Spatial Search*

Xiaoming Cai[†] Pieter Gautier[‡] Ronald Wolthoff[§]

November 29, 2023

Abstract

This paper considers a random search model where some locations provide sellers with better chances of meeting many buyers than other locations (for example popular shopping streets or the first page of a search engine). When sellers are heterogeneous in terms of the quality of their product and/or the probability that a given buyer likes their product, it is desirable that sellers of high-quality niche products sort into the best locations. We show that this does not always happen in a decentralized market. Finally, we allow for endogenous location distributions and show that more trades are realized when locations are similar (in which case the aggregate matching function is urn-ball) but that quality weighted trade can be higher when locations are heterogeneous.

JEL classification: C78, D44, D83.
Keywords: search frictions, spatial equilibrium, sorting

*We thank Sephorah Mangin for insightful comments. Wolthoff gratefully acknowledges financial support from the Social Sciences and Humanities Research Council.

[†]Peking University HSBC Business School. E-mail: xmingcai@gmail.com.
[‡]VU University Amsterdam and Tinbergen Institute. E-mail: p.a.gautier@vu.nl.
[§]University of Toronto. E-mail: ronald.p.wolthoff@gmail.com.

1 Introduction

Good locations for sellers are locations where it is easy to meet buyers. They are often located in densely populated areas (Fifth Avenue in New York, Bond Street London) or they are easily reachable (good parking facilities or near a metro station). For online products, we can think of good locations as being on a top position of a search engine. In this paper, we consider a random search model that takes into account that some locations are better than others. The model offers a framework to study what type of sellers benefit from good locations and how this shapes spatial sorting, how the price of locations is determined and from an urban planning perspective, when is it desirable to make locations similar (in terms of meetings) and when it is better to have heterogeneity in location quality.

In our model, sellers are characterized by the quality of their product, denoted by z, and the probability that a buyer likes their product, which is assumed to be weakly decreasing in z and is denoted by $x(z)$. Furthermore, we assume for simplicity that sellers have one good for sale, but what matters is only that there is some capacity constraint.[1] Locations differ in how many buyers there are per seller, taking into account a buyer-resource constraint. This implies that if we improve the meeting rate for sellers in one location (for example by adding a railway station), the sellers in other locations will meet fewer buyers. In each location, we have a constant-returns-to-scale Poisson meeting technology, but the queue length varies across locations. Formally, we model seller locations as points on a unit circle where uniformly distributed buyers move clockwise to the nearest location. A good location for sellers is then one which is far away from their nearest competitor in a counterclockwise direction.

We start with a simple environment where the location of sellers is exogenous, and let sellers sort into their optimal location, given a competitive rental market. We show that whenever the expected consumer value $zx(z)$ is increasing in z, the planner would always prefer to match good locations with high-quality sellers, i.e., positive assortative matching (PAM) between sellers and locations is desirable. Next, we consider the decentralized market equilibrium. We show that the requirement for PAM is more stringent and a *sufficient* condition is that $zx^2(z)$ is increasing in z, since sellers enjoy a high payoff when there are two or more buyers who want to buy their product. Good locations are particularly valuable for high-quality goods that few buyers like but those who like it, like it a lot. This is consistent with the designer shops that are located on Fifth Avenue and the fact that most of the Michelin three-star restaurants are located in or near big cities.[2]

[1] A second-hand car seller with five cars of a particular type would like to meet at least six buyers to be able to be on the short-side of the market. Here a seller with one unit would like to meet at least two buyers and for both cases this is more likely to happen in a good location than in a bad location.

[2] Another example comes from the following Financial Times article that points out that most manufacturing has moved out of cities, except high-quality niche products that are produced in small quantities like micro-breweries, furniture makers, roboticists and 3D-printing specialists.

We then take an urban planning or regional policy perspective and consider the optimal distribution of locations given the distribution of seller types. We find that although the total number of meetings is maximized when sellers are located at equidistance, this is not always the welfare-maximizing topography. For example, if buyers are located randomly and uniformly on the circle and a small fraction of them has a strong desire for top quality food, then there is no area where it is profitable for top restaurants to enter and even if we would force them to enter, they would create little surplus because they meet relatively few buyers. In contrast, when there is heterogeneity in locations, sellers of high-quality and/or niche products are able to create a lot of value by locating in good spots. So heterogeneous locations can generate more welfare than equidistant locations by allowing for sorting of seller types. This results in fewer trades overall but more high-quality trades.

Interestingly, even though we have random search, heterogeneity in locations can still create heterogeneity in expected numbers of buyers per seller as in directed search models. Random search is relevant for settings where full ex-ante commitment is not possible. For example, a seller may announce a positive reserve price ex ante but ex post after one or more buyers visit, the seller has no incentive to reject values below the reserve price but above the seller's valuation. When search is directed, a seller can increase the expected number of buyers by offering a good deal to the buyers whereas here they can select a good location. The difference is that here, with random search, the price for this good location does not go to the buyers.

When we make the location distribution endogenous, we can think of regional policy in our framework as choosing the optimal topography which must strike a balance between maximizing trade and allowing high-quality sellers to sort into good locations. We show that the optimal topography mimics the directed search equilibrium, which equalizes buyers' marginal contribution to surplus in different submarkets. In contrast, in the random search equilibrium with an endogenous location distribution, high-quality sellers typically overinvest in good locations for rent seeking reasons (they are willing to pay higher rents to receive two or more buyers rather than one). This is socially wasteful because a seller who meets two effective buyers creates the same surplus as a seller who meets one effective buyer. Moreover, a seller who invests in a good location and meets many buyers does not internalize that in other areas more sellers will meet no buyers at all.

Our model is consistent with the finding in Neiman and Vavra (2023) that niche consumption is largest in areas with many buyers per seller like Chicago, Washington DC, Tampa, Los Angeles and Boston and lowest in non-dense, isolated places like Des Moines, Little Rock, Las Vegas, and "West Texas". In the context of the labor market, Gautier and Teulings (2003) create an index that captures per CMSA how many workers are available per job in an area. When many workers are available per job, wages and the cost of living are higher.

Similarly, our model also predicts that rents are higher in good locations with many buyers per seller.

In most of this paper, the quality distribution is exogenous.[3] Menzio (2023) endogenizes the quality distribution by letting ex-ante homogeneous sellers choose their quality in a dynamic Burdett-Judd (1983) framework. He allows search frictions to decline over time. In response, firms offer more specialized products with a higher consumer value. In this environment, it is possible that, despite the decline in search frictions, the economy exhibits a balanced growth path where price dispersion and the extent of competition remain constant. This paper shares the observation that specialized sellers need traffic more than generic ones but it is complementary to Menzio (2023) in the sense that we study what happens when there is cross-sectional rather than time variation in search frictions. This adds a location choice and sorting dimension to the firm's problem. Albrecht et al. (2023) assume that firms' key design choice is vertical rather than horizontal (as it was in Menzio, 2023). In their model, high-quality sellers have higher trading probabilities because buyers visit multiple sellers and high-quality sellers can offer more surplus to buyers than low-quality sellers. In our model, buyers can only visit one seller but sellers can choose a good location to increase their trading probability at the cost of a higher rental price.

There exists a small literature that relates spatial sorting to search frictions. In Helsley and Strange (1990), match quality is higher in large urban areas because workers are more likely to find a job that matches their skills when there are many firms available. They also look at the optimal number of workers and firms to locate in a city subject to a population constraint and find that it is optimal to have cities of equal size. We find the same when products have the same quality. However, when (vertical) quality differs across sellers and when we allow for niche products, the optimal topography involve heterogeneous areas.

Gautier and Teulings (2009) assume a meeting function with increasing returns to scale, which makes large urban areas more efficient search markets. They find that workers with rare skills and firms that need to hire a wide variety of skills benefit most from dense labor markets. Combes et al. (2008) and Dauth et al. (2022) show that high-skilled workers sort into dense areas in respectively France and Germany. Kim (1989) does not have search frictions but in his model, workers do specialize more (rather than invest in general skills) in large markets because the fewer firm types there are, the less likely it is that one of them will demand a particular skill. In the context of the marriage market, Gautier et al. (2010) find that highly-educated singles (more so than couples or singles with less education) locate in big cities to find a partner because the opportunity cost of remaining unmatched are largest for them. Here, we have constant-returns-to-scale and many-on-one meetings. Sellers of

[3] We do discuss how a fixed production cost interacts with the distribution of locations. That is, certain niche products will only be offered when location quality is sufficiently dispersed.

high-quality and niche products create a lot of surplus conditional on matching with a buyer and they locate in areas where the probability that they meet multiple buyers is large. In marriage market terms, our model implies that niche types (types that belong to a niche subculture) sort into large cities.

Pissarides (2000) allows firms and workers to invest in search and recruitment intensity. This is modeled in a reduced form way as a scalar that increases the individual matching rate. At the aggregate level, if firms and workers double their search intensity, their meeting rates are also doubled. In our model, a seller can increase its meeting rate by moving to a good location where there are many buyers per seller but if some locations are better than average, it implies that other locations are worse than average (due to the buyer resource constraint). Since the price of a location is endogenous in our model, so is the price for a higher meeting rate whereas in reduced form models, the cost of increasing search intensity is exogenous.

In the urban economics literature, location choice is driven by the trade-off between positive agglomeration effects which give rise to increasing returns and mobility cost. The first is needed to explain why large cities exist and the second why not all jobs are in one large city.[4] This literature is mostly complementary to this paper. We have not much to say about the size distribution of cities. Our aim is to understand what types of sellers locate in attractive areas characterized by high buyer-seller ratios, why buyer-seller ratios differ across space, how this shapes the quality distribution of products that are offered and whether heterogeneity in location quality is desirable or not.

The paper is organized as follows. In section 2, we introduce the model and define equilibrium. In section 3, we characterize the model for a given location distribution. We first consider homogeneous and then heterogeneous sellers. For the latter case, we analytically solve the model for two examples and then use a first-order approximation to study the effects of making goods more niche. Then, in section 4, we endogenize the distribution of locations and ask from an urban-planning point of view what the optimal distribution of locations is for a given quality distribution and what the market outcome would be if sellers form a coalition which chooses the seller-optimal distribution of locations, as in a real estate investment trust (REIT). Finally, section 5 concludes.

2 The Model

Agents and Preferences. Consider a static environment with a measure 1 of sellers and an endogenous measure λ of buyers, determined by free entry at cost $K > 0$. Both buyers

[4] See for example Ellison and Glaeser (1997), Fujita and Thisse (2002), Rosenthal and Strange (2004) Ellison et al. (2010), Moretti (2012).

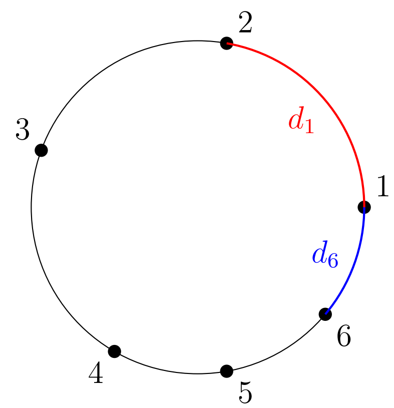

Figure 1: Finite number of sellers

and sellers are risk neutral. Each seller possess a single unit of an indivisible good and each buyer has (inelastic) demand for one unit. Sellers are heterogeneous with respect to the good that they offer. Goods are characterized by their quality z and the corresponding probability $x(z)$ that a buyer likes the good. The distribution of z among sellers is given by a cumulative distribution $F(z)$. We assume that $x(z)$ is weakly decreasing in z, while $zx(z)$, the expected value of good z to a buyer, is weakly increasing in z. We can think of goods with a low value of x (or equivalently, a high value of z) as niche products. Buyers and sellers who do not trade obtain a zero payoff.

Search. Search is random and occurs in space where we allow for heterogeneity in locations. Sellers in good locations meet relatively many buyers. We can think of the good locations as places that are easily reachable, for example because they are close to a station or a highway. Alternatively, they could be in a popular shopping street or mall in the city center. We model locations as points on a circle with circumference 1.

To better explain the search process, we first consider the case where the number of sellers (N_s) and buyers (N_b) are finite, which we illustrate in Figure 1. All sellers randomly arrive on the circle according to a probability distribution that we describe below. After that, buyers arrive on the circle and go clockwise to the nearest seller. We assume that buyers are placed uniformly on the circle, which, as we will see later, is a normalization rather than an assumption.

From the perspective of seller i, the quality of their location depends on the arc distance to their counterclockwise neighbor, denoted by d_i. In Figure 1, d_i takes one of two values: half of all sellers have good spots and the other half have bad spots, where the arc length of a good spot is two times that of a bad spot. Hence, $d_1 = 2d_6 = 2/9$. Since each buyer visits seller i with probability d_i, the probability that seller i meets n buyers is given by

$$\binom{N_b}{n} d_i^n \left(1 - d_i\right)^{N_b - n}. \tag{1}$$

Below, we will allow for a continuum of buyers and sellers with $N_b \to \infty$ and $N_b/N_s \to \lambda$. We can write the expected buyer-seller ratio at location i as $N_b d_i \to \lambda s_i$ where $s_i = d_i N_s$. Equation (1) then converges to $e^{-\lambda s_i}(\lambda s_i)^n/n!$, that is, a Poisson distribution with mean λs_i. So in a large market, seller i's spot is characterized by s_i and we can think of good locations as locations where λs_i is large. The advantage of using s_i as a measure of location quality (for a given N_s) rather than d_i is that $\sum_{i=1}^{N_s} s_i/N_s = 1$, i.e., the mean of s_i is 1 by construction, while the mean of $d_i \to 0$ when we let the market get large. In Figure 1, the distribution of s_i is $\mathbb{P}(s_i = 4/3) = 1/2$ and $\mathbb{P}(s_i = 2/3) = 1/2$.

