Meetings and Mechanisms*

Xiaoming Cai† Pieter Gautier‡ Ronald Wolthoff§

May 30, 2022

Abstract

This paper shows how meeting frictions affect equilibrium trading mechanisms and allocations in an environment where identical sellers post mechanisms to compete for buyers with ex-ante heterogeneous private valuations. Multiple submarkets can emerge, each consisting of all sellers posting a particular mechanism and the buyers who visit those sellers. Under mild conditions, high-valuation buyers are all located in the same submarket, and low valuation buyers can be in: (i) the same submarket, (ii) a different submarket and (iii) a mixture of (i) and (ii). The decentralized equilibrium is efficient when sellers can post auctions with reserve prices or entry fees.

JEL classification: C78, D44, D83.

Keywords: search frictions, directed search, matching function, meeting technology, competing mechanisms, competing auctions.

Running Head Title: Meetings and Mechanisms

*Manuscript received June 2021; revised March 2022.

We thank two anonymous referees, an associate editor, the editor (Dirk Krueger), Philipp Kircher, Michael Peters, Gábor Virág, and various seminar and conference participants for valuable comments. Xiaoming Cai thanks NSFC Grant 71803143 and the Science Foundations of Shenzhen and Guangdong for support. Ronald Wolthoff gratefully acknowledges financial support from the Connaught Fund and the Social Sciences and Humanities Research Council.

†Peking University HSBC Business School, China. E-mail: xmingcai@gmail.com.

‡Vrije Universiteit Amsterdam and Tinbergen Institute, the Netherlands. E-mail: p.a.gautier@vu.nl.

§University of Toronto, Canada. E-mail: ronald.p.wolthoff@gmail.com.

1 Introduction

Real-life markets display a large degree of heterogeneity in the way in which economic agents meet and trade with each other: for example, in traditional bazaars, meetings between buyers and sellers tend to be bilateral; in real estate markets, multiple buyers may bid on the same house; and in labor markets, a typical vacancy receives a large number of applications but only interviews a subset.[1] Similarly, there is variation over time as the internet has made it easier for agents to meet multiple potential trading partners simultaneously; prominent examples of platforms utilizing this feature include eBay in the product market, Match.com in the dating market, CareerBuilder in the labor market, and Google AdWords in the market for online advertising.

In this paper, we present a search model to study how the way in which market participants meet each other affects equilibrium outcomes, including trading mechanisms and allocations. We focus on an environment in which identical sellers post mechanisms to compete for buyers with ex-ante heterogeneous private valuations.[2]

Economic theory has been mostly silent on the question of how market participants get to meet each other and how this meeting process affects equilibrium outcomes. This silence is most apparent in work that sidesteps a detailed description of the meeting process altogether by assuming a Walrasian equilibrium. Perhaps more surprisingly, the search literature—which aims to analyze trade in the absence of a Walrasian auctioneer—does not provide much more guidance: without much motivation, the vast majority of papers in this literature simply assumes one of two specific meeting technologies: either meetings between agents are one-to-one (bilateral) or they are n-to-1, where n follows a Poisson distribution (urn-ball).[3]

[1] See Geertz (1978) for a characterization of the market interaction at a bazaar, Han and Strange (2014) for empirical evidence on bidding wars in real estate markets, and Wolthoff (2018) and Davis and de la Parra (2017) for evidence on applications and interviews in the labor market.

[2] The fact that buyers know their valuation before visiting a seller distinguishes our work from Lester et al. (2015). See below for a more detailed comparison.

[3] Bilateral meetings can be found in e.g., Albrecht and Jovanovic (1986), Moen (1997), Guerrieri et al. (2010), and Menzio and Shi (2011). Urn-ball meetings are used in e.g., Peters (1997), Burdett et al. (2001), Shimer (2005), Albrecht et al. (2014) and Auster and Gottardi (2017). In addition, some papers in the mechanism design literature explore urn-ball meetings in a finite market, making n binomial rather than Poisson, by allowing for entry of buyers into a monopolistic auction (e.g., Levin and Smith, 1994).

This approach seems restrictive for a number of reasons. First, assuming a particular meeting technology inevitably affects equilibrium outcomes like trading mechanisms or allocations: within our environment, bilateral meetings give rise to a separating equilibrium in which buyers sort themselves by type across different posted prices, while urn-ball meetings lead to a pooling equilibrium in which trade is governed by auctions in which all buyers participate (Eeckhout and Kircher, 2010b; Cai et al., 2017). Second, neither bilateral meetings nor urn-ball meetings are necessarily an adequate description of real-life markets; in many cases (e.g., in the labor market described above), it appears necessary to consider alternatives.

We aim to make progress by presenting a unified framework that allows for a wide class of meeting technologies. The class of technologies that we consider allows for various degrees of meeting externalities: a buyer meeting a seller may make it harder for the seller to meet other buyers (as in the labor market, where an applicant may crowd out other, better applicants if firms cannot screen everyone due to time constraints), or have no impact (as in internet auctions), or even facilitate other meetings. The well-known bilateral and urn-ball meeting technologies as well as other meeting technologies in the literature are all special cases of our general setup.

Our framework thus allows us not only to clarify existing results but also to analyze which of them carry over to our more general setting. For example, we establish that the finding of Albrecht et al. (2012) that reserve prices are driven to sellers' valuation in an environment with competing auctions and ex-ante heterogeneous buyers only holds for some meeting technologies but not for others, which complements a similar finding by Lester et al. (2015) in an environment with ex-post buyer heterogeneity. In addition, we show that some meeting technologies give rise to partial separation rather than complete pooling or complete separation, which have been the focus of the literature so far (see Eeckhout and Kircher, 2010b; Cai et al., 2017).

The equilibrium mechanism that we identify includes both auctions without fees or explicit reserve prices (when meetings are urn-ball) and posted prices (when meetings are bilateral)

as special cases.[4] Varying the degree of search frictions in our model thus changes the optimal mechanism. This feature contrasts with much of the search literature, which assumes that the trading mechanism (e.g., bilateral bargaining) is independent of the frictions. However, changes in mechanisms due to changes in the meeting process are frequently observed in real-life. For example, as soon as eBay provided a platform for sellers and buyers to meet, auctions quickly gained popularity for the sale of second-hand products.[5] Similarly, companies like Upwork (previously oDesk) or Freelancer provide online platforms that facilitate many-to-one meetings between employers from high-income countries and contractors from mainly low-income countries.[6] These platforms have created scope for wage mechanisms other than bilateral bargaining. In particular, contractors apply to posted jobs by submitting a cover letter and a bid indicating the compensation that they demand for the job, after which procurers select one of the applicants.[7] These examples nicely illustrate how a new technology can affect the meeting process and how the market responds by adjusting the price or wage mechanism accordingly.

Our paper makes three main contributions. First, we go further than the existing literature on competing mechanisms and characterize equilibrium for a wide class of meeting technologies. The pioneering work in this area (McAfee, 1993; Peters, 1997; Peters and Severinov, 1997; Albrecht et al., 2014) has generally focused on urn-ball meetings. Eeckhout and Kircher (2010b) were the first to show that the meeting technology matters for posted mechanisms and market segmentation by contrasting two classes of technologies. Relative

[4]When meetings are bilateral, buyers either bid the reserve price or pay the entrance fee and bid 0; both are equivalent to a posted price.

[5]Lucking-Reiley (2000) presents various statistics regarding the growing popularity of online auctions in the late 1990s. Although Einav et al. (2017) argue that the popularity of auctions on eBay has declined relative to posted prices in recent years, their study restricts attention to sellers selling multiple units of the same product (mostly retail items). They acknowledge that auctions remain the trading mechanism of choice for most sellers with a single unit, which is the case that we consider here. Various other platforms (e.g., Catawiki or LiveAuctioneers) continue to exclusively use auctions.

[6]See Agrawal et al. (2015) for a detailed description.

[7]A more exotic example from the dating market is the following case where Amir Pleasants, a 21-year woman from New Jersey invited 150 men on a Tinder date to meet in Union square where she organized a pop-up dating competition where first all guys who were shorter than 5 foot 10 were eliminated and after a number of other rounds she ultimately selected a single winner. See https://www.nytimes.com/2018/08/20/style/tinder-dating-scam-union-square.html

to their work, we contribute by characterizing the equilibrium for a wider class of meeting technologies and fully describing the different forms of market segmentation that may arise. Lester et al. (2015) also provide a full characterization of equilibrium for arbitrary meeting technologies, but in a simpler environment in which all buyers are ex ante identical and learn their type only after meeting a seller, which results in all agents participating in the same (sub)market in equilibrium. In contrast, models with ex ante heterogeneity, as we consider here, yield very different equilibrium outcomes: buyers and sellers must determine with whom they are willing to interact and multiple submarkets may arise.

In particular, we show that when a seller can meet multiple but not all buyers—so low-type buyers may crowd out high-type buyers—partial market segmentation can arise: all high-type buyers and a subset of the low-type buyers form a submarket, while the remaining low-type buyers form a separate submarket. Whether this outcome is obtained depends both on the meeting technology and the distribution of buyer types. We provide a precise characterization of when each type of equilibrium (complete pooling, complete separation, partial pooling, or low-type buyers staying out of the market) arises. Motivated by the above real-world life examples, we then show how the meeting technology affects the optimal selling mechanism and market segmentation. For example, we discuss how the equilibrium changes when sellers can screen more buyers or when they can better distinguish the low-valuation from the high-valuation buyers. We also consider how a change in the high type's valuation affects equilibrium. This exercise could be interpreted as an increase in the dispersion of buyers' valuations.

