price auctions. The search process is sequential search mediated by platform technology. Buyers are initially randomly matched to sellers. Upon being matched, buyers observe the type of good being sold on the auction (but not who else is possibly going to participate in this auction). Each buyer can stay in the auction he is matched to, or at a cost $c > 0$, he can draw another match

The success of each draw is determined by a platform-wide parameter p , which we call the *search efficiency*. With probability p , a buyer will be matched with an auction for the good they're looking for; with probability $1 - p$, the match will be random. Hence, $p = 0$ corresponds to completely-random matching as in the standard sequential search model, and $p = 1$ corresponds to a perfectly efficient search technology. Hagiu and Jullien (2011) refer to this search process as platform intermediated search. It nests the standard sequential search model, and allows for comparative statics on the search technology itself (through the search efficiency p) and the search frictions c , which are conceptually distinct. The search stage ends when nobody chooses to draw any more matches.

In the auction stage, we restrict the buyers to playing weakly dominant strategies, so that all buyers bid their valuations. The distribution of types in auctions for each kind of good will be determined in equilibrium by the search decisions.

Discussion: We have deliberately kept the primitives of the model as simple as possible. We are interested in the equilibrium properties of a decentralized auction platform: the platform doesn't match bidders to auctions, nor does it clear the market for each good through a multi-unit auction (which will generally be more efficient). This is in fact how many auction platforms operate, although typically these markets are dynamic so that buyers have multiple bidding opportunities.

An important restriction is to one-dimensional types with horizontal "Hotelling-style" preferences. This simplifies the analysis. But it also emphasizes the importance of matching; lowering search costs would seem to be particularly good for revenues in this environment, since it facilitates assortative matching. The fact that we are able to show that total revenue

can decrease with search even here is thus a strong finding.

4. Analysis

We will proceed by backward induction. First, we develop the analytical tools necessary to characterize expected payoffs and revenues from the auctions for each good. In this step we take as given the distribution of types participating, which will be endogenously determined by search.

Next, we use these expressions to show that the search strategies can be characterized by a pair of thresholds, where those types who prefer good A sufficiently strongly search for it until they find it; those with good B search similarly for it; and the rest don't search. Finally, we show existence and uniqueness of the Nash Equilibrium, consider some comparative statics on search costs, and see what the implications are for efficiency and revenue.

4.1. The Auction Game

The mass and distribution of types bidding on each good is determined by the search strategies. Let μ^j be the mass of types bidding on good j , and F^j be the cumulative distribution function, for $j = A, B$. We would like to characterize the buyer's interim expected surplus, after they have searched into an auction for a particular good, but before they know who they have been matched with and what the auction outcome is. As Myerson (2000) has shown, random matching of buyers to auctions implies that the number of buyers in auctions for good j is a Poisson random variable, with parameter μ^j . In fact, the number of buyers from any subset of X is a Poisson variable with parameter equal to the measure of that subset, and is independent from the number of buyers who show up outside of that subset. We formalize this in a lemma, which is proved in the appendix:

Lemma 1. *Let X_1 and X_2 be disjoint measurable subsets of X , and let k_1 and k_2 be random variables denoting the number of buyers who show up at a given auction for good j with valuations in X_1 and X_2 respectively. Then k_1 and k_2 are independent Poisson random variables with parameters $\mu^j(X_1)$ and $\mu^j(X_2)$ respectively.*

This makes it easy to calculate the probability of winning, and hence the expected surplus. For example, a bidder with valuation x wins an auction for good A iff their valuation is the highest; or equivalently if there are no bidders with valuations higher than x . By Lemma 1,

this latter event is independent of the fact that they themselves are in the auction.² It therefore takes place with chance $e^{\mu^A(1-F^A(x))}$, from the Poisson probability mass function. Using this expression, we can derive expected utility and revenues. The derivation is quite standard (see .e.g., Krishna 2009), applying integration by parts (proof in appendix).

Lemma 2. *The interim expected utility of a bidder with valuation z for good j , and expected seller revenue for a good of type j , are respectively given by*

$$u^j(z) = \int_0^z e^{-\mu^j(1-F^j(s))} ds \quad R^j = m^j(\bar{z}) - \mu^j \int_0^{\bar{z}} (1 - F^j(z)) G^j(z) dz$$

where \bar{z} is the highest valuation for good j , $G^j(z)$ is the distribution of the first-order statistic in j auctions, and $m^j(z) = \int_0^z s g^j(s) ds$ is the expected payment of a buyer with valuation z .

4.2. Sequential search

The preceding lemma provides closed-form expressions for buyer utilities in the second stage of the full model with search, as a function of the aggregate search behavior and resulting distributions of types bidding on each good. This allows us to analyze individual search decisions and from there derive the equilibrium of the full game. A search strategy is a stopping rule: a buyer must decide whether or not to stop searching given the type of good he is currently matched to and the results of previous searches. But since there are a continuum of sellers the result of each search does not depend on the previous results. This implies that the optimal search policy is stationary: it specifies whether or not to stop as a function only of the current good held.

Since there are only two goods, a search policy is just a binary 2-vector, saying whether or not to stop searching when holding each good. Types who stop regardless of whether they are holding good A or good B we will term non-searchers. They will stick with their initial random match. A second group are those who stop when holding A and continue searching when holding B ; and a third are those with the reverse behavior. There are only three groups.³

The probability of finding one's desired good depends on the search technology and the relative proportions of each good. Recall that a fraction α of sellers sell good A . So a buyer

2. One can show this formally by taking X_1 as an open interval $(x - \varepsilon, x + \varepsilon)$, taking X_2 as $(x + \varepsilon, \bar{x})$ and then considering limits as $\varepsilon \rightarrow 0$ by L'Hôpital's rule.

3. The final possible policy — never stop — is clearly not optimal under costly search.

currently holding good B will find good A on the next draw with probability $p_a = p + (1-p)\alpha$; similarly the chance of finding good B when holding A is $p_b = p + (1-p)(1-\alpha)$. Notice how the search technology parameter acts to make the relative proportions either unimportant (with perfect search technology) or very important (with no search technology).