More generally, the distribution of locations can be described by a cdf $L(s)$. The probability distribution for the event that a buyer meets a seller at a location s is $s\,dL(s)$.[5] A few special cases are worth mentioning (see Section 3.1 for more details). When sellers are located at equidistance from each other, $L(s)$ is degenerate at $s = 1$, and the number of buyers that sellers meet follows a Poisson distribution with mean λ. In Appendix A.1, we show that if sellers are placed uniformly on the circle (just like the buyers), then the distribution of locations converges to an exponential distribution with $L(s) = 1 - e^{-s}$ where $s \geq 0$. There we also show how, when the market gets large, the distribution of locations can converge to an arbitrary distribution $L(s)$ which has a mean equal to 1.

We assume that search is random, that is, buyers' arrival on the circle is random, which implies that good locations for sllers are ones that have a long arc-length to their counter-clockwise neighbour. If we allowed buyers to choose locations (search is directed), then the arc length between locations in our model would become irrelevant.

We can relax the assumption that buyers arrive uniformly on the circle and assume that their arrival obeys some other probability distribution. However, this does not make the model more general, since it is $L(s)$, instead of the individual distributions of sellers and buyers, that matters for the market equilibrium.[6]

Market for Locations. After the realization of their locations and before meeting buyers, sellers can trade their locations in a competitive, frictionless market. The price of location s is denoted by $r(s)$.

[5] In the example of Figure 1, the probability that a buyer meets a seller in location $s = 4/3$ is thus $4/3 \cdot 1/2 = 2/3$ and the corresponding probability for $s = 2/3$ is $2/3 \cdot 1/2 = 1/3$.

[6] Note that here we treat the distribution of locations, $L(s)$, as exogenous. In Section 4, we consider the case where the distribution of locations is endogenous.

Trade and Payoffs. After sellers select their preferred location and sell their old location, buyers and sellers meet. Sellers select the buyer they trade with by means of a second-price auction (with a reserve price equal to the seller's valuation of zero).[7] Suppose that a seller offers quality z and has queue length λs (because its location is s). The expected surplus for this seller is given by

$$S(s,z) = z\sum_{n=1}^{\infty} e^{-\lambda s}\frac{(\lambda s)^n}{n!}\left[1-(1-x(z))^n\right] = z\left(1-e^{-\lambda s x(z)}\right), \tag{2}$$

where the summand on the left-hand side denotes the scenario where the seller meets n buyers and at least one likes the product. Alternatively, since the *effective queue length* (the expected number of buyers who value the product) is $\lambda s x(z)$, the probability that the seller meets at least one buyer who likes the product is given by the term in parenthesis on the right-hand side.

The seller's expected payoff from the auction is then

$$\pi(s,z) = z\sum_{n=2}^{\infty} e^{-\lambda s}\frac{(\lambda s)^n}{n!}\left[1-(1-x)^n - nx(1-x)^{n-1}\right] = z\left(1-e^{-\lambda s x}-\lambda s x e^{-\lambda s x}\right), \tag{3}$$

where we have suppressed the argument z from the function $x(z)$ to save space. The summand on the left-hand side denotes the probability that the seller meets n buyers and at least two of the n buyers value the good, in which case the transaction price equals z (if only one buyer arrives who likes the good, this buyer will just bid the reserve price). As in (2), since the effective queue length is is $\lambda s x$, the probability that the seller meets two or more buyers who value the product is given by the term in the parenthesis on the right-hand side.

To simplify exposition later, we define

$$\mathcal{P}(\lambda) = 1 - e^{-\lambda} - \lambda e^{-\lambda} \tag{4}$$

which is the probability of meeting two or more buyers when the expected queue length is λ. Sellers' payoff in equation (3) can then be rewritten as $\pi(s,z) = z\mathcal{P}(\lambda x s)$.

The expected payoff of a buyer who meets this seller is

$$zx\sum_{n=0}^{\infty} e^{-\lambda s}\frac{(\lambda s)^n}{n!}(1-x)^n = zxe^{-\lambda s x}, \tag{5}$$

where, as before, we have suppressed the argument z from the function $x(z)$. The summation

[7] This trading mechanism is payoff equivalent to the case where buyers submit a price to bid for the goods (first-price auction); see Burdett and Judd (1983) for this case.

on the left-hand side denotes the probability that the seller meets no other buyers who value the product, which, as argued above, is simply $e^{-\lambda s x}$. If the seller meets multiple buyers who like the product, the winning bid is z and the entire surplus goes to the sellers.

Sellers' Location Choice. Before meeting the buyers, the problem of seller z is to maximize the expected payoff from choosing location type s, which is given by

$$\max_s \pi(s, z) - r(s), \tag{6}$$

where $\pi(s, z)$ is given by equation (3). Equation (6) abstracts from the earnings associated with sellers selling their location endowment as that is irrelevant for the location choice problem.

Suppose that the support of $L(s)$ is an interval and that, in equilibrium, location s is chosen by sellers of type $z^*(s)$. The first-order condition then implies that the gradient of the location prices equals

$$r'(s) = \lambda z x \mathcal{P}'(\lambda x s) = \lambda^2 s z x^2 e^{-\lambda s x}, \tag{7}$$

where $z = z^*(s)$ and $x = x(z^*(s))$.

Buyer Entry. Free entry of buyers requires that their expected payoff after entering the market is exactly K. That is

$$K = \int_s z^*(s) x^*(s) e^{-\lambda s x^*(s)} s \, dL(s) \tag{8}$$

where $z^*(s)x^*(s)e^{-\lambda s x^*(s)}$ is the expected payoff from meeting a seller at location s, as given by equation (5), and $s \, dL(s)$ is the corresponding probability density.

Sellers' Ex-ante Expected Payoff. The expected payoff of sellers before entering the market is then given by

$$\Pi(z) = \pi(s^*(z), z) - r(s^*(z)) + R, \tag{9}$$

where $s^*(z)$ is the sellers' optimal location choice, and $R = \int_s r(s) \, dL(s)$ is the average location price, which is also the seller's expected payoff from selling their location since they are randomly endowed with a location according to the distribution $L(s)$. Note that the above equation depends on the equilibrium value of λ.

Equilibrium Definition. We can now define an equilibrium as follows.

Definition 1. *An equilibrium is an assignment of sellers to locations, a price schedule for locations $r(s)$ and a measure of buyers λ such that*

1. *Seller optimality: Given $r(s)$, sellers' choices of locations maximize their expected profit. That is, each seller solves the problem given by (6).*
2. *The market for locations clears: The price schedule for locations is such that for each s, the demand for type-s locations equals their supply.*
3. *Free entry of buyers: The expected payoff of buyers equals their entry cost K, i.e., equation (8) holds.*

3 Exogenous Distribution of Locations

3.1 Homogeneous Sellers

Before considering heterogeneous sellers in Section 3.2, we start in this section by assuming that all sellers have the same quality z. We first analyze the planner's problem and then the decentralized equilibrium. Finally, we show that the matching function is quite general and is equivalent to the class of *invariant* meeting technologies, as defined in Lester et al. (2015) and Cai et al. (2017).

Meeting Probabilities. Before locations are realized, the probability that a seller meets n buyers equals

$$P_n(\lambda) \equiv \int_s e^{-\lambda s} \frac{(\lambda s)^n}{n!} \, dL(s), \tag{10}$$

where $n \geq 0$. The probability for the seller to meet at least one buyer is then given by

$$m(\lambda) \equiv 1 - P_0(\lambda) = \int_s \left(1 - e^{-\lambda s}\right) dL(s). \tag{11}$$

Each buyer values the seller's product only with probability $x = x(z)$. Hence, a seller only trades if this seller meets at least one such buyer, which happens with probability

$$1 - \sum_{n=0}^{\infty} P_n(\lambda)(1-x)^n = m(\lambda x), \tag{12}$$

where the summand on the left-hand side represents the scenario where the seller meets n buyers who all do not value the product.[8] The resulting meeting probability is $m(\lambda x)$,

[8] A detailed derivation of equation (12) can be found in Appendix A.2.

which depends only on the effective queue length λx that the seller faces before the location realization. The above equation shows that the meeting process between the seller and buyers who like the seller's product is not affected by how many buyers there are who do not like the product. That is, the latter group does not impose search externalities on the former group. Consider for example a vegan restaurant in a particular location and suppose that there are on average 2 buyers per seat who have a desire for vegan food for whom this restaurant is nearest, then the meeting rate does not depend on how many people there are near this location who do not like vegan food.

Examples. Below we give some examples of $L(s)$ where we can explicitly calculate $m(\lambda)$. Recall that the mean of s is always 1. The examples illustrate that many meeting distributions like the geometric and the gamma distribution can be thought of as spatial mixtures of Poisson distributions.

Example 1. If the distribution of s is degenerate at 1, then $m(\lambda) = 1 - e^{-\lambda}$ which is the standard urn-ball matching function.

Example 2. If $L(s)$ is a discrete distribution with support $\{s_1, \ldots, s_I\}$ and $\mathbb{P}(s = s_i) = \ell_i$, then

$$m(\lambda) = \sum_{i=1}^{I} \ell_i \left(1 - e^{-\lambda s_i}\right).$$

Example 3. If $L(s)$ is an exponential distribution with $L(s) = 1 - e^{-s}$ where $s \geq 0$, then

$$m(\lambda) = \int_0^\infty \left(1 - e^{-\lambda s}\right) dL(s) = 1 - \frac{1}{1+\lambda}, \tag{13}$$

which is the geometric matching function (see e.g. Lester et al., 2015; Cai et al., 2017). More generally, let $L(s)$ be an exponential distribution with support $[s_0, \infty)$ where $s_0 \in [0, 1)$. That is, $L(s) = 1 - e^{-\rho(s-s_0)}$, where $\rho = (1 - s_0)^{-1}$ such that the mean of s is 1. Then

$$m(\lambda) = 1 - \frac{e^{-\lambda s_0}}{1 + \lambda(1 - s_0)}. \tag{14}$$

This matching function is strictly increasing in s_0 (for fixed λ). When $s_0 \to 1$, s gets more and more concentrated around 1 and $m(\lambda) \to 1 - e^{-\lambda}$.

Example 4. If $L(s)$ is a Gamma distribution with density $L'(s) = \rho^\rho s^{\rho-1} e^{-\rho s}/\Gamma(\rho)$ where $s \geq 0$ and $\Gamma(\cdot)$ is the standard Gamma function, then

$$m(\lambda) = 1 - \left(\frac{\rho}{\rho + \lambda}\right)^\rho.$$

When $\rho = 1$, we have the geometric meeting technology, and when $\rho = \infty$, we have the urn-ball meeting technology. For $\rho \in (1, \infty)$, the corresponding $P_n(\lambda)$, defined in equation (10), follows a negative binomial distribution. This distribution is of interest since Davis and de la Parra (2017) provide empirical evidence that the number of job applications that vacancies receive can be well approximated by a negative binomial distribution, adjusted so that zero has a larger weight.

The Planner's Problem. The planner's problem is to select the measure of buyers λ that maximizes total net surplus $Y(\lambda) - \lambda K$, where $Y(\lambda)$ equals

$$Y(\lambda) = \int_s z\left(1 - e^{-\lambda s x(z)}\right) dL(s) = zm(\lambda x(z)), \tag{15}$$

and where we used (12) for the second equality. Since $1 - e^{-\lambda s x(z)}$ and hence $Y(\lambda)$ are strictly concave in λ, the following first-order condition is both necessary and sufficient:

$$K = Y'(\lambda) = \int_s z s x(z) e^{-\lambda s x(z)} \, dL(s). \tag{16}$$

The Decentralized Equilibrium. In the decentralized equilibrium, when a buyer meets a seller at location s, the buyer receives a strictly positive payoff if and only if the buyer likes the product, which happens with probability $x(z)$, and the seller meets no other buyers who like the seller's product. Recall that the buyer's expected payoff in this case is given by equation (5). Given that meetings between buyers and sellers are random, free entry then requires that

$$K = \int_s z x(z) e^{-\lambda s x(z)} s \, dL(s),$$

where the integration represents the uncertainty with respect to locations. The above equation coincides exactly with the planner's first-order condition (16). Thus we have the following result.

Proposition 1. *Suppose that all sellers have the same quality z. Buyer entry in the decentralized equilibrium is then constrained efficient.*

The above result is quite intuitive and is well known in the literature. Suppose that a seller meets multiple buyers. If a buyer is the only one who likes the product and is willing to make a bid, then without this buyer the surplus would be zero. If there are multiple buyers who like the product, then the marginal contribution to surplus by any individual buyer is zero. Furthermore, a buyer does not influence the meeting process between sellers and other buyers (no search externalities). Therefore, a second-price auction guarantees that buyers always receive their marginal contribution to surplus.

Comparative Static. A simple consequence of the above framework is that making the goods simultaneously better and more niche such that the expected buyer value remains the same, increases total surplus for each λ and also increases buyer entry. The intuition is that since there are no meeting externalities among buyers, making the goods more niche takes advantage of the fact that sellers can meet multiple buyers simultaneously and they only need one buyer who likes their product (recall that with niche products, those who like it, like it a lot).