Cai et al. (2017) apply the tools that are developed in this paper (which we discuss in more detail below) to derive conditions on the meeting technology for which the equilibrium features either perfect separation or perfect pooling of different types of buyers. They find that perfect separation occurs if and only if meetings are bilateral, while perfect pooling arises if and only if the meeting technology is jointly concave.[8] They do not discuss more

[8]They also relate those conditions to other properties of meeting technologies that have been derived in the literature, like invariance Lester et al. (2015) and non-rivalry Eeckhout and Kircher (2010b).

realistic meeting technologies where sellers can meet multiple but not all buyers, which are the main focus of our work here. We show that the partial segmentation that may arise for such technologies nests the perfect separation and perfect pooling in Cai et al. (2017) as special cases.

Second, we make a methodological contribution. In particular, we introduce an alternative representation of meeting technologies which keeps the analysis tractable. This representation is the probability ϕ that a seller meets at least one buyer from a given subset; usually, the relevant subset consists of buyers with a valuation above a certain threshold. This probability depends on two arguments: the total queue length λ that the seller faces as well as the queue of buyers μ belonging to the subset. We show that using ϕ instead of the more standard representation of meeting technologies offers a few important advantages. First, the partial derivatives of ϕ have natural interpretations corresponding to key variables such as a buyer's winning probability and the degree of meeting externalities. Second, expected surplus is linear in ϕ, which makes it straightforward to relate the objective of a planner to properties of ϕ.[9] Finally, the use of ϕ guarantees that the expression for a seller's payoff retains a similar structure as in the seminal work by Myerson (1981), i.e., as the integral of buyers' virtual valuation with respect to the distribution of highest valuations, with the difference that this distribution now also depends on how likely each buyer is to meet a seller which in turn depends on the meeting technology. In other words, the introduction of ϕ adds a lot of generality to the competing mechanism literature at relatively low cost.

Finally, our efficiency result contributes to the literature on directed search. In particular, it extends the result by Albrecht et al. (2014) that all agents earn their marginal contribution to surplus in the special case in which meetings are urn-ball and sellers post regular auctions.[10] In that environment, there are no meeting externalities, so a buyer contributes to surplus only if he has the highest valuation among all buyers meeting a seller. The general case is more complicated because now a buyer can also impose positive or negative meeting externalities on

[9] Cai et al. (2017) exploit this feature in their work.

[10] Although we assume a fixed number of sellers to simplify exposition, our results carry over to an environment with free entry of sellers, as in Albrecht et al. (2014), in a straightforward manner.

meetings between the seller and other buyers, which should be reflected in the equilibrium payoffs. We show that an appropriate reserve price or meeting fee/subsidy is both profit maximizing and socially efficient.[11] As a result, all agents continue to receive their marginal contribution to surplus and efficiency survives.

After describing the environment and the alternative representation of the meeting technology in detail in Section 2, we start our analysis in Section 3 by solving the problem of a social planner. It provides a full characterization of the planner's solution and discusses comparative statics. Section 4 shows how the planner's solution can be decentralized. In Section 5, we show how our results can be generalized to N buyer types.

2 Model

2.1 Environment

Agents and Preferences. A static economy is populated by risk-neutral buyers and sellers. Each seller possesses a single unit of an indivisible good, for which each buyer has unit demand. All sellers have the same valuation for their good, which we normalize to zero. Buyers are heterogeneous in their valuation x, which takes one of two different values, satisfying $0 < x_1 < x_2$. We will generally normalize x_1 to 1. The measure of sellers is 1; the measure of buyers with value x_k is B_k for $k \in \{1, 2\}$. Buyers' valuations are private information and the market is anonymous in the sense that buyers and sellers cannot condition their strategies on the identities of their counterparties.

Mechanisms. In the first stage, each seller posts and commits to a direct anonymous mechanism to attract buyers. The mechanism specifies, for each buyer i, a probability of trade and an expected payment as a function of: (i) the total number n of buyers that successfully meet with the seller; (ii) the valuation v_i that buyer i reports; and (iii) the

[11] The reserve price or fee can vary across sellers in equilibrium. This is a key difference with Lester et al. (2015), where the fee is the same for all sellers as it only depends on exogenous parameters.

valuations v_{-i} reported by the $n-1$ other buyers.[12]

Search. After observing all mechanisms, each buyer chooses the one at which he wishes to attempt to match. To capture the idea that coordination is not feasible in a large market, we follow the literature (e.g., Montgomery, 1991; Burdett et al., 2001; Shimer, 2005) and restrict buyers to symmetric strategies. We refer to all buyers and sellers choosing a particular mechanism as a *submarket*.

Meeting Technology. Consider a submarket with a measure b of buyers and a measure s of sellers. The meetings within the submarket are frictional and governed by a *meeting technology*, which we model analogous to Eeckhout and Kircher (2010b). In particular, the meeting technology treats all buyers (sellers) symmetrically, i.e., independent of their identity or type.[13] A buyer can meet at most one seller, while a seller may meet multiple buyers. Define $\lambda = b/s$ as the *queue length* in this submarket.[14] The probability of a seller meeting n buyers, $n = 0, 1, 2, \ldots$, is then given by $P_n(\lambda)$, which is assumed to be continuously differentiable. Because each buyer can meet at most one seller, $\sum_{n=1}^{\infty} nP_n(\lambda) \leq \lambda$. By an accounting identity, the probability for a buyer to be part of an n-to-1 meeting is $Q_n(\lambda) \equiv nP_n(\lambda)/\lambda$ with $n \geq 1$; the probability that a buyer fails to meet any seller is then $Q_0(\lambda) \equiv 1 - \sum_{n=1}^{\infty} Q_n(\lambda)$. Finally, we define $m(\lambda) \equiv 1 - P_0(\lambda)$ as the probability that a seller meets at least one buyer, which is a necessary and sufficient condition for the seller to trade in our environment.[15] To avoid potential pathological cases, we require that $m(\lambda)$ is not linear in any interval.[16]

[12] In line with most of the literature, we abstract from mechanisms that condition on other mechanisms present in the market. See Epstein and Peters (1999) and Peters (2001) for a detailed discussion.

[13] Meetings being type independent is a natural benchmark since it creates a distinction between meetings and matches: matches can be type dependent if the mechanism selects the buyer with the highest valuation. Of course, firms can also increase the likelihood of meeting a particular type by adjusting the mechanisms they post and offering that type a better deal.

[14] This assumes, for simplicity, that a positive measure of buyers and sellers visit the submarket. If this is not the case, we can use Radon-Nykodym derivatives to define queue lengths.

[15] It is straightforward to allow buyers to observe only a fraction of the sellers. If the fraction of sellers that a buyer observes is type independent, this will not change our results.

[16] In other words, $m'(\lambda)$, the marginal effect of a higher queue length on sellers' matching probability, cannot stay constant when we increase λ slightly. This requirement is always satisfied when $m(\lambda)$ is an analytic function (note that $m(\lambda)$ cannot be always linear because it must be between 0 and 1).

2.2 Alternative Representation of Meetings

In most of our analysis, we will remain agnostic about the exact nature of the meeting technology and just impose the minimal structure we need to prove our results. To deal with this level of generality, we first present a transformation of the meeting technology that greatly simplifies the analysis. In particular, we introduce a new function $\phi(\mu, \lambda)$ with $0 \leq \mu \leq \lambda$, defined as

$$\phi(\mu, \lambda) = 1 - \sum_{n=0}^{\infty} P_n(\lambda) \left(1 - \frac{\mu}{\lambda}\right)^n. \tag{1}$$

To understand this function, consider a submarket in which sellers face a queue length λ. Suppose that a fraction μ/λ of the buyers in the submarket has the high value x_2. Since the meeting technology treats different buyers symmetrically, $\phi(\mu, \lambda)$ then represents the probability that a seller meets at least one high-value buyer. Naturally, $\phi(\lambda, \lambda) = m(\lambda)$.

The function $\phi(\mu, \lambda)$ allows us to study competing mechanisms with general meeting technologies in a way that is both more tractable and more intuitive than with $P_n(\lambda)$, $n = 0, 1, \ldots$. Appendix A.1 establishes that no information is lost by considering $\phi(\mu, \lambda)$ instead of $P_n(\lambda)$, since we can always recover one from the other.

To develop intuition for $\phi(\mu, \lambda)$, suppose that $\Delta\lambda$ more buyers visit this submarket, then the probability that the seller meets at least one *incumbent* high-value buyer becomes $\phi(\mu, \lambda + \Delta\lambda)$, where μ is the measure of the incumbent high-value buyers. Therefore, $\phi_\lambda(\mu, \lambda) \equiv \partial\phi(\mu, \lambda)/\partial\lambda$ measures the effect of the new entrants on the meeting probabilities between sellers and incumbent high-value buyers: $\phi_\lambda(\mu, \lambda) < 0$ (resp. > 0) represents negative (resp. positive) meeting externalities. In the special case of $\phi_\lambda(\mu, \lambda) = 0$, there are no meeting externalities among buyers.