Our next step is calculate the optimal search behavior of each buyer, holding the behavior of the other buyers fixed (i.e. it is a best response). This requires only that we classify each type into one of the groups. Let $u^A(x)$ denote the utility a buyer of type x gets from participating in an A auction, and $u^B(x)$ denote his utility from a B auction. These are well-defined given other buyer behavior through Lemma 2. Then the search decision amounts to a simple comparison of the marginal benefit from searching against its marginal cost c . A buyer will search into an A auction if $p_a (u^A(x) - u^B(x)) \geq c$ and will search into a B auction if $p_b (u^B(x) - u^A(x)) \geq c$. If neither condition holds, they will not search.

Now, by Lemma 2, the interim expected utility of a buyer for a good is increasing in their valuation for that good. Moreover, since $y(x)$ is strictly decreasing in x , when a buyer's valuation for good A increases, their valuation for good B decreases; and vice-versa. So the interim utility difference $u^A(x) - u^B(x)$ is strictly increasing in x . This means that the best responses will take the form of threshold rules in the type space:

Lemma 3. *The best responses are characterized by a pair of thresholds (x_b, x_a) such that all buyers with type $x \geq x_a$ search into A , all buyers with type $x \leq x_b$ search into B , and the remaining buyers do not search. Furthermore, for interior x_a and x_b , the thresholds must satisfy*

$$\begin{aligned} u^A(x_a) - u^B(x_a) &= \frac{c}{p_a} \\ u^A(x_b) - u^B(x_b) &= -\frac{c}{p_b} \end{aligned}$$

4.3. Equilibrium

We look for a pure strategy Bayes Nash equilibrium of the game. To prove existence and uniqueness, we exploit the characterization of the best responses in Lemma 3 above. We define a mapping $\mathbf{M} : \mathbf{X}^2 \rightarrow \mathbf{X}^2$ which takes any pair of thresholds (x_b, x_a) — which through search induce distributions of bidders μ^A and μ^B — and returns a new pair of thresholds (x'_b, x'_a) that describe the individual bidder best responses to this aggregate search behavior. A fixed point of this mapping is an equilibrium. In the appendix, we prove continuity of \mathbf{M} . Then since \mathbf{X}^2 is convex (a two-dimensional interval), Brouwer's fixed point theorem

delivers existence.

The uniqueness argument is more involved. We define a mapping $\mathbf{S}(x_b, x_a)$ such that $\mathbf{S}(x_b, x_a) = \mathbf{c}$ iff (x_b, x_a) is an equilibrium (where \mathbf{c} is a constant). We then show that the Jacobian of \mathbf{S} is globally positive, and the principal minors are non-vanishing. Following Gale and Nikaido (1965), this implies that \mathbf{S} is one-to-one, so $\mathbf{S}(x_b, x_a) = \mathbf{c}$ has a unique solution.

Intuitively, we get uniqueness because of the negative relationship between the number of people searching and the incentives to search: if more people search for good A (i.e. x_a is marginally lower), then searching for good A is a little less valuable (i.e. x_a should be marginally higher in response).

Theorem 1 (Existence and Uniqueness). *An equilibrium exists and is unique. A reduction in search costs, either through a decrease in the per-draw cost c , or through an increase in the search efficiency p , implies x_b increases and x_a decreases, so that more types search.*

The comparative statics are intuitive. When search costs decrease from some initial equilibrium, the marginal types now strictly prefer to search, and the thresholds adjust smoothly to restore indifference. The implication of this is that a platform can induce more search and better matches by improving the search technology. The question we now turn too is whether it wants to.

4.4. Welfare

In this section, we establish two basic welfare properties of equilibrium, and give some support to the idea that reducing search costs is generally a good thing overall. First, the equilibrium is efficient, in the sense that the market equilibrium coincides with the solution to the social planner's problem (where the social planner can control the assignment of agents to goods, but not individual auctions). In this exercise we hold the search costs and technology fixed. Second, we argue as a simple corollary that as search costs decrease, social welfare is increasing.

To show this formally, we first define the social planner's problem. The social planner assigns to each buyer x an action in the set $\{a, 0, b\}$, where a and b correspond to searching into A and B respectively, and 0 corresponds to not searching. The objective function of the social planner is total welfare net search costs, which we denote by W . To get a sense of what social welfare looks like, note that each buyer's contribution to total welfare is his valuation of the object that he bids for times his probability of winning. His presence has

no effect on welfare if he does not win, and if he does win, the total surplus generated is his valuation, which is split between himself and the seller according to the price paid. We can then write the surplus generated in the A market as

$$\int_{\underline{x}}^{\bar{x}} x G^A(x) \mu^A f^A(x) dx = E[X_1]$$

where X_1 is the highest bid in a randomly chosen auction of good A . The surplus generated in market B is similarly $E[Y_1]$, for Y_1 the highest bid in a randomly chosen auction of good B . The total search cost incurred across the platform is c times the expected number of draws. Let $\mu^A = \mu(X_a) = \mu \cdot (1 - F(x_a))$ and $\mu^B = \mu(X_b) = \mu \cdot F(x_b)$ be the measures of buyers who search into A and B markets respectively. A fraction $1 - \alpha$ of those in X_a need to search, and expect to search $1/p_a$ times, so the total cost incurred by them is $\frac{(1-\alpha)\mu^A}{p_a}c$. Similarly the total cost incurred by those searching into B is $\frac{\alpha\mu^B}{p_b}c$. Hence, total welfare W can be written as

$$W = \alpha E[X_1] + (1 - \alpha) E[Y_1] - \left(\frac{(1 - \alpha)\mu^A}{p_a} + \frac{\alpha\mu^B}{p_b} \right) c$$

The lemma below establishes that a solution to the social planner's problem must have a threshold form, e.g., it will consist of a pair (x_b^*, x_a^*) such that all $x < x_b^*$ search into B , all $x > x_a^*$ search into A , and the rest do not search. The argument is by contradiction: if the planner assigned type x to search for A but $x' > x$ to not search, then swapping their assignments would be welfare improving.