More precisely, the effect on total surplus of a percentage increase of z combined with a corresponding decrease in x such that the expected value of $zx(z)$ remains constant is equivalent to a percentage increase in the measure of sellers. To see this formally, let the measure of sellers be N_s and the measure of buyers be N_b with $\lambda = N_b/N_s$. Expected total surplus is then given by $N_s z m\left(\frac{N_b z x(z)}{N_s z}\right)$. Thus keeping N_b (the measure of buyers) and the expected value of $zx(z)$ fixed, only the product $N_s z$ matters for the expected surplus, which implies that the effect of a percentage increase of z is equivalent to a percentage increase in the measure of sellers. Note that the same observation applies to buyers' marginal contribution to surplus, which is given by $zx(z)m'\left(\frac{N_b z x(z)}{N_s z}\right)$.

Finally, we can calculate the price of locations. Since all sellers are homogeneous, equation (7) becomes

$$r'(s) = \lambda^2 z x(z)^2 s e^{-\lambda x(z) s}.$$

The above equation implies that $r''(s) = z\lambda^2 x(z)^2 e^{-\lambda s x}(1 - \lambda x(z)s)$. Thus $r''(s) > 0$ when $s < \lambda x(z)$ and $r''(s) < 0$ when $s > \lambda x(z)$. That is, $r(s)$ has an S-shape. Starting from the worst location, increasing the expected number of buyers that a seller meets initially generates little value for the seller because although the probability that a seller receives no buyers declines, the probability to receive more than one buyer remains close to zero (with only one buyer, the seller's payoff is zero). Then, if we further increase s, better locations help to substantially increase the likelihood to get two or more buyers. Finally, when the likelihood of getting multiple buyers is already close to one, better locations add little value and consequently, the price for better locations increases only weakly.

Suppose that the lower support of the distribution $L(s)$ is zero. Solving the differential equation (7) explicitly yields[9]

$$r(s) = z\mathcal{P}(\lambda x(z)s). \tag{17}$$

Note that $\mathcal{P}(\lambda x(z)s)$ is the probability that a seller at location s meets two or more buyers which is independent of the distribution of locations. Since all sellers are homogeneous, in equilibrium there is no arbitrage possible in terms of location choice. This implies that the price difference between locations exactly offsets the corresponding difference in the payoffs from meeting more buyers. That is, given that all sellers are homogeneous, the location prices are such that sellers obtain no gains from trading locations.

Surplus-Maximizing Meeting Technology. For a general distribution of locations $L(s)$, $m(\lambda)$ in equation (11) can be thought of as a mixture of urn-ball processes. It then follows from Jensen's inequality that aggregate matching efficiency can be improved by making locations more equal.

Theorem 1. *Suppose that all sellers have the same quality z. For any given λ, total surplus $Y(\lambda)$ is highest with the urn-ball meeting technology where $m(\lambda) = 1 - e^{-\lambda}$, which is obtained when $L(s)$ is degenerate at 1 (all locations are the same).*

Proof. See Appendix A.3. □

The above result is quite intuitive for our benchmark. However, when combined with Theorem 2 below, it also shows that the urn-ball meeting technology has the highest matching efficiency among all invariant meeting technologies (to be specified below). This latter result is much less obvious. Note that Theorem 1 also carries over to the case where the market for locations does not exist and the matching between sellers and locations is purely random.

Relation with Invariance. In our model, the probability that a seller meets n buyers is given explicitly by equation (10). In earlier literature (see, e.g., Eeckhout and Kircher, 2010), it has been common to start with $P_n(\lambda)$ unspecified. An invariant meeting technology is then defined as one for which equation (12) holds for any λ and x, with (set $x = 1$) $m(\lambda) = 1 - P_0(\lambda)$ (Lester et al., 2015; Cai et al., 2017). Furthermore, the n^{th} derivative of equation (12) with respect to x, evaluated at $x = 1$, equals $P_n(\lambda) = (-1)^{n+1}\frac{\lambda^n}{n!}m^{(n)}(\lambda)$.[10]

Therefore, for any invariant meeting technology, its associated function $m(\lambda)$ has the following properties: i) It is non-negative, and ii) it is infinitely differentiable and $(-1)^{n+1}\frac{d^n}{d\lambda^n}m(\lambda) \geq$

[9]Note that $r(0)$ must be zero because a location with $s = 0$ generates no surplus. When the lower support of distribution $L(s)$ is given by $s_0 > 0$, then as is well known from the literature (see Chiappori (2017) for a textbook treatment), $r(s_0)$ is indeterminate, since the measure of locations equals exactly the measure of sellers. However, since a seller of a product is both a buyer and a seller of locations, this indeterminacy does not affect the payoff of sellers.

[10]Cai et al. (2023) show that $P_0(\lambda)$ is a probability generating function.

0 for $n \geq 1$. That is, $m(\lambda)$ is a *Bernstein function*. In addition, $m(\lambda)$ is bounded between 0 and 1, $m(0) = 0$, and $m'(0) \leq 1$. By Bernstein's theorem, $m(\lambda)$ has the following representation.[11]

Theorem 2. *A function $m(\lambda)$ generates an invariant meeting technology if and only if there exists a probability measure $\widetilde{L}(s)$ on $[0, \infty)$ (the positive real half-line) with $\int_{[0,\infty)} s\, d\widetilde{L}(s) \leq 1$ such that*

$$m(\lambda) = \int_{[0,\infty)} \left(1 - e^{-\lambda s}\right) d\widetilde{L}(s). \tag{18}$$

Proof. See Appendix A.4 □

If $\int_{[0,\infty)} s\, d\widetilde{L}(s) = 1$, then the class of invariant technologies corresponds exactly to the above search process on the circle with $\widetilde{L}(s) = L(s)$.[12] We can also use our model to understand the general case. If $\int_{[0,\infty)} s\, d\widetilde{L}(s) < 1$, then with probability $1 - \int_{[0,\infty)} s\, d\widetilde{L}(s)$ buyers do not arrive on the circle and are passive. This probability is independent of λ. Given the set of buyers who arrive on the circle, the matching process specified by an invariant technology is again equivalent to our search process on the circle. Thus it is without loss of generality to assume that $\int_{[0,\infty)} s\, d\widetilde{L}(s) = 1$. Given the correspondence in Theorem 2, $P_n(\lambda)$ can also be calculated by equation (10) (a mixture of the corresponding Poisson probabilities). This representation of meeting technologies has proved to be useful in other settings as well. For example, Becker and Mangin (2023) use our Bernstein representation to study extreme outcomes in markets with search frictions.

3.2 Heterogeneous Sellers

We now consider the general case where sellers are heterogeneous in quality z, i.e., $F(z)$ is not degenerate. Suppose for example that that there are many chefs but only few of them have the talent to cook at the three-star level. At what locations would they open restaurants? What locations would less-talented chefs choose? And how would the planner assign chefs to locations?

The Planner's Problem. The planner's problem is to first match sellers with locations and to then select the measure of buyers per seller in order to maximize total net surplus. Recall that $S(s, z)$, the expected surplus between a seller z and a location s, is given by equation (2).

[11] See page 21 of Schilling et al. (2012) for the definition of Bernstein functions and Bernstein's Theorem.
[12] We use the notation $\int_{[0,\infty)}$ to emphasize that there can be a mass point at 0.

Lemma 1. *$S(s,z)$, which is defined by equation (2), is strictly increasing in z. Furthermore, it is strictly supermodular in (s,z) for any $\lambda > 0$.*

Proof. See Appendix A.5. □

Since surplus $S(s,z)$ is supermodular in (s,z), the planner's solution is characterized by PAM: better-quality sellers are assigned to better locations. Suppose that both F and L are continuous distributions. The optimal assignment is then captured by the following correspondence,

$$1 - F(z) = 1 - L(s^*(z)), \tag{19}$$

where $s^*(z)$ is the optimal location s for seller type z.

Given the optimal matching between sellers and locations, expected total surplus is given by $Y(\lambda) - \lambda K$, where

$$Y(\lambda) = \int_s z^*(s)\left(1 - e^{-\lambda s x^*(s)}\right) dL(s), \tag{20}$$

$x^*(s) \equiv x(z^*(s))$ and $z^*(s)$ is the inverse of $s^*(z)$. The above equation is strictly concave in λ. As before, the optimal measure of buyers in the market is then given by the FOC:

$$K = Y'(\lambda) = \int_s z^*(s) s x^*(s) e^{-\lambda s x^*(s)}\, dL(s). \tag{21}$$

The Decentralized Equilibrium. We now analyze the decentralized equilibrium. Consider first the matching between sellers and locations. Recall that the expected profit for a seller z from selling the good after choosing location s is $\pi(s,z)$, as given by equation (3).

As is well known since Becker (1973), strict supermodularity of $\pi(s,z)$ in (s,z) implies that the market outcome must be PAM: better-quality sellers choose better locations. However, the following lemma shows that unlike the surplus function $S(s,z)$, sellers' expected profit $\pi(s,z)$ is not necessarily supermodular, unless the elasticity of $x(z)$ is greater than $-1/2$.

Lemma 2. *Consider $\pi(s,z)$ defined by equation (3). Given z, if $\varepsilon_x(z) \equiv zx'(z)/x(z) \geq -1/2$, then the derivatives $\pi_z(s,z) > 0$ and $\pi_{sz}(s,z) > 0$ for any $\lambda > 0$.*

Alternatively, suppose that $\varepsilon_x(z) \in [-1, -1/2)$. Then 1) $\pi_{sz}(s,z) > 0$ if and only if the effective queue length $\lambda x(z) s > 2 + 1/\varepsilon_x(z)$, and 2) there exist a threshold $\Lambda(\varepsilon)$ such that $\pi_z(s,z) > 0$ if and only if the effective queue length $\lambda x(z) s > \Lambda(\varepsilon_x(z))$, where $\Lambda(\varepsilon_x(z))$ is the solution of λ to the following equation

$$\frac{1 - e^{-\lambda} - \lambda e^{-\lambda}}{\lambda} = -\lambda e^{-\lambda} \varepsilon_x(z). \tag{22}$$

$\Lambda(\varepsilon_x(z))$ *is strictly decreasing for* $\varepsilon_x(z) \in [-1, -1/2)$ *with* $\Lambda(-1/2) = 0$.

Proof. See Appendix A.6. □

Supermodularity concerns how the marginal value of s varies with z. Thus when $zx(z)$ is increasing in z, then surplus $S(s, z)$ is supermodular. However, profit $\pi(s, z)$ is supermodular when $zx(z)^2$ (the probability that two buyers are valuable is $x(z)^2$, ignoring higher order terms) is increasing in z, which is equivalent to $zx'(z)/x(z) \geq -1/2$. When the latter condition holds, the quality of sellers who choose location s in equilibrium is also given by equation (19), as in the planner's solution. This condition on the elasticity says that if the chance that a buyer likes the product does not decrease too fast with the quality of the product, high-quality sellers benefit more from good locations than low-quality sellers. If the market for high-quality products becomes thin very fast (the elasticity $zx'(z)/x(z) \leq -1/2$) then even at good locations the probability to get two or more effective buyers is too small for the high-quality sellers to be willing to pay the higher rents there.

Given the equilibrium matching between sellers and locations, buyer entry is such that the expected payoff of buyers equals the entry cost K, i.e., equation (8) must hold. Comparing (8) with the planner's solution shows that buyer entry is constrained efficient.

Proposition 2. *Suppose that* $zx'(z)/x(z) \geq -1/2$ *for any* z. *The decentralized equilibrium is then constrained efficient.*

When $zx'(z)/x(z) < -1/2$, it is possible that the equilibrium matching between sellers and locations does not exhibit PAM and hence the decentralized equilibrium is not efficient. To see this, note that a necessary condition for PAM to hold is that along the equilibrium path, for each s the cross-partial derivative $\pi_{sz}(s, z^*(s)) \geq 0$, where $z^*(s)$ is given by equation (19) and is independent of λ. If the lower support of $L(s)$ is zero, then for small s the effective queue length $\lambda x(z)s$ is close to zero, and by Lemma 2, $\pi_{sz}(s, z^*(s))$ is strictly negative. Hence, the necessary condition does not hold.

The reason that the planner's solution can differ from the market outcome is that in the planner's solution, a higher s increases surplus $S(s, z)$ because it increases the probability that a seller meets at least one effective buyer, whereas in the competitive market of locations, it increases $\pi(s, z)$ (the gross payoff of buyers of locations) because it increases the probability that a seller meets at least two effective buyers who then bid the price up. The investment of a seller in a good location is therefore partly a rent-seeking activity because the returns of receiving more than one buyer are zero from a social welfare point of view but positive from an individual firm's point of view.

Specifically, recall that given z, $\pi(s, z)$ as a function of s has an S-shape: $\mathcal{P}(\lambda x(z)s)$ is first convex and then concave in s. Consider two sellers $z^a < z^b$ who match with locations

$s^a < s^b$, respectively. If the effective queue lengths for the two sellers are low, then assigning location s^b to seller z^a and location s^a to seller z^b will increase the total profits of the two sellers by exploiting the initial convexity of $\mathcal{P}(\cdot)$ but this will decrease social surplus. The *increase* in seller a's probability that two or more buyers visit at the good location exceeds the *decrease* in seller b's probability that two or more buyers visit at the bad location: $\mathcal{P}(\lambda x(z^a)s^b) - \mathcal{P}(\lambda x(z^a)s^a) > \mathcal{P}(\lambda x(z^b)s^b) - \mathcal{P}(\lambda x(z^b)s^a)$, or equivalently $\int_{s^a}^{s^b} \mathcal{P}'(\lambda x(z^a)s)ds > \int_{s^a}^{s^b} \mathcal{P}'(\lambda x(z^b)s)ds$, where we used the fact that $x(z^a) > x(z^b)$ and $\mathcal{P}(\widetilde{\lambda})$ is convex when $\widetilde{\lambda}$ is small. Furthermore, the increase in seller a's probability that two or more buyers visit can be sufficiently large such that it outweights z^b/z^a times the corresponding decrease of seller b's probability. To sum up, in the convex range of $\mathcal{P}(\cdot)$ where the buyer-seller ratio is low, sellers of general products benefit more from good locations then sellers of high quality, niche products because the latter group rarely meets two or more buyers who like their product.