For future reference, note that

$$\phi_\mu(\mu, \lambda) \equiv \frac{\partial\phi(\mu, \lambda)}{\partial\mu} = \sum_{n=1}^{\infty} Q_n(\lambda) \left(1 - \frac{\mu}{\lambda}\right)^{n-1}. \tag{2}$$

That is, $\phi_\mu(\mu, \lambda)$ is the probability for a buyer to be part of a meeting in which all other buyers (if any) have low valuations. In this case, a high-type buyer increases social surplus *directly*, since the good would have been allocated to a low-type buyer in his absence. In a second-price auction, this is also the probability that a high-type buyer wins the auction with strictly positive payoff, which we define to be the *winning probability* of high-type buyers.[17] Since for each n, $(1 - \mu/\lambda)^{n-1}$ is decreasing in μ, $\phi_\mu(\mu, \lambda)$ is then also decreasing in μ, implying that $\phi(\mu, \lambda)$ is concave in μ, which holds strictly ($\phi_{\mu\mu}(\mu, \lambda) < 0$) if and only if $P_0(\lambda) + P_1(\lambda) < 1$.[18] Two special cases of equation (2) are worth mentioning: i) $\phi_\mu(0, \lambda) = 1 - Q_0(\lambda)$, i.e., the probability that a buyer meets a seller, and ii) $\phi_\mu(\lambda, \lambda) = Q_1(\lambda)$, i.e., the probability that a buyer meets a seller without other buyers.

3 Social Planner

3.1 Surplus and Planner's Problem

Surplus. We start our analysis with the following lemma which derives total surplus and agents' marginal contribution to this surplus in a submarket with queue (μ, λ), where μ is the queue length of buyers with value x_2 and λ is the total queue length.

Lemma 1. *Consider a submarket with a measure 1 of sellers and a queue (μ, λ) of buyers. Total surplus in the submarket then equals*

$$S(\mu, \lambda) = m(\lambda) + (x_2 - 1)\,\phi(\mu, \lambda) \tag{3}$$

[17] Because buyers types are discrete, buyers' winning and trading probability are different: a buyer may compete with another buyer with the same value. But as we will see later, this difference is not important for our analysis. The use of this winning probability is a canonical technique developed by McAfee (1993) and Peters and Severinov (1997). They show that buyers' winning probability must be equal at competing sellers. Our function $\phi(\mu, \lambda)$ incorporates their approach by its first partial derivative $\phi_\mu(\mu, \lambda)$ and does more because its second partial derivative $\phi_\lambda(\mu, \lambda)$ represents meeting externalities. Moreover, the function $\phi(\mu, \lambda)$ itself is intimately linked with surplus. See Lemma 1 and also Lemma 3 for the formal statements.

[18] For each $n \geq 0$, $-(1 - \mu/\lambda)^n$ is increasing and concave in μ, and it is strictly concave in μ if and only if $n \geq 2$. Therefore, $\phi(\mu, \lambda)$ is strictly concave in μ if and only if there exists at least one $n \geq 2$ such that $P_n(\lambda) > 0$.

The marginal contribution to surplus of low-type and high-type buyers are, respectively,

$$T_1(\mu,\lambda) = m'(\lambda) + (x_2 - 1)\,\phi_\lambda(\mu,\lambda) \tag{4}$$

$$T_2(\mu,\lambda) = m'(\lambda) + (x_2 - 1)\left(\phi_\mu(\mu,\lambda) + \phi_\lambda(\mu,\lambda)\right). \tag{5}$$

A seller's marginal contribution to surplus equals

$$R(\mu,\lambda) = m(\lambda) - \lambda m'(\lambda) + (x_2 - 1)\left(\phi(\mu,\lambda) - \mu\phi_\mu(\mu,\lambda) - \lambda\phi_\lambda(\mu,\lambda)\right). \tag{6}$$

Proof. See Appendix A.2. □

The first term in equation (3) accounts for the fact that a surplus of (at least) 1 is generated whenever a seller meets at least one buyer. The second term captures that an additional surplus of $x_2 - 1$ is realized when a seller meets at least one high-type buyer.

To understand (4), note that $T_1(\mu,\lambda) = S_\lambda(\mu,\lambda)$ since adding a low-type buyer to the submarket increases λ but has no effect on μ. The first term of (4) reflects the effect of the extra buyer on the number of matches, while the second term represents the externalities that he may impose on meetings between sellers and high-type buyers. Since $m'(\lambda) = \phi_\mu(\lambda,\lambda) + \phi_\lambda(\lambda,\lambda)$, equation (4) can also be written as

$$T_1(\mu,\lambda) = \phi_\mu(\lambda,\lambda) + \phi_\lambda(\lambda,\lambda) + (x_2 - 1)\,\phi_\lambda(\mu,\lambda),$$

where the first term describes the buyer's direct contribution to surplus which arises when there are no other buyers, as discussed below equation (2). The second term and third term represent the externalities that the buyer may impose on sellers' meetings with, respectively, other low-type and high-type buyers.

To understand (5), note that $T_2(\mu,\lambda) = S_\mu(\mu,\lambda) + S_\lambda(\mu,\lambda)$ since adding an additional high-type buyer to the submarket increases both μ and λ. Therefore, $T_2(\mu,\lambda) = T_1(\mu,\lambda) + (x_2 - 1)\,\phi_\mu(\mu,\lambda)$. That is, the additional high-type buyer creates the same meeting exter-

nalities as an extra low-type buyer, but creates additional surplus when there are no other buyers or only low-type buyers, which happens with probability $\phi_\mu(\mu, \lambda)$.

Finally, to understand equation (6), define $z = \mu/\lambda$ to be the fraction of high-type buyers in the queue. If we add λ more buyers to the submarket while keeping z fixed, then adding one more seller increases surplus by $S(\lambda z, \lambda)$. Therefore, $R(\lambda z, \lambda) = S(\lambda z, \lambda) - \lambda \frac{\partial S(\lambda z, \lambda)}{\partial \lambda}$, or, alternatively $R(\mu, \lambda) = S(\mu, \lambda) - \mu T_2(\mu, \lambda) - (\lambda - \mu) T_1(\mu, \lambda)$.

Planner's Problem. One can think of the planner's problem as a three-step optimization problem: first, the planner chooses the number of submarkets to open; second, he determines the allocation of buyers and sellers to the different submarkets; third, he decides on the allocation of the good after meetings have taken place. The third step is trivial: at each seller, the good is always allocated to the buyer with the highest valuation. The first two steps will depend on the meeting technology and the distribution of valuations. Suppose that the planner creates L submarkets with positive seller measures $\alpha^1, \ldots, \alpha^L$, respectively, and potentially an additional submarket with no sellers but only buyers. Of course, this additional submarket generates no surplus but could be useful for reducing meeting externalities. The queue in submarket $\ell = 1, \ldots, L$ is (μ^i, λ^i). The planner's problem is thus

$$\mathcal{S}^*(B_1, B_2) = \sup_{L \geq 1} \sup_{\{(\alpha^\ell, \mu^\ell, \lambda^\ell) \,|\, \ell = 1, \ldots, L\}} \sum_{\ell=1}^{L} \alpha^\ell S(\mu^\ell, \lambda^\ell) \tag{7}$$

subject to the accounting constraints $\sum_{\ell=1}^{L} \alpha^\ell = 1$, $\sum_{\ell=1}^{L} \alpha^\ell \mu^\ell \leq B_2$, and $\sum_{\ell=1}^{L} \alpha^\ell (\lambda^\ell - \mu^\ell) \leq B_1$.[19]

3.2 Preliminary Results

In this section, we present a few preliminary results regarding the planner's solution.

Number of Submarkets. It is not clear a priori that there is an upper bound on the number of submarkets L for any endowment of buyers (B_1, B_2). However, the following

[19] The inequalities reflect that the planner may require some buyers to be inactive and not visit any seller.

Proposition shows that such an upper bound exists. To state the result, we define an *idle* submarket as a market that either contains only buyers or only sellers (as opposed to an *active* submarket in which both buyers and sellers are present).[20]

Proposition 1. *When there are two buyer types, the planner's problem can be solved by opening at most three submarkets. At most one of those submarkets can be idle.*

Proof. See Appendix A.3. □

Proposition 1 serves two goals. First, it is a technical result on existence in the sense that it establishes that the supremum of surplus over all possible allocations can indeed be reached as a maximum. At this level of generality, the planner's solution is not necessarily unique: multiple allocations that generate the same surplus may exist. However, these cases are rather special. Below, we will introduce two mild conditions that hold for all commonly used meeting technologies and that guarantee uniqueness.

Second, Proposition 1 limits the complexity of the planner's problem by bounding the number of submarkets. By equation (7), total surplus is a convex combination of the surpluses generated by individual submarkets. The planner chooses the number of submarkets to find the maximum value that such convex combinations can reach, which simply corresponds to finding the least concave majorant of the surplus function S in equation (3), i.e., the smallest concave function that is greater than S. As a result, the Fenchel-Bunt Theorem provides an upper bound for the number of submarkets needed to solve the planner's problem.[21]

To understand this result, consider first the case of homogeneous buyers, where the planner's problem can be solved by opening at most two submarkets. Suppose that the planner opens three submarkets with queue lengths $\lambda_1 < \lambda_2 < \lambda_3$. Then, the marginal contribution to surplus by both sellers and buyers must be the same across the three submarkets. Since

[20] The planner will of course never simultaneously choose an idle market for buyers and one for sellers.