Lemma 4. *A solution to the social planner's problem must consist of x_b^* and x_a^* such that all $x < x_b^*$ search into B and all $x > x_a^*$ search into A .*

We next show that the socially optimal thresholds must be the same as in the market equilibrium. We prove this directly, showing that the first order conditions for the social planner's maximization problem correspond exactly to the threshold conditions required for the market equilibrium.

Theorem 2 (Efficiency). *Let (x_b^*, x_a^*) be the social planner's solution, and let (x_b, x_a) be the market solution for a given set of search cost parameters, (c, p) . Then $x_b^* = x_b$ and $x_a^* = x_a$. That is, the market equilibrium maximizes total social welfare and is efficient.*

The intuition for this result comes from McAfee and McMillan (1987). They observed that in a VCG mechanism like the second-price auction, a buyer's utility is equal to his

contribution to social welfare, since his payment is exactly equal to the external effect of his presence on the other buyers. As a result, when buyers maximize their own utilities in our platform model with search, they also maximize their contribution to social welfare.

A corollary to this result is that a decrease in search costs always increases total welfare. This follows immediately from the efficiency result: since the market equilibrium maximizes welfare, and higher levels of welfare are possible with decreased search costs – agents, for instance, can simply retain the same actions, but total welfare will be raised since search costs paid are lower – it must be that equilibrium welfare is increasing as search costs decreases.

Corollary 1. *As search costs decrease, total social welfare is increasing.*

This shows that buyers and sellers are jointly better off as a result of decreased search costs. But it does not say how the gains are distributed. The following section examines the effect of search on seller revenues.

4.5. Revenue

As search costs fall and more buyers search, two things are happening for each type of seller: (1) they gain buyers in the middle of the distribution of values, where previous non-searchers begin to search into their auctions; but (2) they lose buyers at the bottom of the distribution of values, where previous non-searchers were getting randomly matched into their auctions, but now choose to search into the other auction.

Hence, there is a “matching effect,” of the type mentioned by Tadelis and Zettelmeyer (2011), in the sense that lower valuation buyers are being “traded” for higher valuation buyers. This is not, however, the entire story: there is no reason that the rate at which buyers leave should generically be equal to that at which they enter. As a result, a seller may lose buyers overall as a result of search, if, for whatever reason, the effect of decreasing search costs on the marginal searchers into the other auction is greater than its effect on those search into the seller’s auction. The loss in competitiveness may result in lower prices.

Figure 1 shows what happens to auction participation in an A auction as search increases, i.e., the search thresholds change from (x_b, x_a) to (x'_b, x'_a) , with $x'_b > x_b$ and $x'_a < x_a$. Recall that auction participation can be conceptualized as a Poisson process, where the density parameter is equal to the population density of participants, which we denote, again with a slight abuse of notations, using $\mu(x)$. Independence of non-overlapping sets allows us to conceptualize participation in a particularly straightforward way.




Figure 1: Effects of search on auction participation

The solid black line is the population density function, which is also the density function in an A auction when no buyers search (since for any set in X , a fraction a of them are placed into an A auction, and we normalize by the measure of A sellers, which is a). Now consider the density of buyers in an A auction under thresholds (x_b, x_a) . In $[\underline{x}, x_b)$, all buyers search into B , which means that no buyers appear in the A market, so the density is zero. Along $[x_b, x_a]$, no buyers search, so the density remains $\mu(x)$. Finally, along $(x_a, \bar{x}]$, the platform density μ of buyers search into A , and we normalize that to a sellers, so the density is $\mu(x)/\alpha$, which is depicted by the dotted line. We can partition X into intervals defined by the points x_b , x'_b , x'_a , and x_a : the number of buyers who show up from each of these intervals is an independent Poisson variable with their respective densities. Letting k_i denote a Poisson variable with measure equal to the shaded region i in the figure, total participation in an A auction can be written as $k_1 + k_2 + k_3 + k_4$.

If we move the thresholds now to (x'_b, x'_a) , two things happen: the density in $[x_b, x'_b)$ becomes zero, and the density in $(x'_a, x_a]$ changes from $\mu(x)$ to $\mu(x)/p_a$. This means that k_2 and k_4 remain the same, and k_1 is replaced by zero. k_3 is replaced with a Poisson random variable of $(1/\alpha)\mu((x'_a, x_a])$. Since the density is only scaled up, however, and Poisson variables are infinitely divisible, this is equivalent to two independent Poisson variables, one with parameter $\mu((x'_a, x_a])$ and the other with parameter $(1/\alpha)\mu((x'_a, x_a])$.⁴ Hence, participation can be summarized as $k_2 + k_3 + k_4 + k_5$.

The net effect is that we remove the Poisson draw from $[x_b, x'_b)$ with a Poisson draw from $(x'_a, x_a]$ (where the value distribution of a buyer in those regions is the conditional distribution of valuations, conditional on being within the region). Obviously, the removal

4. Alternatively, we can consider the Poisson process in $(x'_a, x_a]$ to be the sum of two independent processes, one with density $\mu(x)$ and the other with density $(1/\alpha)\mu(x)$.

of k_1 by itself always hurts revenue, and the addition of k_5 by itself always helps revenue. In general, the net effect cannot be signed, since the measure of buyers sorting in and those sorting out will differ. In the special case where the value distributions of each good are identical (i.e., the platform is symmetric), however, the effect on revenue is unambiguously positive.