3.3 Three Examples

Below, we analytically characterize the the planner's solution for some parametric examples of seller-type and location-quality distributions.

3.3.1 Vertical Quality ($x(z) = 1$ for any z)

In our first example, buyers like each product with certainty, $x(z) = 1$ for all z. As a result, all buyers agree on the ranking of products, such that quality differences between sellers can be viewed as being vertical. We assume that the quality distribution follows a power law and that the location distribution is exponential. Specifically, let $F(z) = 1 - \left(\frac{z_0}{z}\right)^\alpha$ with $\alpha > 1$ and $z \geq z_0 > 0$ and let $L(s) = 1 - e^{-\rho(s-s_0)}$ where $s \geq s_0$, $s_0 \in [0, 1)$ and $\rho = (1 - s_0)^{-1}$.

Expected Total Surplus. Since $x(z) = 1$ for all z, the equilibrium in the market for locations exhibits PAM and thus coincides with the planner's solution (see Lemma 1 and 2). By equation (19), the correspondence between sellers and locations is

$$e^{-(s-s_0)/(1-s_0)} = \left(\frac{z_0}{z}\right)^\alpha,$$

which implies that

$$z^*(s) = z_0 e^{(s-s_0)/(\alpha(1-s_0))}. \tag{23}$$

Substituting the above equation into equation (20) implies that total surplus equals

$$Y(\lambda, \alpha, s_0) = \int_s z^*(s)\left(1 - e^{-\lambda s x^*(s)}\right) dL(s) = \overline{z}\left(1 - \frac{(\alpha - 1)e^{-\lambda s_0}}{(\alpha - 1) + \alpha\lambda(1 - s_0)}\right), \tag{24}$$

where $\overline{z} = z_0\alpha/(\alpha - 1)$ denotes the mean of z and we have added (α, s_0) as arguments of $Y(\lambda)$ to show that total surplus depends on quality and location dispersion. The following lemma characterizes this dependence.

Lemma 3. *Let λ be fixed. If we increase both α and z_0 so that the average quality $\overline{z} = z_0\alpha/(\alpha - 1)$ is constant, then $Y(\lambda, \alpha, s_0)$ decreases.*

If $\alpha\lambda \leq 1$, $Y(\lambda, \alpha, s_0)$ is strictly decreasing in $s_0 \in [0, 1]$. If $\alpha\lambda > 1$, then $Y(\lambda, \alpha, s_0)$ is first increasing and then decreasing in s_0, reaching a maximum at $s_0 = 1 - 1/(\alpha\lambda)$.

Proof. See Appendix A.7. □

Increasing α, while adjusting z_0 to keep the mean $\overline{z}$ fixed, reduces quality dispersion. In the limit $\alpha \to \infty$, all sellers have the same quality z. In this case, the term in parentheses at the right-hand side of (24) reduces to $m(\lambda)$ in equation (14) and surplus is smallest. Hence, the fact that locations are heterogeneous makes quality dispersion desirable because of the complementarity between quality and location.

The effect of location dispersion on total surplus is more complicated. All locations are the same when $s_0 = 1$, while location dispersion is maximized when $s_0 = 0$. By equation (14), a smaller s_0 (more location dispersion) always reduces the total number of trades. When λ is small ($\leq 1/\alpha$), the increase in the matching probability of high-quality sellers dominates the overall decrease in matches so location dispersion is desirable. When λ is large, too much location dispersion leads to a minor increase in the matching probability of high-quality sellers while it reduces the matching probability of low-quality sellers substantially and loation dispersion is not desirable.

Equilibrium. In the above analysis, we have treated λ as exogenous. To determine the equilibrium value of λ and the equilibrium seller payoffs, we consider two extreme cases $s_0 = 0$ and 1. We will use subscripts 0 and 1 to compare variables between the two cases.

When $s_0 = 1$, all locations are the same and the expected output is $Y(\lambda, \alpha, 1) = \overline{z}(1-e^{-\lambda})$. The equilibrium (or equivalently, the socially optimal) λ is given by $K = \partial Y(\lambda, \alpha, 1)/\partial\lambda$, which implies that $\lambda_1 = \log(\overline{z}/K)$, where the subscript 1 indicates the case $s_0 = 1$. Since all locations are the same, sellers will not trade locations and their expected profit is $\Pi_1(z) = z(1 - e^{-\lambda_1} - \lambda_1 e^{-\lambda_1})$. Note that in this special case, varying α (the quality distribution) while holding $\overline{z}$ constant (average value) has no effect on the equilibrium measure of buyers and expected total surplus.

Next, consider the case $s_0 = 0$. The equilibrium (or equivalently, the socially optimal) λ is given by $K = \partial Y(\lambda, \alpha, 0)/\partial\lambda$, where $Y(\lambda, \alpha, 0)$ is given by equation (24) with $s_0 = 0$.

Hence,

$$\lambda_0 = \sqrt{\frac{\overline{z}}{K}\frac{\alpha-1}{\alpha}} - \frac{\alpha-1}{\alpha}, \tag{25}$$

which implies that positive entry of buyers happens if and only if $K < \overline{z}\alpha/(\alpha-1)$.[13] Combined with equation (24), this implies that net social surplus equals

$$Y(\lambda_0(\alpha), \alpha, 0) - K\lambda_0(\alpha) = \left(\sqrt{\overline{z}} - \sqrt{K\frac{\alpha-1}{\alpha}}\right)^2,$$

which is strictly decreasing in α.

Finally, in the proposition below, we derive sellers' expected profits for the two special cases $s_0 = \{0, 1\}$. To simplify exposition, we focus on the case in which $K = z_0$ since this yields a particular simple expression for $s_0 = 0$.

Proposition 3. *Consider two location distributions: $L_0(s) = 1 - e^{-s}$ and $L_1(s)$ is degenerate at $s = 1$. Suppose that $x(z) = 1$ for all z, and $F(z) = 1 - (z_0/z)^\alpha$ with $\alpha > 1$. Furthermore, assume $K = z_0 = 1$.*

Let $\Pi_i(z)$ be the ex-ante seller payoff under location distribution $L_i(s)$ with $i = 0, 1$, which is defined by equation (9). Then for $z \geq z_0 = 1$, we have

$$\Pi_0(z) = \frac{1}{\alpha^2} + z - 1 - \frac{1}{2}\log(z)\left(2 + \log(z)\right), \tag{26}$$

and

$$\Pi_1(z) = z\left(\frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right)\log\left(1 - \frac{1}{\alpha}\right)\right), \tag{27}$$

where $\Pi_0(z)$ is strictly convex with $\lim_{z\to\infty}\Pi_0'(z) = 1$, and $\Pi_1(z)$ is linear with a slope strictly smaller than 1. When z is close to 1 or sufficiently large, $\Pi_0(z) > \Pi_1(z)$. Furthermore, there exists a unique α^* (approximately 1.9) such that $\Pi_0(z) > \Pi_1(z)$ for all z if and only if $\alpha > \alpha^*$.

Proof. See Appendix A.8. □

Two things stand out when comparing sellers' expected payoffs under the two location distributions (heterogeneous locations with $s_0 = 0$ and identical locations with $s_0 = s = 1$). First, when $s_0 = 0$, the sellers with the worst type z_0 have $s = 0$ in equilibrium (they meet no buyers so this is equivalent to being inactive). When $s_0 = s = 1$, each seller expects an

[13] Note that the equilibrium measure of buyers is not necessarily monotonic in α. When $\overline{z}/K \geq 4$, then $\lambda_0(\alpha)$ is always strictly increasing in α. When $\overline{z}/K < 4$, $\lambda_0(\alpha)$ is strictly increasing in α for $\alpha < 4/(4 - \overline{z}/K)$, and strictly decreasing in α for $\alpha > 4/(4 - \overline{z}/K)$.

effective queue length λ. One may expect that sellers with low z are worse off under $s_0 = 0$ than under $s_0 = 1$. However, the above proposition shows that it can be the case that when $s_0 = 0$, the expected payoff from selling one's location endowment is so high that all sellers are better off. When the location distribution is dispersed, sellers with high z are willing to pay a high price for good locations. This not only benefits high-quality sellers but can also benefit low-quality sellers, since those sellers are also location owners.

Second, so far we have assumed that sellers incur no costs of participating in this market. Suppose now that sellers need to pay a production cost $c(z)$ to enter the market, which is assumed to be weakly increasing and convex. Since $\Pi_1(z)$ is linear, it is not profitable for high-quality sellers to enter the market when all locations are the same ($s_0 = 1$) and $c'(z) > \Pi_1'(z) = \frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right) \log\left(1 - \frac{1}{\alpha}\right)$. At the same time, it can still be the case that $c(z) < \Pi_0(z)$ for all z so that all sellers, including the high-quality ones, find it profitable to enter the market when locations are hetergeneous ($s_0 = 0$). In the latter case, the good locations allow the high-quality sellers to meet many buyers (and thus trade and create surplus with high probability), whereas offering a high-quality product is too risky when all locations are identical, because there is a substantial chance that too few buyers arrive.

3.3.2 Power Law for Both Distributions ($F(z)$ and $L(s)$)

In our second example, we allow for heterogeneity in $x(z)$. Neiman and Vavra (2023) document that products that are more niche are offered in large dense cities. That is, niche sellers with a smaller x and larger z will choose a larger s. However, it is not clear whether the effective queue length $\lambda x s$ (or equivalently sellers' trading probability) is increasing or decreasing in seller types. To show that both options are feasible, we construct a knife-edge example where the effective queue length (or sellers' trading probability) remains constant. By perturbating this example, a seller's trading probability can either increase or decreases with z (see our third example below for a first-order approximation approach).

We assume that both the quality distribution and location distribution follow power laws. Specifically, let $F(z) = 1 - \left(\frac{z_0}{z}\right)^\alpha$ with $\alpha > 1$ and $z \geq z_0$ and let $L(s) = 1 - \left(\frac{s_0}{s}\right)^\beta$ where $s \geq s_0$ and $\beta > 1$. Since the mean of s must be 1, we have $s_0 = (\beta - 1)/\beta \in (0, 1)$. Furthermore, let $x(z) = \left(\frac{z_0}{z}\right)^\gamma$ where $0 < \gamma \leq 1$. A smaller γ implies that the probability that a buyer likes the good declines less quickly in z, bringing us closer to the case of vertical quality. Note that $\mathbb{E}z = \overline{z} = z_0\alpha/(\alpha - 1)$ and $\mathbb{E}zx(z) = z_0\alpha/(\alpha - 1 + \gamma)$. We then consider the knife-edge case where $\alpha = \beta\gamma$ or equivalently $\alpha(1 - s_0) = \gamma$, since it yields analytical tractability.

Given that the assignment between sellers and locations is PAM at the planner's solution, the correspondence between sellers and locations is

$$\left(\frac{z_0}{z}\right)^{\alpha} = \left(\frac{s_0}{s}\right)^{\beta}$$

which implies that

$$z^*(s) = z_0 \left(\frac{s}{s_0}\right)^{\beta/\alpha}. \tag{28}$$

The matching probability for a seller of type z is given by $1 - e^{-\lambda x(z)s^*(z)} = 1 - e^{-\lambda s_0}$, which is constant across different sellers. Total expected surplus is then given by

$$Y(\lambda, \alpha, s_0) = \left(1 - e^{-\lambda s_0}\right)\overline{z}. \tag{29}$$

Since all sellers have the same trading probability, the expected total surplus is simply the matching probability times the average quality.

3.3.3 Making Goods More Niche: A First-Order Approximation Approach

When all sellers are homogeneous, we show in Section 3.1 that making the goods more niche (increase z and decrease x simultaneously so that $zx(z)$ is constant) increases total surplus. When sellers are heterogenous, making the goods more niche while fixing the distribution of $q = zx(z)$ again increases total surplus. If we increase $z(q)$ by the same percentage for all q and decrease the corresponding x by the same percentage, then, as we argued in the case of homogeneous sellers, this has the same effect on total surplus as a corresponding percentage change in the measure of sellers.

Recall that $x(z) = \left(\frac{z_0}{z}\right)^{\gamma}$. A more interesting exercise is to see how an increase of γ, the elasticity of $x(z)$, while holding the distribution of q constant increases total surplus. In this case, the percentage increase in z (due to an increase in γ) is higher for a higher fixed expected value $q = zx(z)$. To see this, (with a slight abuse of notation) let $z(q, \gamma)$ and $x(q, \gamma)$ be the value and trading probability associated with q and γ. Since $zx(z) = q$ and $x(z) = \left(\frac{z_0}{z}\right)^{\gamma}$, we have $z(q, \gamma) = z_0(q/z_0)^{1/(1-\gamma)}$ and $x(q, \gamma) = (q/z_0)^{-\gamma/(1-\gamma)}$. The effect of a higher γ on the percentage change of $z(q, \gamma)$ is

$$\frac{\partial \log z(q, \gamma)}{\partial \gamma} = \frac{\log(q/z_0)}{(1-\gamma)^2},$$

which is strictly increasing in q. Thus, the percentage increase in $z(q, \gamma)$ (due to an increase in γ) is higher for a higher fixed expected value q.