[21] The classical Caratheodory theory states that any point in the convex hull of a set $A \subset \mathbb{R}^n$ (i.e., the smallest convex set containing A) can be represented as a convex combination of $n+1$ points of A. The Fenchel-Bunt Theorem states that if the set A is connected, then for the above construction we only need n points instead of $n+1$. Since the graph of $S : \mathbb{R}^2 \to \mathbb{R}$ is a connected subset in $\mathbb{R}^3$, the Fenchel-Bunt Theorem implies that we only need three points to construct the planner's solution since it belongs to the convex hull of the graph.

the meeting technology exhibits constant returns to scale, total surplus is simply the sum of buyers' and sellers' marginal contributions to surplus multiplied by their respective measure (Euler's theorem). This implies that the planner can reallocate the buyers and the sellers in the submarket with queue length λ_2 to the other two submarkets such that the queue lengths in those two submarkets remain λ_1 and λ_3, respectively, and total surplus stays the same.

When there are two types of buyers, the above logic still works if, for example, all three submarkets have the same queue length but different fractions of high-type buyers. However, if the submarket with an intermediate fraction of high-type buyers has the largest queue length, then the above reallocation is not possible and the planner needs three submarkets. Opening four submarkets will make one submarket redundant because the planner can then always reallocate buyers and sellers in one of the two submarkets with an intermediate fraction of high-type buyers to the other three submarkets such that total surplus remains constant.

Concavity. If the surplus function $S(\mu, \lambda)$ is jointly concave in (μ, λ) then its least concave majorant is of course S itself. In this case, merging any two submarkets always increases total surplus and the planner's unique solution is simply to pool all buyers and sellers into a single submarket (see Cai et al., 2017). However, as we show below, joint concavity of $S(\mu, \lambda)$ is often violated. In these cases, the planner needs to solve a non-concave optimization problem, which is notoriously difficult. As mentioned, we make progress below by formulating two weak restrictions on the meeting technology. Under those restrictions, the first-order conditions are both necessary and sufficient, and a simple algorithm provides their solution.

Even if concavity of $S(\mu, \lambda)$ fails globally, it still needs to hold locally in any submarket (μ, λ) satisfying $0 < \mu < \lambda$. Otherwise, by definition, we can break the submarket into two and reallocate sufficiently small measures of buyers $\Delta\mu$ and $\Delta\lambda$ to increase total surplus, i.e.,

$$\frac{1}{2}S(\mu - \Delta\mu, \lambda - \Delta\lambda) + \frac{1}{2}S(\mu + \Delta\mu, \lambda + \Delta\lambda) > S(\mu, \lambda).$$

The Hessian matrix of the surplus function $S(\mu, \lambda)$ must therefore be negative semi-definite at the point (μ, λ). This normally requires two inequalities to hold since the Hessian is a

2×2 matrix. However, the surplus function $S(\mu, \lambda)$ is linear in $\phi(\mu, \lambda)$, which in turn is always concave in μ. The only remaining condition therefore is that the determinant of the Hessian is positive, which implies the following result.

Lemma 2. *A submarket (μ, λ) with $0 < \mu < \lambda$ can be part of the planner's solution only if i) $\phi_{\mu\mu}(\mu, \lambda) < 0$ (or equivalently $P_0(\lambda) + P_1(\lambda) < 1$) and ii) the Hessian matrix of the surplus function (3) is negative semi-definite at (μ, λ), which is equivalent to:*

$$(x_2 - 1)\left(\phi_{\lambda\lambda}(\mu, \lambda) - \frac{\phi_{\mu\lambda}(\mu, \lambda)^2}{\phi_{\mu\mu}(\mu, \lambda)}\right) \leq -m''(\lambda) \tag{8}$$

Proof. See Appendix A.4. □

To understand equation (8), recall that $\phi(\mu, \lambda)$ is the probability that high-type buyers generate additional surplus $(x_2 - 1)$. If ϕ is concave in (μ, λ), or equivalently $\phi_{\lambda\lambda} - \phi_{\mu\lambda}^2/\phi_{\mu\mu} \leq 0$, then merging different submarkets will increase the probability that high-type buyers generate additional surplus (see Cai et al., 2017). However, when ϕ is not concave in (μ, λ), then the left-hand side of (8) is positive and represents the *maximal marginal gain* that can be achieved by increased separation of the two types of buyers.[22] This gain is small if x_2 is close to 1, i.e., the buyer types are almost identical. Since $m''(\lambda)$ is negative for all common meeting technologies (see Assumption 1 below), we can interpret the right-hand side as the marginal gain of pooling buyers (in terms of increasing the aggregate number of matches). When condition (8) fails, the planner prefers to break up a submarket since that would increase total surplus.[23]

[22] To see this, suppose that we divide a submarket (μ, λ) into two submarkets $(\mu - \Delta\mu, \lambda - \Delta\lambda)$ and $(\mu + \Delta\mu, \lambda + \Delta\lambda)$ with $\Delta\lambda > 0$ and $\Delta\mu$ of indeterminate sign. Then, the reduction in the number of matches is $m(\lambda) - \frac{1}{2}m(\lambda + \Delta\lambda) - \frac{1}{2}m(\lambda - \Delta\lambda)$, which equals $-m''(\lambda)\Delta\lambda^2 > 0$. The marginal gain from a higher probability of matching with high-type buyers is $(x_2 - 1)\left[\frac{1}{2}\phi(\mu - \Delta\mu, \lambda - \Delta\lambda) + \frac{1}{2}\phi(\mu + \Delta\mu, \lambda + \Delta\lambda) - \phi(\mu, \lambda)\right] = (x_2 - 1)\left(\phi_{\mu\mu}\Delta\mu^2 + 2\phi_{\mu\lambda}\Delta\mu\Delta\lambda + \phi_{\lambda\lambda}\Delta\lambda^2\right)/2$. The gain is maximized when $\Delta\mu = -\Delta\lambda\phi_{\mu\lambda}/\phi_{\mu\mu}$, and the maximum gain is $(x_2 - 1)(\phi_{\lambda\lambda} - \phi_{\mu\lambda}^2/\phi_{\mu\mu})$.

[23] We thank an anonymous referee for suggesting the above interpretation of equation (8).

3.3 Some Mild Restrictions on the Meeting Technology

Assumptions. We now present two weak assumptions regarding $\phi(\mu, \lambda)$ which provide the minimal structure we need to prove our results. To introduce the assumptions, we apply a change of notation and define $z = \mu/\lambda$ as the fraction of high-type buyers in the queue.

Assumption 1. *$\phi(\lambda z, \lambda)$ is concave in λ for any given z. Furthermore, $\lim_{\lambda \to 0} m'(\lambda) = 1$ and $\lim_{\lambda \to \infty} m(\lambda) - \lambda m'(\lambda) = 1$.*

The first part of Assumption 1 states that if we hold the fraction of high-type buyers constant, the marginal effect of an extra buyer on the seller's probability of meeting at least one high-type is decreasing in the total queue length. This assumption has a number of important implications. First, since $\phi(\lambda z, \lambda)$ is always between 0 and 1, it implies that $\phi(\lambda z, \lambda)$ is increasing in λ for given z. Second, because $m(\lambda) \equiv \phi(\lambda, \lambda)$, it implies that $m(\lambda)$ is concave; the concavity is strict, i.e. $m(\lambda) > \lambda m'(\lambda)$, since we require that $m(\lambda)$ is not linear in any interval. Third, Assumption 1 implies that $\phi(\lambda z, \lambda) \geq \lambda \frac{\partial \phi(\lambda z, \lambda)}{\partial \lambda} = \mu \phi_\mu(\mu, \lambda) + \lambda \phi_\lambda(\mu, \lambda)$ and thus that sellers' marginal contribution to surplus R in equation (6) is always strictly positive, which means that no sellers should be idle in the planner's solution. Fourth, the assumption implies that $R(\lambda z, \lambda)$ is strictly increasing in λ for a given z, since $\frac{\partial}{\partial \lambda} R(\lambda z, \lambda) = -\lambda m''(\lambda) + (x_2 - 1) \frac{\partial^2}{\partial \lambda^2} \phi(\lambda z, \lambda) > 0$. Fifth, Assumption 1 also implies that $T_2(\lambda z, \lambda)$, the marginal contribution to surplus by high-type buyers given by equation (5), is strictly positive, because $m'(\lambda) > 0$ and $\phi_\mu(\lambda z, \lambda) + \phi_\lambda(\lambda z, \lambda) \geq z \phi_\mu(\lambda z, \lambda) + \phi_\lambda(\lambda z, \lambda) = \frac{\partial \phi(\lambda z, \lambda)}{\partial \lambda} \geq 0$.

The second part of Assumption 1 is more a normalization than an assumption. It implies that in a submarket with only buyers with valuation x_k, the marginal contribution of these buyers is x_k when $\lambda \to 0$, while for sellers it is x_k when $\lambda \to \infty$.