Definition 1. *A platform is symmetric if permuting the labels A and B (e.g., calling the A good the B good and vice versa) does not change any of the value distributions or proportions of goods in the seller population.*

Note that the conditions for this to be true are quite specific: it must be that $a = 1/2$, $\bar{x} = y(x)$, $\underline{x} = y(\bar{x})$, $y(x) = \bar{x} - x$ and $H(x) = F(x)$ for all $x \in X$. That is, the distributions of valuations for A and B goods are identical. Under symmetry, the two equilibrium thresholds will also always be symmetric around the mean of the value distribution (since relabeling the goods does not change the fundamentals, this is implied by uniqueness). Hence, the measure of new searchers when the search cost decreases will be the same in each market; from the point of view of any seller, this means that the measure exiting is equal to the measure entering. That is, there is no net change in measure, but simply a replacement of all the lower buyer types with higher buyer types.⁵ The improved matching must lead to higher revenue:

Theorem 3 (Revenue under Symmetry). *For a symmetric platform, as search costs decrease, either through a reduction in c or an increase in p , expected revenue is increasing for all sellers.*

By considering the symmetric case, it is also easy to see how revenue might decrease: the measure of buyers leaving may be much larger than the measure of buyers entering, and so even though those entering may have higher valuations, their number may be insufficient to counteract the negative effects of decreased buyer density. Put another way, lower search costs lead to better matching — a *valuation effect* — but also to a change in market thickness — a *segmentation effect*. The segmentation effect will exert downward pressure on revenues in one market and additional upward pressure on revenues in the other. Surprisingly, it is possible that the segmentation effect can hurt revenues in the one market so much that the increased revenues in the other market fails to compensate (i.e. sellers are worse off overall).

5. In terms of the figure, k_1 and k_5 have the same parameter — since they are independent, this is equivalent to replacing each buyer in $[x_b, x'_b]$, should at least one show up, with a buyer in $(x'_a, x_a]$, which is always revenue improving.

Theorem 4 (Revenue under Asymmetry). *Overall seller revenue may decrease as search costs decrease.*

We show this by constructing a numerical counterexample. Consider a platform where $a = 0.05$, $X = [0, 1]$, $y(x) = 5 - x$, and $\mu = 3$. In this market, the relatively abundant good is also the one that people prefer, the B good, for which all buyers have valuations between two and three. Under perfectly random matching, it is also the good that most buyers will draw into. The low valuation buyers, however, have little chance of winning the good, and hence would prefer to receive the A good with a higher chance of winning and higher surpluses.

As the search efficiency p increases, the effect is most dramatic for those who would prefer to search into the A market, since that is the one that is harder to draw. When $p = 0$, most buyers would still be able to draw into the B market relatively easily, but would find it costly to draw A (the expected number of draws is 20). The expected number of draws for a buyer who wants to get in to a B auction, by contrast, is only $10/19 \approx 1.052$. A change from $p = 0$ to $p = 1$, then, represents a tenfold decrease in the effective search cost for those searching into A , and only an eleven percent decrease in search costs for those searching into B . Hence, cost decreases through p will induce relatively more buyers to search into A than into B .

Figure 2a shows the threshold values as p changes, and figure 2b shows how the measure of buyers in each market changes. The dotted lines show results for the A market, and the solid lines show results for the B market. As expected, x_b rises and x_a falls, since lower search costs induce more searching. Although the threshold x_b moves more than the threshold x_a , B loses buyers overall, which is shown in figure 2b. The reason for this is that the additional “searchers” into B would mostly have ended up in B anyway, as would the additional searchers into A : searching *into* B has relatively little impact on the measure of buyers in B , and searching *out* of B has a large impact.

These changes in participation are reflected in the market revenues, which are shown in figure 2c. Revenue in market B drops, since B auctions are overall losing buyers, and in this case the valuations of the buyers that they gain are relatively low. In contrast, A revenues increase. Finally, figure 2d shows *total* revenues. Although the A revenues do increase, there are relatively few A sellers, so the effect on total revenues is dominated by the decrease in B revenues.

Our counterexample was constructed around a particular type of asymmetry: the existence of a small market with lower overall valuations and relatively low per-draw search

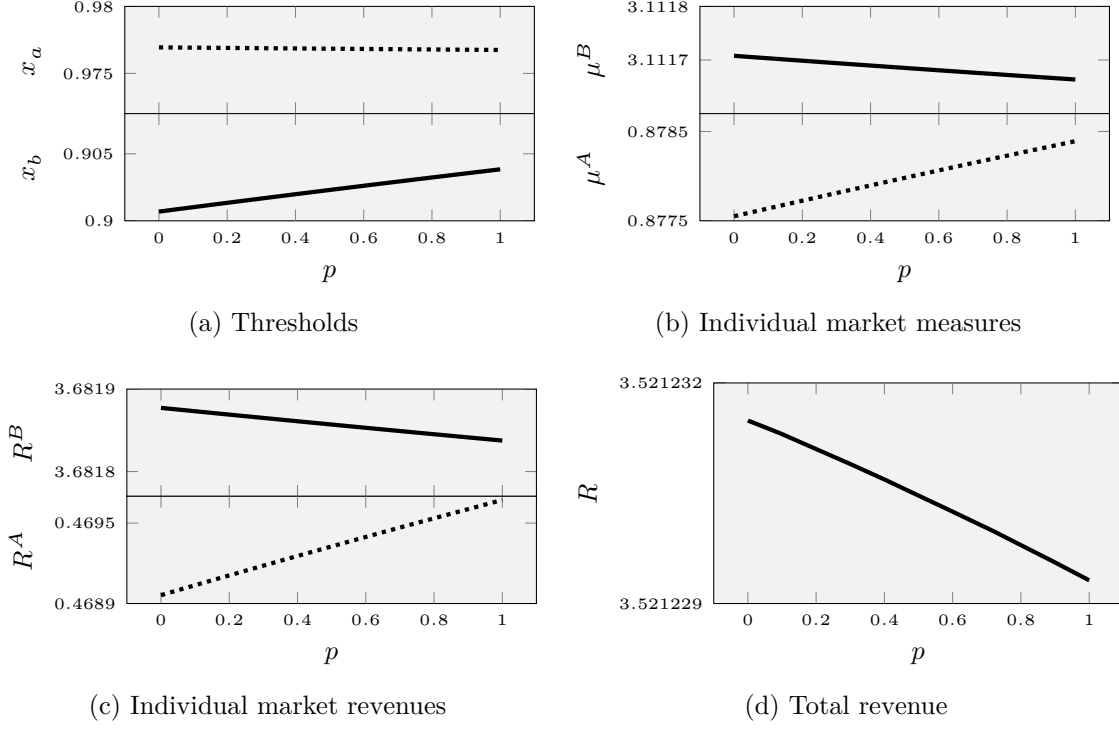


Figure 2: Increasing p under asymmetric conditions

costs. We have run a number of numerical simulations, which seem to indicate that these counterexamples are relatively rare. Overall revenue generally increases as search costs fall, except when there are significant asymmetries. Numerical results are similar when, instead of varying the search success probability p we directly vary the search cost c .