In the above two examples, γ equals 0 or $\alpha(1 - s_0)$, respectively; an analytic expression for total surplus for general γ is difficult to obtain. We thus adopt a first-order approximation approach for the examples above to study the effects of increasing γ on surplus.[14] Furthermore, we analyze how this effect depends on α, which measures the dispersion of seller quality. Recall that in the two examples above, we assumed that the distribution of z follows a power law, i.e. $F(z) = 1 - \left(\frac{z_0}{z}\right)^{\alpha}$ with $\alpha > 1$ and $z \geq z_0$. In Appendix A.9 we show that in both examples above, the effect is smaller when α is larger, that is when quality dispersion is smaller.

4 Endogenous Distribution of Locations

We now consider the case where the distribution of locations is endogenous. To simplify exposition, we focus on the matching between sellers and locations by assuming that the measure of buyers is exogenously given. After all, conditional on the matching between sellers and locations, buyer entry is always constrained efficient because second-price auctions are used as trading mechanism.

4.1 The Planner's Problem

The planner's objective is to choose a distribution of locations $L(s)$ to maximize total net surplus, given the measure of buyers. Equivalently, we can assume that the planner allocates a location space $s(z)$ for each seller type, subject to the constraint that the total space is 1, i.e., the circumference of the circle.[15]

The planner's problem is thus given by

$$\max_{s(z)} \quad Y = \int_{z_0}^{\infty} z\left(1 - e^{-\lambda s(z) x(z)}\right) dF(z), \tag{30}$$

$$\text{s.t.} \int_{z_0}^{\infty} s(z)\, dF(z) = 1. \tag{31}$$

As in equation (2), the integrand in the first line is the surplus generated by a seller of type z. Total expected surplus Y then follows by integrating across different seller types. The constraint in the second line specifies that the total space (on the circle) allocated to the sellers should add up to 1.

[14]We let γ be slightly above 0 and $\alpha(1 - s_0)$, respectively, for these examples, while holding other factors fixed.

[15]As with homogeneous sellers, the planner will choose the same s for the same type of sellers (urn-ball is optimal when sellers are homogeneous).

The Lagrangian of this problem is

$$\mathcal{L} = \int_{z_0}^{\infty} z\left(1 - e^{-\lambda s(z)x(z)}\right)\, dF(z) + \xi\lambda\left(1 - \int_{z_0}^{\infty} s(z)\, dF(z)\right),$$

where ξ is the modified multiplier (multiplied by λ). Let the planner's solution be $s_p(z)$. For each z, it must satisfy the first-order condition

$$zx(z)e^{-\lambda x(z)s_p(z)} = \xi \quad \text{if} \quad \xi < zx(z), \tag{32}$$

with complementary slackness (if $zx(z) \leq \xi$, then $s_p(z) = 0$). Since the objective function is strictly concave in $s(z)$ and the constraint is linear in $s(z)$, the FOC is both necessary and sufficient for the planner's solution. Solving the FOC yields that for $zx(z) \geq \xi$,

$$s_p(z) = \frac{1}{\lambda x(z)}\left(\log zx(z) - \log \xi\right). \tag{33}$$

Since $zx(z)$ is increasing and $x(z)$ is decreasing, it follows that $s_p(z)$ is increasing. In other words, PAM between sellers and locations continues to hold.

Interestingly, the planner's solution is equivalent to the outcome in a directed search equilibrium (buyers observe the sellers' locations and the terms of trade ex ante), where the Lagrangian multiplier ξ corresponds to the market utility of buyers.

Special Cases. We consider two special cases where the planner's solution can be solved explicitly. First, we keep expected quality $zx(z)$ is constant. By equation (32), when $s(z) = 0$, the marginal benefit of increasing $s(z)$ is $zx(z)$. Therefore, all sellers should be active ($s(z) > 0$). Furthermore, the first-order condition (32) implies that $x(z)s_p(z)$ should be constant across different seller types, which implies that $s_p(z)$ must be proportional to z. Since the mean of s is 1, $s_p(z)$ is simply z divided by the mean of z. Even though all sellers offer the same expected value, niche goods should be offered in a location with a higher s. For future use, the following lemma summarizes the above results.

Lemma 4. *Suppose that $zx(z)$ is constant, i.e., $zx(z) = z_0x(z_0)$ for all $z \geq z_0$. Then $s_p(z) = z/\overline{z}$, where $\overline{z} = \int_{z_0}^{\infty} z dF(z)$ and the optimal location distribution is given by $L(s) = F(s\overline{z})$ with $s \geq z_0/\overline{z}$.*

The second special case is $x(z) = 1$ for all z (pure vertical quality). In this case, the first-order condition (33) becomes

$$\lambda s_p(z) = \log z - \log \xi. \tag{34}$$

The optimal s_i is linear in the log of z_i. Below, we solve the planner's problem explicitly by assuming that $F(z)$ is a power distribution with $F(z) = 1 - (z_0/z)^\alpha$, where $\alpha > 1$. In this case, the planner may exclude low-type sellers from participation (i.e., set $s(z) = 0$). Let z_1 be the minimum type among all active sellers. If $z_1 > z_0$, then $s(z) = 0$ for $z \leq z_1$ and $s(z) > 0$ for $z > z_1$. The type distribution of active sellers is then $F_1(z) = 1 - (z_1/z)^\alpha$ where $z \geq z_1$. Let $L_1(s)$ be the location distribution among active sellers. The optimal distribution $L_1(s)$ is then given by $1 - L_1(s) = \mathbb{P}(\log z - \log \xi \geq \lambda s) = \mathbb{P}(z \geq \xi e^{\lambda s}) = (\frac{z_1}{\xi})^\alpha e^{-\alpha\lambda s} = e^{-\alpha\lambda(s-s_0)}$, where $s_0 = s_p(z_1) = (\log z_1 - \log \xi)/\lambda$.

We now separately consider the two cases $z_1 = z_0$ and $z_1 > z_0$. First, when $z_1 = z_0$, all sellers are active. Since the mean of s must be one, we have $\alpha\lambda = (1 - s_0)^{-1}$, i.e., set $\rho = \alpha\lambda$ in the example in Section 3.3.1. Therefore, $z_1 = z_0$ *if and only if* $\alpha\lambda \geq 1$. Furthermore, the buyers' marginal contribution to surplus ξ is then determined by the condition: $1/(\alpha\lambda) = 1 - s_0 = 1 - (\log z_0 - \log \xi)/\lambda$, which implies that $\xi = z_0 e^{-(\lambda - 1/\alpha)}$.

Second, when $\alpha\lambda < 1$ or equivalently $z_1 > z_0$, we have $z_1 = \xi$, and $1 - L_1(s) = e^{-\alpha\lambda s}$. Since now only a fraction $(z_0/z_1)^\alpha$ of sellers are active, the unconditional distribution of locations is $1 - L(s) = 1 - (z_0/z_1)^\alpha(1 - L_1(s))$. Since the mean of $L(s)$ must be 1, we have $(z_0/z_1)^\alpha = \alpha\lambda$ and $z_1 = \xi = z_0(\alpha\lambda)^{-1/\alpha}$.

4.2 The Decentralized Equilibrium

To make the equilibrium distribution of locations endogenous, we assume that sellers form a coalition, as in a real estate investment trust (REIT). The coalition first chooses the seller-optimal distribution of locations, i.e., the one that maximizes the expected total seller profit. Then it charges competitive market rents which are redistributed back to the sellers in a lump-sum way. The constraint remains that the total supply of arc length is fixed and equals 1. After the purchase of locations, sellers randomly meet with buyers and each product is sold by a second-price auction.[16]

The competitive nature of the market for locations has two implications. First, it is without loss of generality (by the first welfare theorem) to assume that—instead of the two-stage process described above—the coalition chooses an arc length for each seller directly to maximize the expected total seller profit, which will be the approach that we follow below. Second, at the coalition's optimal solution, the price of locations $r(s)$ must be linear. To see this, note that coalition optimality requires that the marginal values of arc length are constant across different sellers. Since the marginal value of arc length should be equal to the gradient of location price (see equation (7)), the competitive price of locations must be

[16] As in the benchmark model, we allow sellers to be initially endowed with some location. As before, this will not affect the equilibrium outcome.

linear in s (see the discussion around equation (38) below). Therefore, an equivalent way of endogenizing the equilibrium location distribution is to allow sellers to buy and sell arc lengths at a unit price in a competitive market, so the total price is linear. The equilibrium unit price is such that the total demand of arc lengths equals the total supply.

Homogeneous Sellers. To understand the coalition's problem, we first consider the special case where all sellers are homogeneous. As we illustrate below, since the maximization problem is not concave, we allow for the possibility that identical sellers have different s. Suppose that the coalition assigns s_i to a fraction of ℓ_i sellers where $i = 1, \dots, I$, then the coalition's problem is

$$\max_{(s_i,\ell_i)} \quad \sum_{i=1}^{I} \ell_i \cdot z\mathcal{P}(\lambda x(z) s_i)$$
$$\text{s.t.} \quad \sum_{i=1}^{I} \ell_i s_i = 1.$$

Recall that $\mathcal{P}(\lambda)$ denotes the probability that seller z meets two or more effective buyers, as defined in equation (4). Since $\mathcal{P}(\lambda)$ is non-concave, the above problem is to maximize a convex combination of a non-concave function under a linear constraint, which implies that we need to find the concave hull of $\mathcal{P}(\lambda)$ (i.e., the least concave function that is greater than it).

Since $\mathcal{P}(\lambda)$ has an S-shape, its derivative is first increasing and then decreasing. So, its concave hull is given by

$$\widehat{\mathcal{P}}(\lambda) = \begin{cases} \frac{\lambda}{\Lambda_1}\mathcal{P}(\Lambda_1) & \text{if } \lambda \leq \Lambda_1 \\ \mathcal{P}(\lambda) & \text{if } \lambda \geq \Lambda_1. \end{cases} \tag{35}$$

That is, when $\lambda \leq \Lambda_1$, $\widehat{\mathcal{P}}(\lambda)$ is the line segment between point $(0, 0)$ and $(\Lambda_1, \mathcal{P}(\Lambda_1))$, and when $\lambda \geq \Lambda_1$, $\widehat{\mathcal{P}}(\lambda)$ coincides with $\mathcal{P}(\lambda)$. The threshold Λ_1 is determined by the condition that at $\lambda = \Lambda_1$, the slope of $\mathcal{P}(\lambda)$ equals the slope of the line segment, which implies that $\widehat{\mathcal{P}}(\lambda)$ has a continuous, decreasing derivative. Formally, Λ_1 solves $\mathcal{P}(\Lambda_1)/\Lambda_1 = \mathcal{P}'(\Lambda_1) = \Lambda_1 e^{-\Lambda_1}$, which is exactly $\Lambda(-1)$ defined in Lemma 2 (see equation (22)) and its value is close to 1.8.

The coalition's solution for the case of homogeneous sellers is then straightforward. If $\lambda x(z) \geq \Lambda_1$, then all sellers will have the same s. In contrast, if $\lambda x(z) < \Lambda_1$, then the coalition will allocate a fraction of ℓ sellers to $s_1 = 0$ (equivalent to excluding those sellers from participation) and a fraction $1 - \ell$ sellers to $s_2 = \Lambda_1/(\lambda x(z))$ where ℓ is such that $(1 - \ell)s_2 = 1$. Therefore, the coalition's solution coincides with the planner's solution from the last subsection if and only if $\lambda x(z) \geq \Lambda_1$.

Heterogeneous Sellers. We now consider the general case where sellers are heterogeneous. To simplify exposition, we assume that $F(z)$ is continuous. Suppose that the coalition optimally assigns arc length $s_c(z)$ to a seller of type z. The above result shows that for any active seller, the coalition will set the effective queue length $\lambda x(z)s_c(z) \geq \Lambda_1$. Furthermore, by Lemma 2, $\pi(s,z)$ is strictly increasing in z since the effective queue length is always at least $\Lambda_1 = \Lambda(-1)$, which implies that the coalition chooses a threshold $\widehat{z}$ below which $s(z)$ equals to zero. Of course, the coalition can set $\widehat{z} = z_0$, the lowest seller type, so that all sellers are active in this case.

Furthermore, since sellers with $z > \widehat{z}$ have an effective queue length of at least Λ_1, by Lemma 2 we have $\pi_{sz}(s^c(z), z) > 0$.[17] Thus at the coalition's solution, PAM between sellers and locations always holds. That is, $s_c(z)$ is increasing in z.[18]

Formally, the coalition chooses $s(z)$ for each seller type z to maximize total surplus:

$$\max_{s(z)} \quad \int_z z\widehat{\mathcal{P}}(\lambda x(z)s(z))dF(z) \tag{36}$$

subject to the constraint (31), where $\widehat{\mathcal{P}}(\cdot)$ is given by equation (35). The Lagrangian of this problem is,

$$\mathcal{L} = \int_z z\widehat{\mathcal{P}}(\lambda x(z)s(z))dF(z) + \zeta\lambda\left(1 - \int_z s(z)dF(z)\right)$$

where ζ is the modified multiplier (multiplied by λ). For each z, the optimal solution $s_c(z)$ must satisfy the following first-order condition,

$$zx(z)\widehat{\mathcal{P}}'(\lambda x(z)s_c(z)) = \zeta \quad \text{if} \quad \zeta \leq zx(z)\mathcal{P}'(\Lambda_1) \tag{37}$$

with complementary slackness (if $zx(z)\mathcal{P}'(\Lambda_1) < \zeta$, then $s_c(z) = 0$). Note that the above first-order condition is both necessary and sufficient for the coalition's solution.