Level Curves. If the planner creates multiple active submarkets, then sellers' contribution to surplus must be equal in all those submarkets. In other words, these submarkets must lie

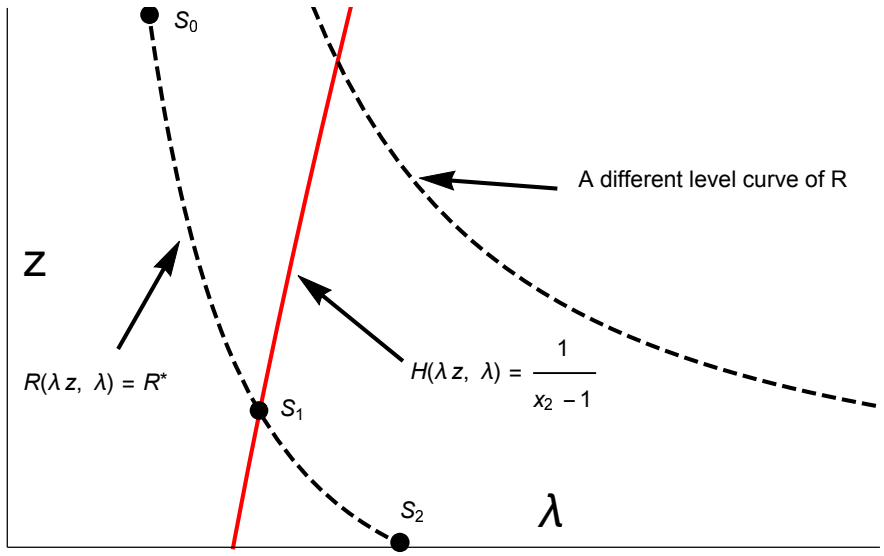

Figure 1: Illustration of Assumption 2 (Single Crossing)

on some level curve $R(\lambda z, \lambda) = R^*$. Figure 1 illustrates two of such level curves.[24]

Similarly, buyers' marginal contribution to surplus must be equal in all submarkets that they visit. A key question is therefore how a buyer's marginal contribution $T_1(\lambda z, \lambda)$ or $T_2(\lambda z, \lambda)$ varies along level curves of $R(\lambda z, \lambda)$. Lemma 5 in Appendix A.6 establishes that the answer depends on the *sign* of the determinant of the Hessian of the surplus function. In Figure 1, the red solid curve is where the determinant is zero; the area to its left is where the determinant is positive and the area to its right is where the determinant is negative. As a result, high-type buyers' marginal contribution $T_2(\lambda z, \lambda)$ is non-monotonic along the level curve $R(\lambda z, \lambda) = R^*$. It increases as we move from point S_0 to S_1 and decreases as we move from S_1 to S_2. In other words, it reaches its maximum at point S_1. Intuitively, when z is large, a high-type buyer is unlikely to be the only high type that sellers meet even though the queue is short. When z is small, the queue length is large. The large number of low-type buyers reduces the likelihood that a seller meets a high-type buyer and hence reduces the marginal contribution to surplus by high-type buyers. The reverse holds for $T_1(\lambda z, \lambda)$.

[24]For illustrative purposes, the level curves in Figure 1 are downward sloping in the λ-z plane. This is accurate under the weak additional assumption that $\phi_\mu(\lambda z, \lambda)$ is strictly decreasing in λ for $0 < z < 1$, i.e., Assumption 3 introduced later. Although it is satisfied by all common meeting technologies, we abstain from making this assumption here because it is not required for our main results.

Inspection of Figure 1 shows that for $T_1(\lambda z, \lambda)$ and $T_2(\lambda z, \lambda)$ to have a unique extremum, the level curves of $R(\lambda z, \lambda)$ must intersect with the red curve only once and from the left. For this to hold in general, we require one additional assumption on the meeting technology, which is weak in the sense that it is trivially satisfied for the two most common classes of meeting technologies: bilateral and invariant (see the discussion of Figure 2 below).

Assumption 2 (Single Crossing)**.** *For λ with $P_0(\lambda)+P_1(\lambda) < 1$ (or equivalently $\phi_{\mu\mu}(\mu, \lambda) < 0$), define*

$$H(\mu, \lambda) \equiv \frac{1}{-m''(\lambda)} \left(\phi_{\lambda\lambda}(\mu, \lambda) - \frac{\phi_{\mu\lambda}(\mu, \lambda)^2}{\phi_{\mu\mu}(\mu, \lambda)} \right). \tag{9}$$

At any point (z, λ) where $H(\lambda z, \lambda) > 0$, we have i) $\partial H(\lambda z, \lambda)/\partial\lambda > 0$ and ii)

$$-\frac{\partial \phi_\mu(\lambda z, \lambda)/\partial z}{\partial \phi_\mu(\lambda z, \lambda)/\partial \lambda} < -\frac{\partial H(\lambda z, \lambda)/\partial z}{\partial H(\lambda z, \lambda)/\partial \lambda}. \tag{10}$$

It is worth highlighting that—like Assumption 1—Assumption 2 concerns the meeting technology only. As discussed after (8), the function H represents the relative magnitude of marginal gain from separation and that from pooling. The first part of Assumption 2 requires that whenever there is marginal gain from separation ($H(\lambda z, \lambda) > 0$), then a longer queue (while fixing the fraction of high-type buyers) will make the relative gain from separation even higher. The left-hand side and right-hand side of (10) denote the slope of the level curves of $\phi_\mu(\lambda z, \lambda)$ and $H(\lambda z, \lambda)$, respectively (in the λ-z plane). Thus, the second part of Assumption 2 states that any level curve of $\phi_\mu(\lambda z, \lambda)$ crosses any positive level curve of $H(\lambda z, \lambda)$ at most once and from the left.

Recall that $\phi_\mu(\lambda z, \lambda)$ is the probability that a high-type buyer contributes to the surplus directly. One can increase λ and adjust z accordingly to make ϕ_μ constant, so that the difference between marginal contributions to surplus by the two types of buyers also remains constant. Assumption 2 states that in doing so, the relative gain from separation will increase. This single-crossing condition implies that each level curve of $R(\lambda z, \lambda)$ can be

divided into two segments: the determinant of the Hessian is positive when z is large and negative when z is small, as in Figure 1. At point S_1, the determinant is exactly zero, so that the maximum marginal gain from separation exactly cancels out the corresponding marginal loss in matching probability. The maximal marginal gain is obtained by equalizing sellers' marginal contribution to surplus in the two submarkets. Therefore, at point S_1 along the level curve of $R(\lambda z, \lambda)$, the marginal changes in both the low-type and the high-type buyers' marginal contributions to surplus are exactly zero. In other words, the three level curves of $R(\lambda z, \lambda)$, $T_1(\lambda z, \lambda)$ and $T_2(\lambda z, \lambda)$ are tangent to each other at point S_1. Note that since $T_1(\lambda z, \lambda)$ and $T_2(\lambda z, \lambda)$ have the same slope along the level curve, this common slope must be given by the left-hand side of (10), the slope of $T_2(\lambda z, \lambda) - T_1(\lambda z, \lambda)$. These claims are made precise in Lemma 6 in Appendix A.6.

Examples of Meeting Technologies. We now describe three classes of meeting technologies that satisfy Assumption 1 and 2. The first two classes have been the focus of earlier literature (see e.g., Cai et al., 2017) and are presented for completeness. The third class is novel and bridges the distance between the first two classes. Figure 2 presents a Venn diagram illustrating this idea.

1. *Bilateral.* With a bilateral technology, each seller meets at most one buyer, i.e., $P_0(\lambda) + P_1(\lambda) = 1$ with $P_1(\lambda)$ strictly concave, such that $\phi(\mu, \lambda) = m(\lambda)\mu/\lambda$. A well-known example is telephone-line matching, where $P_1(\lambda) = \lambda/(1+\lambda)$, i.e., $\phi(\mu, \lambda) = \mu/(1+\lambda)$.

2. *Joint Concavity.* Cai et al. (2017) define joint concavity as the $\phi(\mu, \lambda)$ of a meeting technology being concave in (μ, λ), i.e., $\phi_{\mu\mu}\phi_{\lambda\lambda} - \phi_{\mu\lambda}^2 \geq 0$, which they show implies either positive ($\phi_\lambda > 0$) or no meeting externalities ($\phi_\lambda = 0$). The absence of meeting externalities ($\phi_\lambda(\mu, \lambda) = 0$ for all μ and λ) is known as *invariance*, which is a special case of joint concavity since $\phi_\lambda = 0$ implies $\phi_{\mu\lambda} = 0$ and $\phi_{\lambda\lambda} = 0$.[25] Perhaps the best-known example of a technology satisfying invariance—and thus joint concavity—is urn-ball (Butters, 1977). A second example is the geometric technology of Lester

[25] Lester et al. (2015) first introduced invariant meeting technologies in terms of $P_n(\lambda)$. Cai et al. (2017) show that their definition is equivalent to $\phi_\lambda(\mu, \lambda) = 0$.

et al. (2015), where $P_n(\lambda)$ is a geometric probability distribution with mean λ, i.e., $P_n(\lambda) = \left(\frac{1}{1+\lambda}\right)\left(\frac{\lambda}{1+\lambda}\right)^n$ or $\phi(\mu, \lambda) = \mu/(1+\mu)$. A micro-foundation for this meeting technology is that finitely many buyers and sellers are uniformly and independently positioned on a circle and buyers walk clockwise to the nearest seller; then let the number of buyers and sellers approach infinity while keeping their ratio constant at λ.