Although our numerical investigations do suggest that negative overall revenue effects are “rare”, it should be kept in mind that we have constructed our model as a “best case” scenario in terms of generating positive revenue effects as a result of matching: buyer valuations for each good are deterministically negatively correlated with their valuations for the other good. In more general settings, the negative revenue effects can only be exacerbated, since there would no longer be a guarantee that the buyers leaving a market as a result of search are always those with the lowest valuations. The fundamental lesson in terms of revenue is that lowering search costs affects revenue in two ways, and that the net effect is not a Pareto improvement: except under perfect symmetry, some sellers will lose.

5. Conclusion

Recent years have seen substantial investments by platforms in reducing search costs. Yet, there has been relatively little work on understanding what the implications of decreased search costs are for the parties involved. The continuum framework developed here simplifies auction participation in a large network to a Poisson process, should help to overcome many of the technical barriers associated with analyzing platforms with large numbers of agents on both sides of the market in situations where search matters. We have also provided positive results to shed some light into the effects of search on platform welfare and revenue, which may help to explain and guide platform policy related to search investments. In particular, in a second-price auction platform – and by revenue equivalence, in all efficient single-good auction mechanisms (Riley and Samuelson 1981) – total social welfare is increasing as search costs decrease. Moreover, in a symmetric platform, all sellers will benefit from search, so that total seller revenues and individual seller revenues will increase as search costs decrease. If platform interests are aligned with seller interests, then (as would be the case if the sellers were being charged, for instance), then a platform would want to invest in decreasing search costs in symmetric situations. On the other hand, we have also shown that in a generic setting, there are distributional consequences to increased search: one set of sellers will lose, while the other gains, with ambiguous effects on total revenue. A formal examination of platform incentives will require a richer model, which we hope to develop in future work.

References

- Ambrus, A., and R. Argenziano. 2009. “Asymmetric Networks in Two-Sided Markets.” *American Economic Journal: Microeconomics* 1 (1): 17–52.
- Armstrong, M. 2006. “Competition in two-sided markets.” *RAND Journal of Economics* 39 (3): 668–693.
- Diamond, P. A. 1971. “A model of price adjustment.” *Journal of Economic Theory* 3 (2): 156–168.
- Ellison, G., and D. Fudenberg. 2003. “Knife-Edge or Plateau: When Do Market Models Tip?” *The Quarterly Journal of Economics* 118 (4): 1249–1278.

- Gale, D., and H. Nikaido. 1965. "The Jacobian matrix and global univalence of mappings." *Mathematische Annalen* 159 (2): 81–93.
- Hagiu, A. 2009. *Multi-sided platforms: From microfoundations to design and expansion strategies*. Technical report. Harvard Business School.
- Hagiu, A., and B. Jullien. 2011. *Search diversion, rent extraction and competition*. 11-124. Technical report. HBS.
- Hagiu, A., and R. Lee. 2011. "Exclusivity and Control." *Journal of Economics and Management Strategy* 20 (3): 679–708.
- Halaburda, H., and Y. Yehezkel. 2011. "Platform Competition under Asymmetric Information."
- Klemperer, P. 2004. *Auctions: Theory and Practice*. Princeton University Press.
- Krishna, V. 2009. *Auction Theory*. Academic Press.
- McAfee, R. P., and J. McMillan. 1987. "Auctions with entry." *Economics Letters* 23 (4): 343–347.
- Mortensen, D. T. 1970. "Job Search, the Duration of Unemployment, and the Phillips Curve." *The American Economic Review* 60 (5): 847–862.
- Myerson, R. 2000. "Large Poisson Games." *Journal of Economic Theory* 94:7–45.
- Riley, J. G., and W. F. Samuelson. 1981. "Optimal auctions." *The American Economic Review* 71 (3): 381–392.
- Rochet, J., and J. Tirole. 2006. "Two-sided markets: a progress report." *The RAND Journal of Economics* 37 (3): 645–667.
- Tadelis, S., and F. Zettelmeyer. 2011. "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment." *SSRN eLibrary*.
- Weyl, G. 2010. "A Price Theory of Multi-Sided Platforms." *American Economic Review* 100 (4): 1642–1672.
- Weyl, G., and A. White. 2012. "Insulated Platform Competition."

A. Appendix

This appendix includes the full proofs of details omitted over the course of the main text.

Proof of Lemma 1 (Independence). Let n denote the number of buyers who shows up in a given auction, so $n \sim \text{Poisson}(\mu)$. For any given $n \geq k$, the number with valuations in some measurable set X_1 , $k_1 | n$, is binomially distributed, since there are n buyers, each with an independent probability of being in X_1 of probability $p(X_1)$. We know n is Poisson distributed, so the prior probability of exactly k buyers from X_1 is given by

$$\begin{aligned} P(k_1 = k) &= \sum_{n=k}^{\infty} P(n) \cdot P(k_1 = k | n) = \sum_{n=k}^{\infty} \frac{e^{-\mu} \mu^n}{n!} \binom{n}{k} p(X_1)^k (1 - p(X_1))^{n-k} \\ &= \frac{e^{-\mu(X_1)} \mu(X_1)^k}{k!} \end{aligned}$$

To see independence, consider the probability that $k_1 = l$ and $k_2 = m$, again letting n denote the total number of buyers in an auction. Any given buyer is in X_1 with probability $p(X_1)$, in X_2 with probability $p(X_2)$, and in neither with probability $1 - p(X_1) - p(X_2)$. Given n , The number of buyers in each of these groups – X_1 , X_2 , and X ($X_1 \cup X_2$) – is multinomially distributed, so

$$P(k_1 = l \text{ and } k_2 = m | n) = \frac{n!}{k! l! (n - k - l)!} p(X_1)^l p(X_2)^m (1 - p(X_1) - p(X_2))^{n-l-m}$$

Taking the probability weighted sum over all n and rearranging gives the desired expression, $P(k_1 = l \text{ and } k_2 = n) = P(k_1 = l) \cdot P(k_2 = n)$. \square

Proof of Lemma 2. These expressions come from standard auction theory results — the only difference in our case is that the distribution of the order statistics *faced by sellers* is one further step removed from the buyer value distributions. In a second-price auction, buyers bid their valuations, so the distribution of the first-order statistic, $G^j(z)$ is also the probability that z wins the good. Hence, the surplus of buyer z is $\int_0^z G^j(s) ds$ (e.g., Klemperer 2004, p. 41).