Since $\widehat{\mathcal{P}}(\lambda)$ is first linear and then strictly concave, and $zx(z)$ is (weakly increasing), the coalition's solution is characterized by a threshold $\widehat{z}$ below which $s_c(z) = 0$. If $\zeta \leq z_0x(z_0)\mathcal{P}'(\Lambda_1)$, then $\widehat{z} = z_0$; if $\zeta > z_0x(z_0)\mathcal{P}'(\Lambda_1)$, then $\widehat{z}$ is determined by the condition $\zeta = \widehat{z}x(\widehat{z})\mathcal{P}'(\Lambda_1)$ (in this case $\lambda x(\widehat{z})s_c(\widehat{z}) = \Lambda_1$). When $z > \widehat{z}$, the above first-order condition

[17] Since the effective queue length is greater than 1, $\pi_{sz}(s^c(z), z) > 0$ irrespective of whether $\varepsilon_x(z) \geq -1/2$.

[18] This is different from the case of exogenous location distribution, where PAM between locations and sellers fails when the effective queue length for some sellers is small.

becomes

$$\zeta = zx(z)\widehat{\mathcal{P}}'(\lambda x(z)s_c(z)) = zx(z)\mathcal{P}'(\lambda x(z)s_c(z)) = zx(z)\cdot\lambda x(z)s_c(z)e^{-\lambda x(z)s_c(z)}. \tag{38}$$

Comparing the above equation with (7) implies that the derivative of rental price $r'(s) = \zeta$. The rental price in the coalition's solution is thus linear.

We now analyze under what condition the planner's solution coincides with the coalition's solution. Since the planner's problem is strictly concave, its solution $s_p(z)$ is always continuous. As we showed above, the coalition's problem is not concave, so that its solution $s_c(z)$ can have a jump. To see this, recall that z_0 is the lowest seller type. If $z_0 = 0$ (or sufficiently small), then at the planner's solution there exists a threshold z' such that the optimal $s_p(z') = 0$ and $x(z)s_p(z)$ then increase strictly and continuously for $z > z'$. At the coalition's solution, there exists a threshold z'' such that $\lambda x(z'')s_c(z'') = \Lambda_1$ and $x(z)s_c(z)$ increases strictly and continuously for $z > z''$. In this case, the two solutions must differ. Therefore, a necessary condition for the planner's solution to coincide with the coalition's solution is that z_0 is sufficiently large so that $\lambda x(z_0)s_p(z_0) \geq \Lambda_1$. Conditional on this, the two solutions coincide if and only if the planner' solution also satisfies the coalition's first-order condition (38). Comparing (38) with the planner's first-order condition implies that the two solutions coincide if and only if $zx(z)$ is constant for all $z \geq z_0$. The following proposition summarizes the above result.

Proposition 4. *The planner's solution coincides with the coalition's solution (the decentralized equilibrium outcome) if and only if* $zx(z) = z_0x(z_0)$ *for all* $z \geq z_0$ *and* $\lambda x(z_0)z_0/\overline{z} \geq \Lambda_1$, *where* $\overline{z}$ *is the mean of* z *(* $\overline{z} = \int_{z_0}^{\infty} zdF(z)$*). Furthermore, in this case* $s(z) = z/\overline{z}$.

Proof. See Appendix A.10. □

Suppose that the planner considers assigning additional Δs to a certain seller z. This increases surplus if and only if the seller meets no valuable buyers, in which case an additional buyer creates an expected surplus $zx(z)$. Since $zx(z)$ is constant, at the planner's optimum the probability that sellers meet no valuable buyers must be the same across different sellers. That is, the effective queue length $\lambda x(z)s_p(z)$ is constant across sellers. The coalition faces a different trade-off when considering to assign additional Δs to a certain seller z. It creates value for the seller if and only if the seller already meets with exactly one valuable buyer. In that case, an additional buyer creates an expected surplus $zx(z)$. In the special case that the effective queue length is constant across sellers, the probability that the sellers meet exactly one valuable buyer is also constant. Hence the marginal tradeoffs in terms of allocating arc length to sellers (intensive margin) for the planner and the coalition are exactly the same. However, the coalition would like to exclude some sellers from participation to take advantage of increasing returns to scale when the effective queue length is too small (we can think of determining who is active or not as adjustments along the extensive margin). If the measure of buyers is large enough, the coalition does not exclude sellers from the market. Then the two solutions coincide.

Next, we compare the planner's solution with the coalition's solution for the case where $zx(z)$ is strictly increasing.

Proposition 5. *Suppose that $zx(z)$ is strictly increasing. Then there exists some $z_1 \geq \widehat{z}$ such that $s_c(z) > s_p(z)$ for $z > z_1$, and $s_c(z) \leq s_p(z)$ for $z < z_1$ (the latter "$\leq$" holds as "$=$" only when $s_p(z) = 0$), where $\widehat{z}$ is the minimal type of active sellers in the coalition's solution.*

Proof. See Appendix A.11. □

When $zx(z)$ is strictly increasing, the planner will assign a longer effective queue length to sellers with a higher z such that the marginal contribution to surplus of arc lengths is the same across sellers. Consider two sellers a and b with $z_a < z_b$. At the planner's solution, the marginal value of arc-lengths must be the same across the two values: $z_a x(z_a) e^{-\widetilde{\lambda}_{p,a}} = z_b x(z_b) e^{-\widetilde{\lambda}_{p,b}}$, with effective queue lengths $\widetilde{\lambda}_{p,a} = \lambda x(z_a) s_p(z_a)$ and $\widetilde{\lambda}_{p,b} = \lambda x(z_b) s_p(z_b)$, which implies that $\widetilde{\lambda}_{p,a} < \widetilde{\lambda}_{p,b}$. From the coalition's point of view, the marginal value of arc length is higher at seller b in the planner's solution since $z_a x(z_a) \widetilde{\lambda}_a e^{-\widetilde{\lambda}_a} < z_b x(z_b) \widetilde{\lambda}_b e^{-\widetilde{\lambda}_b}$, since $\widetilde{\lambda}_b > \widetilde{\lambda}_a$. As a result, compared with the planner's solution, the coalition assigns a longer arc length to high-z sellers (more buyers has a relatively large effect on moving from one to two effective buyers) and a shorter arc length to low-z sellers, while the total arc length is fixed.

Comparison with Directed Search. In our model, sellers can increase their trading probability by investing in their queue length. This feature of our model is reminiscent of directed search models. However, the equilibrium in our model is in general not efficient, while directed search models of competing auctions are (see e.g. Albrecht et al., 2014). What explains this difference? An important difference is that under spatial sorting, high-quality sellers can attract more buyers by paying more for a good location but this extra payment does not go to the buyers. In a directed search equilibrium, high-quality sellers can attract more buyers by offering them a better deal. That is, they can transfer value directly to the buyers. In the terminology of Cai et al. (2023), directed search allows sellers to buy queues in a competitive market. The inefficiency in the spatial sorting equilibrium is driven by the fact that sellers want to meet two or more buyers (if they meet only one buyer, all the surplus goes to the buyer). If λ and or x are very small, most sellers either meet one or zero buyers who like their product. Moving to a better spot increases the probability that a seller meets with one rather than no buyer which creates social surplus but sellers are not willing to pay for the better location because all the surplus goes to the buyers in this case. In contrast, when $\lambda x(z)$ is large, individual sellers are willing to pay for a better location to increase the likelihood of meeting two or more buyers rather than one but from a social point this is just rent-seeking and in the decentralized equilibrium there will be too much dispersion in locations which generates too few matches.

5 Conclusion

In this paper, we developed a tractable random search model that takes into account that some locations are better than others. We then use our framework to study what type of sellers benefit from good locations and how this translates into spatial sorting. Having heterogeneous regions leads to less trade but possibly to more quality-weighted trade. The resulting equilibrium can be inefficient because of a rent seeking externality. When there are few buyers per seller, sellers of general products benefit most from the good locations because when two or more buyers arrive, they receive the full surplus while for a social planner, one buyer who likes the product is enough and any additional buyer visit adds nothing to surplus. For sellers of niche products, the good spots in this case, mainly increase the probability of one buyer who likes their product but the probability of two effective buyer arrivals remains close to zero.

We defined x to be the probability that a buyer likes the product but alternatively, we can think of x as the probability that a credit-constrained buyer can afford the product. Within cities, sellers of expensive high-quality products would still have incentives to move to good locations. At those good locations, expensive niche products will be offered that the poor cannot afford. In the other areas, more standard goods will be offered. The buyers and sellers in other areas (with fewer buyers per seller) will trade more standard products (H&M, McDonalds) but they can also benefit because those areas will be cheaper. The welfare effects also depend on who receives the rents from the good locations, but we leave a more detailed analysis of that for future research.

Another novel implication of our model is that both the aggregate meeting and matching function, which maps the measures of buyers and sellers into respectively meetings and trades, is ultimately driven by the distribution of product quality. Finally, we showed that when the location distribution is endogeneous, the market outcome need not be efficient because sometimes sellers invest too much in good locations in order to increase the likelihood of multiple buyers which increases the price.

Appendix A Additional Results and Omitted Proofs

A.1 Finite foundations for continuous $L(s)$

A.1.1 $L(s)$ is an exponential distribution.

In this case, sellers are independently and uniformly allocated on the circle. Consider a particular seller i. The arc distance between seller i and any other seller is then a uniform distribution on $[0,1]$. Therefore, d_i is the minimum of $N_s - 1$ independent random variables which all follow $U[0,1]$, which implies that $\mathbb{P}(d_i \leq d) = 1-(1-d)^{N_s-1}$ and thus $\mathbb{P}(d_i N_s \leq s) = 1-(1-s/N_s)^{N_s-1}$. Since $\lim_{N_s\to\infty} 1-(1-s/N_s)^{N_s-1} = 1-e^{-s}$, as the market gets large, $L(s) = 1-e^{-s}$.

A.1.2 $L(s)$ is an arbitrary distribution.

Given any distribution $L(s)$ with mean 1, we now construct a model which determines d_i for each seller as follows. Sellers first independently draw a random variable D from the distribution $L(\cdot)$. Set $d_i = D_i/\sum_{j=1}^{N_s} D_j$. Next, let the market get large and derive the limit distribution of $N_s d_i$:

$$N_s d_i = N_s \cdot \frac{D_i}{\sum_{j=1}^{N_s} D_j} = \frac{D_i}{\frac{1}{N_s}(\sum_{j=1}^{N_s} D_j)} \to D_i$$

By the law of large numbers, the denominator approaches 1, the mean of L. Hence $N_s d_i$ approaches the distribution of D_i, which is $L(\cdot)$.

A.2 Proof of Equation 12

Since $P_n(\lambda)$ is given by equation (10), we have

$$\begin{aligned} 1-\sum_{n=0}^{\infty} P_n(\lambda)(1-x)^n &= 1-\sum_{n=0}^{\infty}\left(\int_s e^{-\lambda s}\frac{(\lambda s)^n}{n!}\,dL(s)\right)(1-x)^n \\ &= 1-\int_s\left(\sum_{n=0}^{\infty} e^{-\lambda s}\frac{(\lambda s)^n(1-x)^n}{n!}\right)dL(s) = 1-\int_s e^{-\lambda x s}\,dL(s) = \int_s 1-e^{-\lambda x s}\,dL(s) \end{aligned}$$

where for the first equality in the second row, we interchanged summation and integration, and the second equality follows from the fact that the Poisson probabilities sum up to 1. The last term on the right-hand side is exactly $m(\lambda x)$ by definition. □

A.3 Proof of Theorem 1

By equation (15), we have

$$Y(\lambda) = \int_s z\left(1 - e^{-\lambda s x(z)}\right) dL(s) \leq z\left(1 - e^{-\lambda x(z)\int_s s dL(s)}\right) = z\left(1 - e^{-\lambda x(z)}\right)$$

where for the inequality we used Jensen's inequality, and the second equality follows from the fact that the mean of s is always 1. The inequality holds with equality if and only if $L(s)$ is degenerate. The last term on the right-hand side is exactly the total surplus generated by the urn-ball meeting technology. □

A.4 Proof of Theorem 2

Invariance implies (18). For the first part of the proof, consider an invariant meeting technology defined by the condition that $\{P_0(\lambda), P_0(\lambda), \dots\}$ satisfies equation (12). As we argued before Theorem 2, $m(\lambda)$ is a Bernstein function. By Bernstein's theorem, the function $m(\lambda)$ has the following Lévy-Khintchine representation:

$$m(\lambda) = a_1 + a_2\lambda + \int_{(0,\infty)} \left(1 - e^{-\lambda s}\right) dL(s),$$

where $a_1, a_2 \geq 0$ and L is a measure on $(0, \infty)$ satisfying $\int_{(0,\infty)} \min\{1, t\}\, dL(s) < \infty$ (see Theorem 3.2 of Schilling et al., 2012).

Since $m(0) = 0$, it follows that $a_1 = 0$. Moreover, since $m(\lambda)$ is bounded from above by 1, a_2 must equal 0 as well. Further, if $\lambda \to \infty$, we have $1 - e^{-\lambda s} \nearrow 1$ for any $t > 0$, and therefore $m(\lambda) \to \int_{(0,\infty)} 1\, dL(s)$ by the monotone convergence theorem. Since $m(\lambda)$ cannot exceed 1, the total measure of $L(\cdot)$ must be less or equal to 1: $\int_{(0,\infty)} 1\, dL(s) \leq 1$. If the total measure is strictly less than 1, without loss of generality we can assign measure $1 - \int_{(0,\infty)} 1\, dL(s)$ on point $s = 0$. Therefore, it is without loss of generality to assume that $L(\cdot)$ is a probability measure on $[0, \infty)$.