3. *Stochastic Capacity.* This class is similar to the geometric technology, except sellers may face time or capacity constraints, preventing them from meeting all buyers that try to visit them. The maximum number of buyers that a seller can meet follows a geometric distribution with parameter σ and support $\{1, 2, 3, \dots\}$. The number of meetings taking place is therefore the minimum of two geometric variables, the number of buyers that try to meet/contact the seller and the seller's capacity. This technology reduces to a bilateral one (telephone-line) when $\sigma = 0$ and an invariant one (geometric) when $\sigma = 1$.[26] We show in Appendix B.2 that $\phi(\mu, \lambda) = \mu/(1 + \sigma\mu + (1-\sigma)\lambda)$.

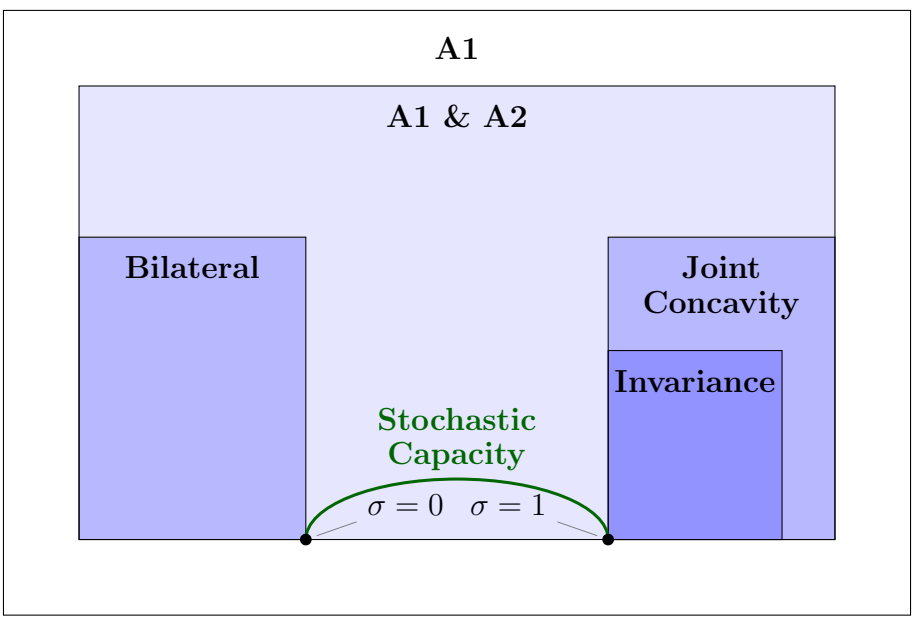

Figure 2: Venn Diagram of Meeting Technologies

Although it is possible to construct exotic meeting technologies that violate Assumption 1,

[26] For a similar technology with Poisson applications, see Wolthoff (2018).

we view it as a basic and reasonable requirement. That is, this assumption defines the universe of meeting technologies that we consider (see Figure 2). In Appendix A.5, we show that the above examples all satisfy Assumption 1. Assumption 2 becomes void for bilateral technologies (since $P_0(\lambda) + P_1(\lambda) = 1$ for all λ) and jointly concave technologies (since $H(\lambda z, \lambda) \leq 0$ for all λ and z), so it is trivially satisfied in those cases.[27] We show in Appendix A.5 that Assumption 2 is satisfied by the stochastic capacity technology as well, but also provide an example of a technology that violates it.

3.4 Full Characterization

Two Submarkets. We can now further tighten the bound on the number of submarkets. We illustrate the argument in Figure 1. Suppose that at the social optimum, the marginal contribution to surplus of sellers is R^* and the black dashed level curve $R(\lambda z, \lambda) = R^*$ intersects the red curve $H(\lambda z, \lambda) = 1/(x_2 - 1)$ at point $S_1 = (\lambda^* z^*, \lambda^*)$. The first part of Assumption 2 ensures that the second-order condition (8) is satisfied left of the red line and violated right of the red line. In other words, submarkets with $0 < z < z^*$ (i.e., points on the S_1S_2 trajectory) cannot be part of the planner's solution. The only feasible submarket on this side is therefore the corner S_2 where $z^* = 0$.

In contrast, the second-order condition is satisfied in submarkets with $z \geq z^*$, i.e., points on the S_0S_1 trajectory. However, by Lemma 5, $T_2(\lambda z, \lambda)$ is strictly decreasing in z along this trajectory. Since the marginal contribution of high-type buyers must be the same among all submarkets containing such buyers, there can therefore only exist one submarket with $z \geq z^*$. To sum up, there exist at most two submarkets in the social optimum: one with $z \geq z^*$ and one with $z = 0$.

The above observation greatly simplifies the analysis because it implies that there are only three possible solutions: (i) complete pooling, i.e., all agents are in one market; (ii) complete separation, i.e., there is one submarket for all high-type buyers and one (possibly idle) for

[27] By equation (8), $H(\lambda z, \lambda) > 0$ if and only if $\phi_{\mu\mu}\phi_{\lambda\lambda} - \phi_{\mu\lambda}^2 < 0$ at the point (z, λ). In this paper, we consider the more realistic case where ϕ is not always or never concave in (μ, λ).

all low-type buyers; (iii) mixing, i.e., there is one submarket that contains all high-type and some low-type buyers, and one (possibly idle) submarket with the remaining low-type buyers. From Cai et al. (2017), we know that jointly concave technologies imply complete pooling, while bilateral technologies imply complete separation. The third possibility, which spans the range between these extremes, is new. It allows the planner to take advantage of multilateral meetings to screen ex post by pooling high types with some low types, while reducing the degree of crowding out by separating other low types.

The optimal extent of separation depends on the magnitude of the meeting externalities, the measures of high and low types, and the relative valuations, i.e., x_2/x_1, which is simply x_2 since we normalize $x_1 = 1$. To solve for it, assume, without loss of generality, that the planner opens two submarkets, one with a high average valuation containing all high-type buyers, and one with a low average valuation without high-type buyers. Two decisions then remain: i) how to allocate low-type buyers and ii) how to allocate sellers. We solve these decisions sequentially.

Allocation of Sellers. Suppose the planner assigns a measure b_1 of low-type buyers to the submarket with the high average valuation. The optimal allocation of sellers, denoted by $\alpha^*(b_1)$, then solves

$$\overline{S}(b_1) = \max_{\alpha} \quad \alpha S\left(\frac{B_2}{\alpha}, \frac{B_2 + b_1}{\alpha}\right) + (1-\alpha) S\left(0, \frac{B_1 - b_1}{1-\alpha}\right). \tag{11}$$

Of course, $\alpha^*(B_1) = 1$. For $b_1 < B_1$, both terms on the right-hand side are concave in α by Lemma 4 in Appendix A.6, such that $\alpha^*(b_1)$ is uniquely characterized by the first-order condition, i.e.,

$$\begin{aligned} \alpha^* &= 1 && \text{if } R(B_2, B_2 + b_1) \geq 1 \\ R\left(\frac{B_2}{\alpha^*}, \frac{B_2 + b_1}{\alpha^*}\right) &= R\left(0, \frac{B_1 - b_1}{1-\alpha^*}\right) && \text{if } R(B_2, B_2 + b_1) < 1 \end{aligned} \tag{12}$$

The first case in (12) describes a corner solution, where sellers' marginal contribution to surplus is higher in the submarket with the high average valuation even when all sellers are allocated to this submarket. The second case describes an interior solution, where sellers' marginal contributions must be the same across the two submarkets. Note that $\alpha^*(b_1)$ can never be 0, because the planner will never leave high-type buyers idle.

Allocation of Low-Type Buyers. Having solved for $\alpha^*(b_1)$, we next consider the allocation of low-type buyers. By the envelope theorem, $\overline{S}(b_1)$ is differentiable and for $b_1 < B_1$,

$$\overline{S}'(b_1) = T_1\left(\frac{B_2}{\alpha^*(b_1)}, \frac{B_2 + b_1}{\alpha^*(b_1)}\right) - T_1\left(0, \frac{B_1 - b_1}{1 - \alpha^*(b_1)}\right). \tag{13}$$

That is, the additional surplus generated by moving one low-type buyer from the low-average-valution to the high-average-valuation submarket is simply the difference between the buyer's marginal contributions to surplus in the two submarkets.

The special case $b_1 = B_1$ warrants discussion as it is not defined by the above equation (because it leads to 0/0 in the final argument). In this case, the planner allocates all sellers and buyers to the submarket with the high average valuation, and considers the welfare effect of moving an ε number of low-type buyers to a separate submarket. Whether the planner also moves sellers to this separate submarket depends on $R(B_2, B_2 + B_1)$. In particular, if $R(B_2, B_2 + B_1) \geq 1$, then the planner will keep all sellers in the submarket with the high average valuation, because sellers' contribution to surplus in the second submarket is bounded by 1; in contrast, if $R(B_2, B_2 + B_1) < 1$, then the planner will move a small number of sellers to the separate submarket to equalize sellers' contribution to surplus across the two submarkets. Regardless,

$$\overline{S}'(B_1) = T_1(B_2, B_2 + B_1) - T_1(0, \lambda) \tag{14}$$

where λ is such that $R(B_2, B_2 + B_1) = R(0, \lambda)$ if $R(B_2, B_2 + B_1) < 1$, otherwise $\lambda = \infty$.