To find $G^j(s)$, let $\nu^j(s) = \mu^j(1 - F^j(s))$ be the measure of buyers with valuations above s . Then the probability that s is greater than the valuations of all buyers in an auction is simply the probability that the number of buyers with valuations above s in an auction is zero, and by independence, that number is Poisson distributed with parameter $\nu^j(s)$; hence, $G^j(s) = e^{-\nu^j(s)}$. Expected revenue is equal to the expectation of the second highest

bid, which we denote Z_2^j . The cumulative distribution function for $Z_2^j - G_2^j(z)$ — can be calculated as the probability that zero or one buyer shows up from the subset $(z, \bar{z}]$:

$$G_2^j(z) = G^j(z) + \mu(1 - F^j(z))G^j(z)$$

which has density

$$\begin{aligned} g_2^j(z) &= g^j(z) + [-\mu^j f^j(z)G^j(z) + \mu^j(1 - F^j(z))g^j(z)] \\ &= \mu^j(1 - F^j(z))g^j(z) \end{aligned}$$

The expectation can then be directly calculated

$$E[Z_2] = \int_0^{\bar{z}} z g_2^j(z) dz = m^j(\bar{z}) - \mu^j \int_0^{\bar{z}} (1 - F^j(z))G^j(z) dz$$

□

Proof of Theorem 1. We prove our main theorem in two parts. First, we show existence with Brouwer's fixed point theorem. We can characterize equilibria by ordered threshold pairs $(x_b, x_a) \in X^2$ — such that $x < x_b$ search into B , $x > x_a$ search into A and the rest do not search — and let $T : X^2 \rightarrow X^2$ be a mapping from a pair of buyer thresholds to a new pair that describes those buyers indifferent between searching (either into A or B) and not searching. Since the buyer distribution is atomless and utilities vary continuously and monotonically with type, this mapping must also be continuous. The set X^2 is compact, so Brouwer's theorem ensures the existence of a fixed point, which is an equilibrium.

Uniqueness and monotonicity are slightly more involved. First we introduce a bit of notation. Let $\mathbf{x} = (x_b, x_a)$, and define $s(x; \mathbf{x}) = u^A(x) - u^B(x)$ be the difference between A and B market utilities for a type x buyer — since this depends on the behaviors of other buyers, we parameterize this by the thresholds. Then let $\mathbf{S}(\mathbf{x}) \equiv (s(x_b; \mathbf{x}), s(x_a; \mathbf{x}))$ give the market differences to the threshold types.

By Lemma 3, equilibria are characterized by $\mathbf{S}(\mathbf{x}) = (-c/p_b, c/p_a) = \mathbf{c}$. Below we show that the Jacobian of \mathbf{S} is everywhere positive and all principal minors are non-vanishing. Then by Gale-Nikaido, the mapping is one-to-one and hence $\mathbf{S}(\mathbf{x}) = \mathbf{c}$ has a unique solution. Moreover, by Cramer's rule, we can sign the terms of the inverse of the Jacobian matrix. Applying the implicit function theorem to $\mathbf{S}(\mathbf{x}) = \mathbf{c}$ and substituting in the known signs of the inverse Jacobian matrix gives monotonicity. □

Lemma 5 (Useful expressions). *Let F^A and F^B denote the buyer distributions in each market, and let G^A and G^B denote the distributions of the first-order statistics in the respective auctions. Then*

$$G^A(x) = e^{-\mu^A(1-F^A(x))} \quad \text{and} \quad G^B(y) = e^{-\mu^B(1-F^B(y))}$$

where

$$F^A(x) = \begin{cases} \frac{\mu}{\mu^A} (F(x) - F(x_b)) & \text{for } x_b \leq x < x_a \\ \frac{\mu}{\mu^A} \left[\frac{1}{\alpha} F(x) - \left(\frac{1}{\alpha} - 1 \right) F(x_a) - F(x_b) \right] & \text{for } x_a \leq x < \bar{x} \end{cases}$$

$$F^B(y) = \begin{cases} \frac{\mu}{\mu^B} (H(y) - H(y_a)) & \text{for } y_a \leq y < y_b \\ \frac{\mu}{\mu^B} \left[\frac{1}{1-\alpha} - \left(\frac{1}{1-\alpha} - 1 \right) H(y_b) - H(y_a) \right] & \text{for } y_b \leq y < \bar{y} \end{cases}$$

where

$$\mu^A = \frac{1}{\alpha} - \left(\frac{1}{\alpha} - 1 \right) F(x_a) - F(x_b) \quad \mu^B = F(x_a) + \frac{\alpha}{1-\alpha} F(x_b)$$

$$\bar{y} = y(\underline{x}) \quad \underline{y} = y(\bar{x}) \quad y_b = y(x_b) \quad y_a = y(x_a)$$

Proof. As before, we will prove the result for the A market. The total measure of buyers in the A market, normalized by sellers, is given by μ^A , since

$$\mu^A = \mu(X_0) + \frac{1}{\alpha} \mu(X_a) = \mu \left[\frac{1}{\alpha} - \left(\frac{1}{\alpha} - 1 \right) F(x_a) - F(x_b) \right]$$