Next, the probability that a worker meets a firm is $m(\lambda)/\lambda = \int_0^\infty \left(1 - e^{-\lambda s}\right)/\lambda\, dL(s)$, which cannot exceed 1 for any $\lambda \geq 0$. One can easily verify that when $\lambda \searrow 0$, we have $\left(1 - e^{-\lambda s}\right)/\lambda \nearrow s$. Therefore, $\lim_{\lambda \to 0} m(\lambda)/\lambda = \int_0^\infty s\, dL(s)$ by the monotone convergence theorem. Hence, $L(\cdot)$ must satisfy $\int_0^\infty s\, dL(s) \leq 1$.

(18) implies invariance. For the second part of the proof, assume that $m(\lambda)$ is given by equation (18) where L is a probability measure on $[0, \infty)$ satisfying $\int_{[0,\infty)} s dL(s) \leq 1$. Since this corresponds to our meeting process on the circle, the resulting meeting technology is invariant (see equation (12)) and $P_n(\lambda)$ is given by equation (10). □

A.5 Proof of Lemma 1

Since $S(s,z)$ is given by equation (2), we have

$$\frac{\partial S(s,z)}{\partial s} = zx(z)\lambda e^{-\lambda s x(z)} > 0$$

Similarly,

$$\begin{aligned}\frac{\partial S(s,z)}{\partial z} &= e^{-\lambda s x(z)}\left(\lambda s z x'(z) - 1 + e^{\lambda s x(z)}\right) > e^{-\lambda s x(z)}\left(\lambda s z x'(z) - 1 + 1 + \lambda s x(z)\right) \\ &= e^{-\lambda s x(z)}\lambda s x(z)\left(\frac{zx'(z)}{x(z)} + 1\right) \geq 0\end{aligned}$$

where for the last inequality we used the fact that $zx(z)$ is weakly increasing. Furthermore,

$$\frac{\partial^2 S(s,z)}{\partial s \partial z} = \lambda e^{\lambda(-s)x(z)}\left(1 + \frac{zx'(z)}{x(z)} - \frac{zx'(z)}{x(z)}\lambda s x(z)\right) \geq \lambda e^{\lambda(-s)x(z)}\left(-\frac{zx'(z)}{x(z)}\lambda s x(z)\right) \geq 0$$

where for the first inequality we used the fact that $zx(z)$ is weakly increasing and for the second inequality that $x(z)$ is weakly decreasing. Apparently, the two weak inequalities can not hold with equality at the same time, which implies that $S(s,z)$ is *strictly* supermodular. □

A.6 Proof of Lemma 2

Note that $\pi(s,z)$ is given by equation (3). We have

$$\frac{\partial \pi(s,z)}{\partial s} = \lambda^2 s z x(z)^2 e^{-\lambda s x(z)} > 0$$

Similarly,

$$\frac{\partial \pi(s,z)}{\partial z} = e^{\lambda(-s)x(z)}\left(e^{\lambda s x(z)} - 1 - \lambda s x(z) + (\lambda s x(z))^2 \varepsilon_x(z)\right)$$

where $\varepsilon_x(z) = zx'(z)/x(z) \in [-1, 0]$. The above equation is positive if and only if

$$\frac{e^{\lambda s x(z)} - 1 - \lambda s x(z)}{(\lambda s x(z))^2} > -\varepsilon_x(z)$$

Define $\widetilde{\lambda} = \lambda s x(z)$ (the effective queue length). Then the derivative of the left-hand side with respect to $\widetilde{\lambda}$ is given by $(2 + \widetilde{\lambda} + e^{\widetilde{\lambda}}(\widetilde{\lambda} - 2))/\widetilde{\lambda}^3$. Note that $2 + \widetilde{\lambda} + e^{\widetilde{\lambda}}(\widetilde{\lambda} - 2)$ is a convex function whose second-order derivative is $e^{\widetilde{\lambda}}\widetilde{\lambda}$ and its derivative at $\widetilde{\lambda} = 0$ equals zero.

Therefore, $2+\widetilde{\lambda}+e^{\widetilde{\lambda}}(\widetilde{\lambda}-2)>0$ and the left-hand side is strictly increasing in $\widetilde{\lambda}$. Therefore, for each given $\varepsilon_x(z)$, there exists a threshold $\Lambda(\varepsilon_x(z))$ such that $\pi_z(s,z)>0$ if and only if $\lambda x(z)s>\Lambda(\varepsilon_x(z))$. The threshold $\Lambda(\varepsilon_x(z))$ is such that the above inequality holds with equality, i.e., it is the solution to (22).

Next, we have

$$\frac{\partial^2\pi(s,z)}{\partial s\partial z}=\lambda^2 sx(z)^2 e^{\lambda(-s)x(z)}\left(1+\left(2-\lambda x(z)s\right)\varepsilon_x(z)\right) \tag{39}$$

Since $\varepsilon_x(z)\le 0$, the above is strictly positive if $\varepsilon_x(z)\ge -1/2$. If $\varepsilon_x(z)\in[-1,-1/2)$, then the above is strictly positive if and only if $\lambda x(z)s>2+1/\varepsilon_x(z)$. □

A.7 Proof of Lemma 3

Since $Y(\lambda,\alpha,s_0)$ is given by equation (24),

$$\frac{\partial Y(\lambda,\alpha,s_0)}{\partial\alpha}=\overline{z}\frac{-(1-s_0)\lambda e^{-\lambda s_0}}{\left((\alpha-1)+\alpha\lambda(1-s_0)\right)^2}<0,$$

where $s_0<1$. Note that in the above analysis, we change the quality distribution while holding fixed the mean quality $\overline{z}$ and the measure of buyers.

Similarly, we have

$$\frac{\partial Y(\lambda,\alpha,s_0)}{\partial s_0}=\overline{z}\frac{(\alpha-1)\lambda e^{-\lambda s_0}}{\left((\alpha-1)+\alpha\lambda(1-s_0)\right)^2}\left((1-s_0)\alpha\lambda-1\right).$$

The sign of $\partial Y(\lambda,\alpha,s_0)/\partial s_0$ is thus completely determined by the term $(1-s_0)\alpha\lambda-1$. When $\alpha\lambda\le 1$, total surplus is maximal at $s_0=0$. When $\alpha\lambda>1$, then total output is maximal at $s_0=1-1/(\alpha\lambda)$. □

A.8 Proof of Proposition 3

When $s_0=1$, as we argued before Proposition 3, $\Pi_1(z)=z(1-e^{-\lambda_1}-\lambda_1 e^{-\lambda_1})$ where $\lambda_1=\log(\overline{z}/K)$. Plugging $\overline{z}=z_0\alpha/(\alpha-1)$ and $z_0=K=1$ into $\Pi_1(z)$ then yields equation (27). Since $\alpha>1$, the slope of $\Pi_1(z)$ is strictly smaller than 1.

Next, consider the case $s_0=0$. First, we calculate the price of locations explicitly.[19]

[19] When $s_0\in(0,1)$, $r(s)$ can also be calculated explicitly with a more complicated expression.

Plugging $z^*(s)$ from equation (23) into equation (7) yields,

$$r_0'(s) = z_0 \lambda_0^2 s e^{-s(\lambda_0 - 1/\alpha)}.$$

where the subscript 0 denotes the case $s_0 = 0$. Solving the above differential equation yields:

$$r_0(s) = z_0 \left(1 + \frac{1}{\alpha(\lambda_0 - \frac{1}{\alpha})}\right)^2 \left(1 - e^{-s(\lambda_0 - \frac{1}{\alpha})} - s(\lambda_0 - \frac{1}{\alpha})e^{-s(\lambda_0 - \frac{1}{\alpha})}\right)$$

where $\lambda_0 - \frac{1}{\alpha} = \sqrt{\frac{z_0}{K}} - 1$. The average location price is given by

$$R_0 = \int_0^\infty r_0(s)\, d(1 - e^{-s}) = z_0 \left(\frac{\alpha\lambda_0}{\alpha\lambda + \alpha - 1}\right)^2 = K\left(\sqrt{\frac{z_0}{K}} - \frac{\alpha - 1}{\alpha}\right)^2,$$

Since we assume that $K = z_0 = 1$, the above two equations imply $r(s) = s^2/(2\alpha^2)$ and $R_0 = 1/\alpha^2$. By equation (23), we have $s^*(z) = \alpha \log(z)$. Plugging $s^*(z)$, $r_0(s^*(z))$ and R_0 into equation (9) then yields $\Pi_0(z)$ in equation (26).

Since $\Pi_0(z)$ is given by equation (26), we have $\Pi_0'(z) = \frac{z - \log z - 1}{z}$. When $z \to \infty$, $\Pi_0'(z) \to 1$. Furthermore, $\Pi_0''(z) = \log z / z^2 > 0$ since $z > z_0 = 1$; hence $\Pi_0(z)$ is strictly convex.

Define $\Delta\Pi(z) = \Pi_0(z) - \Pi_1(z)$. Since $\Pi_0(z)$ is strictly convex and $\Pi_1(z)$ is linear, $\Delta\Pi(z)$ is strictly convex. When $z = 1$, $\Pi_0(1) = 1/\alpha^2$ and $\Pi_1(1) = \frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right) \log\left(1 - \frac{1}{\alpha}\right)$, which implies that $\Delta\Pi(1) = \left(1 - \frac{1}{\alpha}\right)\left[-\log\left(1 - \frac{1}{\alpha}\right) - \frac{1}{\alpha}\right] > 0$ since $-\log(1 - x) > x$ for $x \in (0, 1)$. When z is sufficiently large, $\Delta\Pi(z) \to \infty$ since $\lim_{z\to\infty} \Pi_0'(z) = 1$, and $\Pi_1(z)$ is linear with a slope strictly smaller than 1, .

Next, consider the derivative of $\Delta\Pi(z)$:

$$\Delta\Pi'(z) = \frac{z - \log z - 1}{z} - \left(\frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right) \log\left(1 - \frac{1}{\alpha}\right)\right), \tag{40}$$

where the second term on the right-hand side is $\Pi_1'(z)$. When $z = 1$, $\Delta\Pi'(1) = -\Pi_1'(z) < 0$. Hence there exists some $z_m > 1$ such that $\Delta\Pi'(z_m) = 0$ or equivalently $\Delta\Pi(z_m)$ reaches its minimum at z_m. $\Delta\Pi(z) > 0$ for all z if and only if $\Delta\Pi(z_m) > 0$. Since $\Delta\Pi'(z_m) = 0$,

$$\Delta\Pi(z_m) = \Pi_0(z_m) - z_m \frac{z_m - \log z_m - 1}{z_m} = \frac{1}{\alpha^2} - \frac{(\log z_m)^2}{2}$$

which implies that $\Delta\Pi(z_m) > 0$ if and only if $z_m < e^{\sqrt{2}/\alpha}$, which is the case if and only if

$\Delta\Pi'(e^{\sqrt{2}/\alpha}) > 0$. By equation (40),

$$\Delta\Pi'(e^{\sqrt{2}/\alpha}) = \mathcal{T}(\alpha) \equiv 1 - \left(\frac{\sqrt{2}}{\alpha} + 1\right) e^{-\sqrt{2}/\alpha} - \left(\frac{1}{\alpha} + \left(1 - \frac{1}{\alpha}\right) \log\left(1 - \frac{1}{\alpha}\right)\right),$$

When $\alpha = 1$, the value of $\mathcal{T}(\alpha)$ defined in the above equation is strictly negative, and when $\alpha \to \infty$, $\mathcal{T}(\alpha) \to 0$. Note that the derivative of $\mathcal{T}(\alpha)$ is given by

$$\mathcal{T}'(\alpha) = \frac{-2e^{-\frac{\sqrt{2}}{\alpha}} - \alpha \log\left(1 - \frac{1}{\alpha}\right)}{\alpha^3}.$$

The numerator above is strictly convex since its second-order derivative is given by $\frac{4\left(\sqrt{2}\alpha - 1\right)e^{-\frac{\sqrt{2}}{\alpha}}}{\alpha^4} + \frac{1}{(\alpha-1)^2\alpha}$, which is strictly positive since $\alpha > 1$. When $\alpha = 1$, the numerator is strictly positive; when $\alpha \to \infty$, the numerator is strictly negative (the limit equals -1). Since the numerator is strictly convex, $\mathcal{T}'(\alpha)$ is first positive and then negative, which implies that $\mathcal{T}(\alpha)$ is first increasing and then decreasing. Since its limit is zero, its maximum is strictly positive. Since at $z = 1$ its value is strictly negative, thus there exists a unique α^* such that $\mathcal{T}(\alpha^*) = 0$; when $\alpha > \alpha^*$, $\mathcal{T}(\alpha)$ first increases and then decreases (with a limit equal zero). Thus $\mathcal{T}(\alpha) > 0$ if and only if $\alpha > \alpha^*$. Numerical results suggest that α^* is approximately 1.9.

A.9 Derivations for Section 3.3.3

As in the case of homogeneous sellers, we can analyze the effect of making goods more niche. To do so, we again increase z and simultaneously reduce $x(z)$, keeping the distribution $G(q)$ of expected buyer value $q = zx(z)$ fixed. Because we fix the distribution of q, the correspondence between sellers and locations does not depend on γ and can thus be denoted by $q^*(s)$, which is given by a variant of equation (19), i.e., $1 - G(q^*(s)) = 1 - L(s)$.