The optimal b_1, b_1^*, must then satisfy the first-order condition, i.e.,

$$\overline{S}'(b_1^*) \leq 0 \text{ if } b_1^* = 0; \quad \overline{S}'(b_1^*) = 0 \text{ if } 0 < b_1^* < B_1; \quad \overline{S}'(b_1^*) \geq 0 \text{ if } b_1^* = B_1. \tag{15}$$

It turns out that the first-order condition is sufficient even though $\overline{S}(b_1)$ is not necessarily concave. The following proposition formalizes our results.[28]

Proposition 2. *Under Assumption 1 and 2, at the social optimum, all sellers are active, and there are at most two submarkets, one of which contains all high-type buyers and has a shorter queue. Furthermore, the planner's solution is unique, and the first-order conditions* (12) *and* (15) *are necessary and sufficient.*

Proof. See Appendix A.7. □

The role of Assumption 1 is technical. It implies that if multiple submarkets were to have the same fraction of high-type buyers, then the planner should always merge the market since total surplus is always concave in λ. Hence, different submarkets must have different fractions of high-type buyers. In other words, Assumption 1 implies that the sellers' marginal contribution to surplus $R(\lambda z, \lambda)$ is strictly decreasing in λ for a given z, such that the level curves of $R(\lambda z, \lambda)$ are well defined.

The fact that all high-type buyers visit the same submarket then follows from Assumption 2. As discussed above, this assumption implies that each level curve of $R(\lambda z, \lambda)$ consists of two segments. Submarkets in the segment where z is small cannot be part of an equilibrium because there is always marginal gain from separation (the determinant of the Hessian matrix of the surplus function is negative which means that the second-order condition is violated). In the segment where z is large, the marginal contribution to surplus of high-type buyers is strictly decreasing in z along the level curve, which means that they can be present in at most one submarket.

[28] As we show later, the market equilibrium decentralizes the planner's solution and therefore features endogenous market segmentation where both sellers and low-type buyers are indifferent between different segments. Barro and Romer (1987) give a nice example that illustrates how sellers can promise utility by either a low price or fewer other buyers: the Paris metro used to sell expensive first-class tickets for wagons which were physically similar to the second-class ones but which were less crowded in equilibrium.

Algorithm. Our analysis suggests a simple numerical algorithm to solve the planner's problem: start with $b_1 = B_1$ (i.e., pooling) and compute $\overline{S}'(B_1)$ according to equation (14); if $\overline{S}'(B_1) \geq 0$, then $b_1 = B_1$ is the solution, otherwise decrease b_1 until the first-order condition is satisfied or $b_1 = 0$.

The knife-edge case $\overline{S}'(B_1) = 0$ deserves special attention because it pins down the boundary between areas of pooling and partial segmentation. The detailed analysis of this special case is technical and delegated to Appendix B.1.

3.5 Example: Stochastic Capacity

We now briefly discuss the social planner's solution for the stochastic capacity technology; the details are relegated to Appendix B.2.[29] For this technology, the model has 4 parameters: x_2, which measures the dispersion of buyer values, B_1 and B_2, the number of low-type and high-type buyers (recall that the number of sellers is normalized to be 1), and σ, which captures the extent to which sellers can meet multiple buyers. We distinguish between three regions that follow from two cutoff values for σ, i.e., $\sigma_0(x_2)$ and $\sigma_1(x_2)$, where

$$\sigma_0(x_2) \equiv \frac{\sqrt{x_2} - 1}{\sqrt{x_2} + 1} < \frac{\sqrt{x_2}}{\sqrt{x_2} + 1} \equiv \sigma_1(x_2) \tag{16}$$

Case 1: Low σ. Consider first the case in which $\sigma \leq \sigma_0(x_2)$, where the cutoff $\sigma_0(x_2)$ follows from the condition that the infimum of $H(\mu, \lambda)$ equals $1/(x_2 - 1)$. The second-order condition (8) can never be satisfied in this case, i.e., a submarket (μ, λ) where $0 < \mu < \lambda$ cannot be part of the planner's solution. Instead, perfect separation is obtained: to prevent strong congestion effects of low-type buyers on high-type buyers, the planner sends the two types of buyers to different submarkets. We thus prove analytically that for any given x_2, there exists a meeting technology such that full separation is always optimal for any endowment of buyers. This proves the conjecture in Section 5.3 of Eeckhout and Kircher (2010b) who only showed the existence of such a meeting technology numerically.

[29] In that appendix, we also characterize the planner's solution for a second example of a meeting technology.

Case 2: Intermediate σ. The case $\sigma \in (\sigma_0(x_2), \sigma_1(x_2)]$ is illustrated in Figure 3a. First, consider the scenario where $(B_2, B_1 + B_2)$ belongs to the brown area. In that case, the planner's solution consists of two active submarkets: 1) a submarket located on the line segment AC with only low-type buyers, and 2) a submarket located on the curve AB with all high-type buyers and possibly some low-type buyers. Denote the queue length in the first submarket by λ_0 (which is between 2 and 7 in the figure). The queue lengths in the second submarket are then a function of λ_0, i.e., $(\mu(\lambda_0), \lambda(\lambda_0))$, such that the marginal contribution to surplus of sellers and that of low-type buyers are equal across both submarkets. Graphically, both submarkets and $(B_2, B_1 + B_2)$ all lie on the same straight line.

If $(B_2, B_1 + B_2)$ belongs to the blue area, then the optimum is full separation where one submarket contains all high-type buyers and the other contains all low-type buyers. Whether the submarket with low-type buyers contains sellers depends on B_2: if the marginal contribution to surplus by sellers in the segment with high-type buyer is greater than 1, then this submarket contains all sellers, otherwise both submarkets contain sellers. Finally, when $(B_2, B_1 + B_2)$ belongs to the white area, then the optimum is pooling where one market contains all sellers and buyers.

As σ approaches $\sigma_1(x_2)$, point C goes to $\lambda_0 = \infty$, making the line BC horizontal in the limit. Finally, the green curve represents the set of queues where the marginal contribution to surplus by the low-type buyers equals zero, which only becomes relevant when $\sigma > \sigma_1(x_2)$, as we discuss next.

Case 3: Large σ. The case $\sigma > \sigma_1(x_2)$ is illustrated in Figure 3b. If $(B_2, B_1 + B_2)$ belongs to the brown area in Figure 3b, then it is again optimal to open two active submarkets, where the queue in the first submarket lies on the curve AB and the second submarket lies on the horizontal axis, characterized by the same marginal contribution to surplus of sellers and of low-type buyers (same as case 2). If $(B_2, B_1 + B_2)$ belongs to the yellow area, then it is optimal to have one active submarket and one inactive submarket. The queue in the first submarket lies on the curve BD, where the marginal contribution to surplus by the low-type buyers is

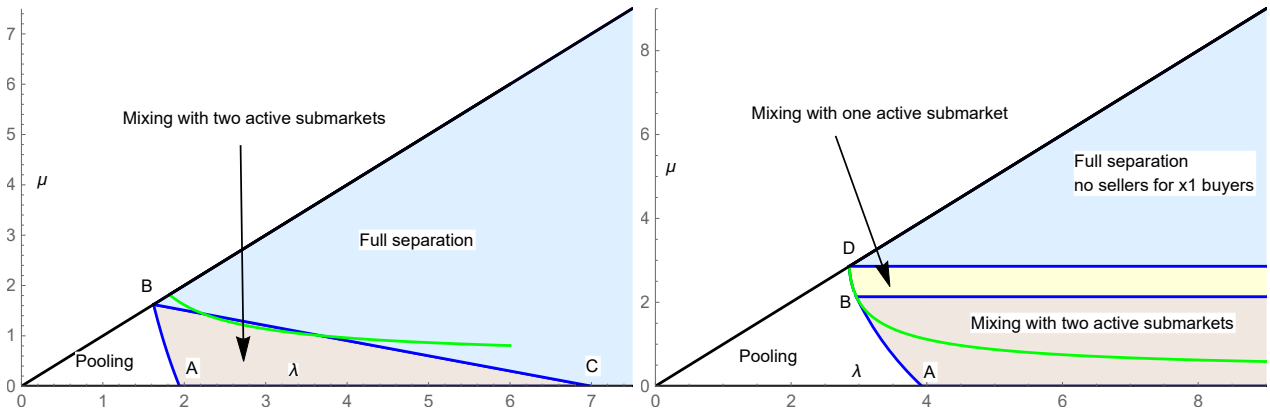

(a) $\sigma = 0.45 < \sigma_1(x_2)$ (b) $\sigma = 0.65 > \sigma_1(x_2)$

Figure 3: Planner's Solution for the Stochastic Capacity Technology with $x_2 = 2$

zero, and the second submarket contains only low-type buyers and no sellers. If $(B_2, B_1 + B_2)$ belongs to the blue area, then the optimum is full separation where one submarket contains all sellers and high-type buyers and the other contains all low-type buyers and no sellers. Finally, when $(B_2, B_1 + B_2)$ belongs to the white area, then the optimum is pooling where one market contains all sellers and buyers.