Let $\nu(x)$ be the measure of buyer above x for $x \in [x_b, \bar{x}]$, then

$$\nu(x) = \begin{cases} \frac{\mu}{\alpha} [1 - F(x)] & \text{for } x > x_a \\ \mu [F(x_a) - F(x)] + \frac{\mu}{\alpha} [1 - F(x_a)] & \text{for } x \leq x_a \end{cases}$$

The probability that x is above the highest bid in a given auction is the probability that zero buyers with valuation $x' > x$ shows up, or $e^{-\nu(x)}$.⁶ \square

6. Note also the role that independence plays in our analysis: since the number of buyers showing up from each interval is independent, the above is equivalent to thinking of A auctions as a single-good platform where the total measure of buyers is μ^A and valuations are independently and identically distributed according to

Proof of Theorem 1: Uniqueness and Monotonicity. To complete the proof, we show that the Jacobian of \mathbf{S} is positive, and then use that fact to sign the inverse of the Jacobian matrix. Since $u^A(x) = \int_0^x G^A(z) dz$, and changes to either threshold, x_b or x_a , only affect the winning probabilities of buyers with types below those thresholds. The derivatives of u^A and u^B with respect to the threshold values will follow almost immediately from the derivatives of G^A and G^B with respect to x_b and x_a . Again letting $\nu(x)$ be the measure of buyers in A markets with valuations above x , and taking the derivatives with respect to x_b and x_a yields

$$\begin{aligned}\frac{d\nu(x)}{dx_b} &= -\mu f(x_b) \text{ for } x \leq x_b \text{ and } 0 \text{ otherwise} \\ \frac{d\nu(x)}{dx_a} &= -\mu \left(\frac{1}{\alpha} - 1 \right) f(x_a) \text{ for } x > x_a \text{ and } 0 \text{ for } x < x_a\end{aligned}$$

Both inequalities are strict in the case of $d\nu(x)/dx_a$ since the derivative at precisely x_a depends on the direction in which x_a is changing.

From this, we can see how G^A changes with x_b and x_a :

$$\begin{aligned}\frac{dG^A(x)}{dx_b} &= \begin{cases} \mu f(x_b) G^A(x_b) & \text{for } x \leq x_b \\ 0 & \text{otherwise} \end{cases} \\ \frac{dG^A(x)}{dx_a} &= \begin{cases} \left(\frac{1}{\alpha} - 1 \right) \mu f(x_a) G^A(x) & \text{for } x < x_a \\ 0 & \text{for } x > x_a \end{cases} \quad (\ddagger)\end{aligned}$$

Using these expressions in u^A , we get

$$\begin{aligned}\frac{d}{dx_b} u^A(x_b) &= G^A(x_b) + \mu f(x_b) u^A(x_b) \\ \frac{d}{dx_a} u^A(x_b) &= \frac{1-p}{p} \mu f(x_a) u^A(x_b) \\ \frac{d}{dx_b} u^A(x_a) &= \mu f(x_b) u^A(x_b) \\ \frac{d}{dx_a} u^A(x_a) &= G^A(x_a) + \frac{1-p}{p} \mu f(x_a) u^A(x_a)\end{aligned}$$

Utilities on the B side can be expressed equivalently using B valuations (i.e., writing utilities

the population frequencies (since F^A is simply the measure function scaled by μ^A , the total measure.

as functions of y_b and y_a) and B densities; the derivatives with respect to x will involve an additional y' term. Further simplifications can be made by observing that $h(y)|y'(x)| = f(x)$, which we substitute in to the expressions for $ds(\cdot)/dx_i$ to get:

$$\begin{aligned}\frac{d}{dx_b}s(x_b) &= G^A(x_b) + |y'(x_b)|G^B(x_b) + \mu f(x_b) \left[u^A(x_b) + \frac{p}{1-p}u^B(x_b) \right] \\ \frac{d}{dx_a}s(x_b) &= \mu f(x_a) \left[\frac{1-p}{p}u_a(x_b) + u_b(x_a) \right] \\ \frac{d}{dx_b}s(x_a) &= \mu f(x_b) \left[\frac{p}{1-p}u_b(x_a) + u_a(x_b) \right] \\ \frac{d}{dx_a}s(x_a) &= G^A(x_a) + |y'(x_a)|G^B(x_a) + \mu f(x_a) \left[\frac{1-p}{p}u^A(x_a) + u^B(x_a) \right]\end{aligned}$$

The Jacobian is equal to

$$\frac{ds(x_b)}{dx_b} \frac{ds(x_a)}{dx_a} - \frac{ds(x_b)}{dx_a} \frac{ds(x_a)}{dx_b} \quad (*)$$

This is strictly positive if $ds(x_b)/dx_b > ds(x_a)/dx_b$ and $ds(x_a)/dx_a > ds(x_b)/dx_a$. Using K_b and K_a to stand in for large repeated terms above, we find:

$$\begin{aligned}\frac{d}{dx_b}s(x_b) &= K_b + \mu f(x_b) \left[u^A(x_b) + \frac{p}{1-p}u^B(x_b) \right] \\ &\geq K_b + \mu f(x_b) \left[u^A(x_b) + \frac{p}{1-p}u^B(x_a) \right] = K_b + \frac{ds(x_a)}{dx_b} \\ \frac{d}{dx_a}s(x_a) &= K_a + \mu f(x_a) \left[\frac{1-p}{p}u^A(x_a) + u^B(x_a) \right] \\ &\geq K_a + \mu f(x_a) \left[\frac{1-p}{p}u^A(x_b) + u^B(x_a) \right] = K_a + \frac{ds(x_b)}{dx_a}\end{aligned}$$

where K_b and K_a are strictly positive. Substituting this into equation (*) above immediately delivers the strictly positive Jacobian. Finally, a strictly positive Jacobian implies that the inverse of the Jacobian matrix is signed as follows (Cramer's rule):

$$\text{sgn}(J^{-1}) = \begin{pmatrix} + & - \\ - & + \end{pmatrix}$$

By the implicit function theorem, $d(x_b, x_a)/dc = J^{-1} \cdot d(-c/p_b, c/p_a)/dc$, so $dx_b/dc < 0$ and $dx_a/dc > 0$, as claimed. \square

Proof of Lemma 4. Suppose not. Consider the set of individuals that is assigned to sort into A and those that are assigned to sort into B (we can consider any sets of two possible strategies and the analysis follows similarly). If the optimal assignment does not have the structure described in the lemma, then we can find some set X^A of buyers who sort into A , and another set X^B of buyers who sort into B , each of measure $\epsilon > 0$, such that $\inf X^B > \sup X^A$. That is, there is a set of buyers who sort into B each with strictly higher x than another set of buyers who sort into A of same measure.