Let $\gamma^* = 0$ for the first example, and $\gamma^* = \alpha/\beta$ for the second example; then the first-order approximation of equation (20) is

$$Y(\lambda, \gamma^* + \Delta\gamma) \approx Y(\lambda, \gamma^*) + \int_s q^*(s) \frac{1 - e^{-\lambda s x^*(s)} - \lambda s x^*(s) e^{-\lambda s x^*(s)}}{(x^*(s))^2} \Delta x^*(s)\, dL(s), \qquad (41)$$

where $\Delta x^*(s) = x(q^*(s), \gamma^* + \Delta\gamma) - x(q^*(s), \gamma^*)$. Given $x(q, \gamma) = (q/z_0)^{-\gamma/(1-\gamma)}$, we have

$$\Delta x^*(s) = -\frac{1}{(1-\gamma^*)^2} \left(\frac{z_0}{q}\right)^{\frac{\gamma^*}{1-\gamma^*}} \log\left(\frac{q}{z_0}\right) \Delta\gamma. \qquad (42)$$

First-order approximation for Example 1. We now consider a first-order approximation around $\gamma = 0$ for Example 1 above. To simplify the analysis we set $s_0 = 0$, so

$L(s) = 1 - e^{-s}$. Furthermore, we fix the distribution of expected buyer value ($q = zx(z)$): $G(q) = 1 - \left(\frac{z_0}{q}\right)^{\alpha}$ with $\alpha > 1$ and $q \geq z_0$. By equation (23), the assignment between sellers and locations is then $q^*(s) = z_0 e^{s/\alpha}$, which, by equation (42), implies that $\Delta x^*(s) \approx s/\alpha \cdot \Delta\gamma$. The first-order approximation in (41) then becomes

$$Y(\lambda, \Delta\gamma) = \overline{z}\frac{\alpha\lambda}{\alpha\lambda + \alpha - 1} + \overline{z}\frac{\alpha^2\lambda^2(\alpha\lambda + 3(\alpha - 1))}{(\alpha - 1)(\alpha\lambda + \alpha - 1)^3}\Delta\gamma,$$

where $\overline{z} = z_0\alpha/(\alpha - 1)$, and the first term on the right-hand side corresponds to Y in equation (24) with $s_0 = 0$. Note that we have included $\Delta\gamma$ as an argument of Y to emphasize its dependence on γ. By the above equation, the percentage increase in Y is then given by

$$\frac{\partial \log Y(\lambda, \gamma)}{\partial\gamma}\Big|_{\gamma=0} = \frac{\alpha\lambda(\alpha\lambda + 3(\alpha - 1))}{(\alpha - 1)(\alpha\lambda + \alpha - 1)^2}, \tag{43}$$

which is strictly positive since $\alpha > 1$. Note that the above equation is strictly decreasing in α.[20] Recall that a higher γ leads to a higher percentage change of $z(q)$ for a higher q. Therefore, the effect of γ on total surplus is stronger when α is smaller (when the quality distribution $G(q)$ is more dispersed). When the quality distribution is concentrated at z_0 ($\alpha \to \infty$), equation (43) converges to zero, since $x(z_0)$ is always 1 for any γ.

First-order approximation for Example 2. At $\gamma^* = \alpha/\beta$, the distribution of the expected buyer value $q = zx(z)$ is given by $1 - G(q) = \mathbb{P}(zx(z) \geq q) = \mathbb{P}(z/z_0 \geq (q/z_0)^{1/(1-\gamma^*)}$, which implies that $G(q) = 1 - (z_0/q)^{\alpha/(1-\gamma^*)}$. Furthermore, by equation (28), we have $q^*(s) = z_0(\frac{s}{s_0})^{\frac{1-\gamma}{\gamma}}$ and $x^*(s) = (q^*(s)/z_0)^{-\gamma/(1-\gamma)} = s_0/s$.

We now fix $G(q)$ and increase γ from γ^* to $\gamma^* + \Delta\gamma$. Since $q^*(s) = z_0(\frac{s}{s_0})^{\frac{1-\gamma}{\gamma}}$, by equation (42) we have

$$\Delta x^*(s) = -\frac{s_0 \log(\frac{s}{s_0})}{s\gamma^*(1 - \gamma^*)}\Delta\gamma$$

which implies that the percentage reduction $\Delta x^*(s)/x^*(s) = -\Delta\gamma \cdot \log(\frac{s}{s_0})/\gamma^*(1 - \gamma^*)$, which is higher for sellers with higher quality.

[20] To see this, note that the derivative of (43) with respect to α is given by

$$-\frac{\Delta\gamma\lambda^3}{(\alpha - 1)^2(\alpha\lambda + \alpha - 1)^3} - \frac{\Delta\gamma\lambda\left((\alpha - 1)^2(\lambda + 1)(\lambda + 3) + (\alpha - 1)(3\lambda^2 + 7\lambda + 6) + 3\lambda(\lambda + 1)\right)}{(\alpha - 1)(\alpha\lambda + \alpha - 1)^3}$$

which is strictly negative since $\alpha > 1$.

Plugging the above expression of $\Delta x^*(s)$ into (41) yields

$$Y(\lambda, \gamma^* + \Delta\gamma) \approx \left(1 - e^{-\lambda s_0}\right)\overline{z} + \frac{\overline{z}\left(1 - e^{-\lambda s_0} - \lambda s_0 e^{-\lambda s_0}\right)}{(1-\gamma^*)(\alpha - 1)}\Delta\gamma$$

Again, increasing γ while holding the distribution of q constant increases the expected total surplus. The effect is smaller when α is higher.

A.10 Proof of Proposition 4

As we argued before Proposition 4, a necessary condition for the planner's solution to coincide with the coalition's solution is that $\lambda x(z_0)s_p(z_0) > \Lambda_1$. Furthermore, since $s_p(z)$ needs to also satisfy equation (38), comparing (38) with the planner's FOC (32) shows that $zx(z)$ needs to be constant for $z \geq z_0$. If the above two necessary conditions are satisfied, they are also sufficient since the first-order conditions are not only necessary but also sufficient. Finally, by Lemma 4 we have $s_p(z) = z/\overline{z}$, which is also the coalition's solution $s_c(z)$.□

A.11 Proof of Proposition 5

The planner's solution $s_p(z)$ is a continuously increasing function which is determined by the first-order condition (32): $\mathcal{F}^p(s,z) = \xi$, where $\mathcal{F}^p(s,z) = zx(z)e^{-\lambda x(z)s}$. The coalition's solution $s_c(z)$ has one possible jump point $\widehat{z}$ below which $s_c(z)$ is zero and above which $s_c(z)$ is continuously increasing with $\lambda x(z)s_c(z) \geq \Lambda_1$ and it is given by the coalition's first-order condition (38): $\mathcal{F}^c(s,z) = \zeta$, where $\mathcal{F}^c(s,z) = zx(z)\lambda x(z)se^{-\lambda x(z)s}$. Of course, it may be the case that $\widehat{z}$ does not exist in which case we set $\widehat{z}$ to be z_0.

Next, we show that the level curves of $\mathcal{F}^p(s,z)$ and $\mathcal{F}^c(s,z)$ satisfy the single-crossing property.

$$-\frac{\partial\mathcal{F}^c(s,z)/\partial z}{\partial\mathcal{F}^c(s,z)/\partial s} - \left(-\frac{\partial\mathcal{F}^p(s,z)/\partial z}{\partial\mathcal{F}^p(s,z)/\partial s}\right) = \frac{1+\varepsilon_x(z)}{\lambda zx(z)\left(\lambda x(z)s - 1\right)}$$

where $\varepsilon_x(z) = zx'(z)/x(z) > -1$ (we assume that $zx(z)$ is strictly increasing). Since at the coalition's interior solution, $\lambda x(z)s \geq \Lambda_1 > 1$, the above equation is always strictly positive.

If $\widehat{z} = z_0$, then $s_c(z)$ is continuous for $z \geq z_0$. Since $1 = \int_z s_p(z)dF(z) = \int_z s_c(z)dF(z)$, by continuity there exists some $z_1 > z_0$ such that $s_p(z_1) = s_c(z_1)$, denote which by s_1. As we showed above, the level curve $\mathcal{F}^p(s,z) = \mathcal{F}^p(s_1,z_1)$ crosses the level curve $\mathcal{F}^c(s,z) = \mathcal{F}^c(s_1,z_1)$ once and from below. Thus $s_c(z) > s_p(z)$ for $z > z_1$ and $s_c(z) < s_p(z)$ for $z < z_1$.

Next, suppose that $\widehat{z} > z_0$. If $s_c(\widehat{z}) \geq s_p(\widehat{z})$, then $s_c(z) > s_p(z)$ for all $z > \widehat{z}$, otherwise the level curve $\mathcal{F}^c(s,z) = \zeta$ must cross the level curve $\mathcal{F}^p(s,z) = \xi$ from above at some point. If $s_c(\widehat{z}) < s_p(\widehat{z})$, by continuity there exists some $z_1 > \widehat{z}$ such that $s_p(z_1) = s_c(z_1)$ since

$1 = \int_{z \geq \widehat{z}} s_c(z) dF(z) \geq \int_{z \geq \widehat{z}} s_p(z) dF(z)$. Again by the single-crossing property, $s_c(z) > s_p(z)$ for $z > z_1$ and $s_c(z) < s_p(z)$ for $\widehat{z} \leq z < z_1$. When $z < \widehat{z}$, $s_c(z) = 0$ so that $s_c(z) \leq s_p(z)$ continuous to hold, and in this case $s_c(z) = s_p(z)$ if and only if $s_p(z) = 0$. □

References

Albrecht, J., Menzio, G., and Vroman, S. (2023). Vertical differentiation in frictional product markets. *Journal of Political Economy Macroeconomics*, forthcoming.

Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2014). Efficient entry with competing auctions. *American Economic Review*, 104(10):3288–3296.

Becker, G. S. (1973). A theory of marriage: Part I. *Journal of Political Economy*, 81(4):813–846.

Becker, L. and Mangin, S. (2023). The effect of search frictions on extreme outcomes. mimeo.

Burdett, K. and Judd, K. L. (1983). Equilibrium price dispersion. *Econometrica*, 51:955–969.

Cai, X., Gautier, P., and Wolthoff, R. (2023). Meetings and mechanisms. *International Economic Review*, 64(1):155–185.

Cai, X., Gautier, P. A., and Wolthoff, R. P. (2017). Search frictions, competing mechanisms and optimal market segmentation. *Journal of Economic Theory*, 169:453 – 473.

Chiappori, P.-A. (2017). *Matching with transfers: The economics of love and marriage*. Princeton University Press.

Combes, P.-P., Duranton, G., and Gobillon, L. (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics*, 63(2):723–742.

Dauth, W., Findeisen, S., Moretti, E., and Suedekum, J. (2022). Matching in Cities. *Journal of the European Economic Association*, 20(4):1478–1521.

Davis, S. J. and de la Parra, B. S. (2017). Application flows. mimeo.

Eeckhout, J. and Kircher, P. (2010). Sorting vs screening - search frictions and competing mechanisms. *Journal of Economic Theory*, 145:1354–1385.

Ellison, G. and Glaeser, E. L. (1997). Geographic concentration in u.s. manufacturing industries: A dartboard approach. *Journal of Political Economy*, 105(5):889–927.

Ellison, G., Glaeser, E. L., and Kerr, W. R. (2010). What causes industry agglomeration? evidence from coagglomeration patterns. *American Economic Review*, 100(3):1195–1213.

Fujita, M. and Thisse, J.-F. (2002). *Economics of Agglomeration: Cities, Industrial Location, and Regional Growth*. Cambridge University Press.

Gautier, P. A., Svarer, M., and Teulings, C. N. (2010). Marriage and the city: Search frictions and sorting of singles. *Journal of Urban Economics*, 67(2):206–218.

Gautier, P. A. and Teulings, C. N. (2003). An Empirical Index for Labor Market Density. *The Review of Economics and Statistics*, 85(4):901–908.

Gautier, P. A. and Teulings, C. N. (2009). Search and the city. *Regional Science and Urban Economics*, 39(3):251–265.

Helsley, R. W. and Strange, W. C. (1990). Matching and agglomeration economies in a system of cities. *Regional Science and Urban Economics*, 20(2):189–212.

Kim, S. (1989). Labor specialization and the extent of the market. *Journal of Political Economy*, 97(3):692–705.

Lester, B., Visschers, L., and Wolthoff, R. (2015). Meeting technologies and optimal trading mechanisms in competitive search markets. *Journal of Economic Theory*, 155:1–15.

Menzio, G. (2023). Optimal product design: Implications for competition and growth under declining search frictions. *Econometrica*, 91(2):605–639.

Moretti, E. (2012). *The new geography of jobs.* Houghton Mifflin Harcourt.

Neiman, B. and Vavra, J. (2023). The rise of niche consumption. *American Economic Journal: Macroeconomics*, forthcoming.

Pissarides, C. A. (2000). *Equilibrium Unemployment Theory.* Oxford University Press.

Rosenthal, S. S. and Strange, W. C. (2004). Chapter 49 - evidence on the nature and sources of agglomeration economies. In Henderson, J. V. and Thisse, J.-F., editors, *Cities and Geography*, volume 4 of *Handbook of Regional and Urban Economics*, pages 2119–2171. Elsevier.

Schilling, R., Song, R., and Vondraček, Z. (2012). *Bernstein Functions: Theory and Applications.* De Gruyter studies in mathematics. W. De Gruyter, 2nd edition.