Consider what happens if we change the assignments of X^A and X^B ; that is, we construct an alternative assignment that leaves all other search decisions the same, but tells those in X^A to search into B and those in X^B to search into A . Since X^A and X^B have the same measure, the alternative assignment results in the same measure of bidders in each market. Furthermore, the value distribution in each auction first order stochastically dominates the value distribution in the original assignment, since the change is to have replaced a measure of low valuation buyers with an equal measure of higher valuation buyers. This means that the distribution of the first-order statistics in the alternative distribution also first-order stochastically dominate the first-order statistic distributions of the original assignment, since

$$G^A(x) = e^{-\mu^A(1-F^A(x))}$$

and μ^A stays the same in the alternative assignment. This means that $E[x_1]$ and $E[y_1]$ are greater under the alternative assignment, and search costs remain the same, so total welfare is raised by switching strategies. This contradicts the original assumption that the allocation was socially optimal. \square

Proof of Theorem 2. The expression for welfare is

$$W = a \left(\bar{x} - \int_0^{\bar{x}} G^A(x) dx \right) + (1 - \alpha) \left(\bar{y} - \int_0^{\bar{y}} G^B(y) dy \right) - \left(\frac{(1 - \alpha)\mu^A}{p_a} + \frac{a\mu^B}{p_b} \right) c$$

Taking the derivative of this with respect to x_b and x_a , using expressions for dG^j/dx_i from equation (‡) above in the uniqueness proof, gives

$$\begin{aligned} \frac{dW}{dx_b} &= \alpha \mu f(x_b) \left[- \int_0^{x_b} G^A(x) dx + \int_0^{y_b} G^B(y) dy - \frac{c}{p_b} \right] \\ &= \alpha \mu f(x_b) \left[- (u^A(x_b) - u^B(x_b)) - \frac{c}{p_b} \right] \end{aligned}$$

$$\begin{aligned}\frac{dW}{dx_a} &= (1 - \alpha)\mu f(x_a) \left[-\int_0^{x_a} G^A(x) dx + \int_0^{y_b} G^B(y) dy - \frac{c}{p_a} \right] \\ &= (1 - \alpha)\mu f(x_a) \left[- (u^A(x_a) - u^B(x_a)) - \frac{c}{p_a} \right]\end{aligned}$$

The terms in brackets are easily signed: when x_b is below the threshold of the market solution, dW/dx_b is positive since $u^A(x_b) - u^B(y_b) < c/p_b$, and similarly, when x_a is above the market solution threshold, dW/dx_a is negative. A quick inspection shows that the first order conditions for the social planner are equivalent to the threshold equilibrium conditions, so we have ($x_b^* = x_b$ and $x_a^* = x_a$, i.e., the social optimum coincides with the market solution). Hence, the search equilibrium is efficient in the sense that it maximizes total welfare. \square

B. Numerical Implementation (not for publication)

This section contains details regarding our implementation of the numerical results. The data was generated using code written in C++; performance and accuracy considerations ruled out the use of Matlab or Mathematica, whose built-in optimization routines performed poorly for our setting. To avoid rewriting widely used mathematical operations, we used two third-party open-source libraries: GSL for numerical integration and NLOpt for minimization. Both libraries have been extensively tested for correctness, so we can be confident of our results even for small magnitudes.

The implementation is quite straightforward: for sets of parameters (x_b, x_a) , we define utility functions, whose forms are given in the main text, and an error function, $\epsilon(x_b, x_a)$ as

$$\epsilon(x_b, x_a) = \epsilon_b(x_b, x_a) + \epsilon_a(x_b, x_a) \quad (1)$$

where

$$\epsilon_b(x_b, x_a) = \begin{cases} |s_b(x_b, x_a) - c_b| & \text{for } x_b > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

and

$$\epsilon_a(x_b, x_a) = \begin{cases} |s_a(x_b, x_a) - c_a| & \text{for } x_a < 1 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

The error function is minimized via the SBPLX algorithm, and the numerical minimization procedure is terminated when $\epsilon < 10^{-9}$.

C. General Sequential Search (not for publication)

This section describes a more general version of our search framework. In particular, we allow for a continuum of types on *both* sides of the market, which allows for greater flexibility and may better accommodate other sequential search settings, such as job search.

We focus on describing match outcomes *given search decisions*. Individual optimal behavior may vary greatly depending on the context and allocation mechanisms involved; however, our framework provides a general and powerful way to describe search outcomes, which are often necessary to characterize agent utilities.

The set of buyer types is X , Borel-measurable, with measure function $M(\cdot)$ and density $m(\cdot)$; the set of seller types is Y with measure $N(\cdot)$, density $n(\cdot)$, and total measure 1. Utilities depend only on types, so we assume that buyer actions are completely determined by their types. Following the sequential search literature, we can characterize search decisions by acceptance sets $y(x)$ for each x — that is, a buyer of type x keeps searching until he draws some $y \in y(x)$, at which point he stops. We also let $|z|$ denote the measure of a set z .

Lemma 6. *For a seller of type y , the density of buyers he faces is given by*

$$f_y(x) = \frac{m(x)}{|y(x)|}$$

where $y(x)$ is the set of all sellers that x buyers accept.

Sketch of proof. Buyers of type x search into $y(x)$, where the relative density of y sellers is $n(y)/|y(x)|$. Hence the density of x buyers is $m(x)n(y)/|y(x)|$, and the relevant density for an individual seller is the buyer-seller ratio, so we divide by $n(y)$ to get the result. \square