3.6 Comparative Statics

Having solved the planner's problem, we now analyze comparative statics. To simplify exposition, we focus again on the example of the stochastic capacity meeting technology. In Appendix B.3, we show that the same results hold for general meeting technologies under certain assumptions. We are particularly interested in how the optimal allocation of buyers and sellers varies with (i) the screening parameter σ, and (ii) the high-type buyer's value x_2. The first comparative static can be thought of as analyzing the effect of new technologies like automated resume screening, while the second comparative static can be thought of as analyzing the effect of an increase in the dispersion of buyers' values.

Changes in Screening Capacity. For the stochastic capacity technology, the extent to which sellers can screen buyers is captured by the parameter σ. Note that $\phi(\mu, \lambda)$ is increasing in σ, while $\phi(\lambda, \lambda)$ is independent of σ. That is, a better screening technology increases the

probability that a seller finds a high-type buyer, but does not change the probability that a seller meets at least one buyer. The following results show that when the screening technology improves, the separation area shrinks and the area where pooling is optimal increases.

Proposition 3. *Given (x_2, σ, B_1, B_2), if the optimal allocation is complete pooling, then for (x_2, σ', B_1, B_2) with $\sigma' > \sigma$, the optimal allocation is again complete pooling.*

Given (x_2, σ, B_1, B_2), if the optimal allocation is complete separation, then for (x_2, σ', B_1, B_2) with $\sigma' < \sigma$, the optimal allocation is again complete separation.

Proof. See the general results of Proposition 8 and 9 in Online Appendix B.3. □

When $\sigma = 0$, meetings are always bilateral and the optimal allocation is always complete separation, since a low-type buyer meeting a seller always crowds out high-type buyers. When $\sigma = 1$, low-type buyers do not impose negative meeting externalities on high type buyers and therefore it is optimal to pool all buyers and sellers in one market. Thus, for any (x_2, B_1, B_2), the above proposition shows that the optimal allocation changes smoothly from complete separation to complete pooling as σ increases from 0 to 1.

Changes in the Dispersion of Buyer Values. Since we are mainly interested in the optimal allocation of buyers and sellers, any variation of buyer values is equivalent to a corresponding change in x_2 while fixing $x_1 = 1$ for this purpose, which implies that we can use the parameter x_2 to measure the dispersion in buyer values. As we increase x_2, the output loss due to low-type buyers crowding out high-type buyers becomes larger. One may therefore expect that complete separation becomes a more likely outcome while complete pooling becomes less likely. The following proposition presents the formal results.

Proposition 4. *Given (x_2, σ, B_1, B_2), if the optimal allocation is complete pooling, then for (x_2', σ, B_1, B_2) with $x_2' < x_2$, the optimal allocation is again complete pooling.*

Given (x_2, σ, B_1, B_2), if the optimal allocation is complete separation, then for (x_2', σ, B_1, B_2) with $x_2' > x_2$, the optimal allocation is again complete separation.

Proof. This result follows from Proposition 10 and 11 in Online Appendix B.3. □

When $x_2 \to x_1 = 1$ and $\sigma > 0$, then the optimal allocation is complete pooling because the gain from partial separation is negligible, and the planner prefers to pool all buyers and sellers in one place to maximize the matching probability. When x_2 is sufficiently large and $\sigma < 1$, it is optimal to exclude the low-type buyers from participating and set up one market for all sellers and high-type buyers. Thus, for any (σ, B_1, B_2) with $0 < \sigma < 1$, the above proposition shows that the optimal allocation always changes from complete pooling to complete separation as x_2 increases.

As we will show below, when the optimal allocation features complete separation, it can be decentralized by sellers posting fixed prices, while when the optimal allocation is complete pooling, then the decentralized equilibrium necessarily involves auctions with reserve prices or fees. Thus as σ changes, the optimal trading mechanism changes accordingly.

4 Market Equilibrium

In this section, we show that the equilibrium is constrained efficient and that no seller can do better than posting a second-price auction combined with either a reserve price or a meeting fee. The reserve price can be positive or negative, where the latter just means that the seller is willing to sell the good at a price below his valuation, which we normalized to 0. Similarly, the meeting fee can be positive, in which case it is paid by each buyer meeting the seller, or negative, in which case payments take place in the opposite direction.

4.1 Equilibrium Definition

Strategies. Let D be the set of all direct anonymous mechanisms equipped with some natural σ-algebra $\mathcal{D}$. A seller's strategy is a probability measure δ^s on $(D, \mathcal{D})$. A buyer needs to decide on whether or not to participate in the market, and if he does, which sellers (who are characterized by the mechanisms they post) to visit. To acknowledge that a buyer's strategy depends (only) on his value x_k and the fact that—due to the lack of coordination—buyers treat all sellers who post the same mechanism symmetrically, we denote his strategy

by δ_k^b, a measure on $(D, \mathcal{D})$. If $\delta_k^b(D) < 1$, then buyers with value x_k will choose not to participate in the market with probability $1 - \delta_k^b(D)$, in which case their payoff will be zero.[30] To simplify the exposition, we require that the measure δ_k^b is absolutely continuous with respect to δ^s.[31] The Radon-Nikodym derivative $d\delta_k^b/d\delta^s$ determines the queue length and queue composition—i.e., how many buyers and what types of buyers are available per seller—for each mechanism (almost surely) in the support of δ^s. Formally, for (almost every) mechanism ω in the support of δ^s and $k = 1, 2$, the queue length of buyers with value x_k, $q_k(\omega)$, is given by

$$q_k(\omega) = B_k \frac{d\delta_k^b}{d\delta^s}. \tag{17}$$

The queue (μ, λ) in the submarket is thus given by $\boldsymbol{q}(\omega) \equiv (q_2(\omega), q_1(\omega) + q_2(\omega))$.

Payoffs. Note that for any mechanism $\omega \in D$, the expected payoff of a seller who posts mechanism ω is completely determined by ω and its queue $\boldsymbol{q}(\omega)$. Therefore, we can denote it by $\pi(\omega, \boldsymbol{q}(\omega))$. Similarly, let $V_k(\omega, \boldsymbol{q}(\omega))$ denote the expected payoff of a buyer with value x_k from visiting a submarket with mechanism ω which has queue $\boldsymbol{q}(\omega)$.

Market Utility and Beliefs. We now define conditions on buyers' and sellers' strategy $(\delta^s, \delta_1^b, \delta_2^b)$ which need to be satisfied in equilibrium. First, consider the optimality of buyers' strategies. The *market utility function* U_k is defined to be the maximum utility that a buyer with value x_k can obtain by visiting a seller or being inactive.

$$U_k = \max\left(\max_{\omega \in \text{supp}(\delta^s)} V_k(\omega, \boldsymbol{q}(\omega)), 0\right),$$

[30] The assumption that all sellers post a mechanism is without loss of generality, because they can stay inactive by posting a sufficiently inattractive mechanism, e.g., a reserve price above x_2.

[31] This rules out a scenario in which a zero measure of sellers attracts a positive measure of buyers of type x_k, in which case the probability for each of these buyers to obtain the good is zero, which violates optimality.

where $\boldsymbol{q}(\omega)$ is given by equation (17). Of course, optimality of buyers' choices requires that buyers choose the mechanism that yields the highest payoff. Formally, we have

$$V_k(\omega, \boldsymbol{q}(\omega)) \leq U_k \quad \text{with equality if } \omega \text{ is in the support of } \delta_k^b.$$

Next, we consider the optimality of sellers' strategies. All posted mechanisms should generate the same expected payoff π^* and there should be no profitable deviations. A seller considering a deviation to a mechanism $\widetilde{\omega}$ not in the support of δ^s needs to form beliefs regarding the queue $\boldsymbol{q}(\widetilde{\omega})$ that he will be able to attract. We call a queue $\boldsymbol{q}(\widetilde{\omega})$ *compatible* with the mechanism $\widetilde{w}$ and the market utility function U_k if for any $k \in \{1, 2\}$,

$$V_k(\widetilde{\omega}, \boldsymbol{q}(\widetilde{\omega})) \leq U_k \quad \text{with equality if } q_k(\widetilde{\omega}) > 0. \tag{18}$$

Of course, for any mechanism ω in the support of δ^s, $\boldsymbol{q}(\omega)$ is compatible with mechanism ω and the market utility function because of the optimal search behavior of buyers. The literature usually assumes that when posting $\widetilde{w}$, the seller is optimistic and expects the most favorable queue among all queues that are compatible with $\widetilde{\omega}$ and the market utility function (see, for example, McAfee, 1993; Eeckhout and Kircher, 2010a,b). That is,

$$\boldsymbol{q}(\widetilde{\omega}) = \arg\max_{\widetilde{\boldsymbol{q}}} \pi(\widetilde{\omega}, \widetilde{\boldsymbol{q}}) \tag{19}$$

where the choice of $\widetilde{\boldsymbol{q}}$ is subject to the constraint in equation (18).[32] For simplicity of exposition, we adopt this convention, but we show in Appendix C that this assumption is largely unnecessary because there is only one possible queue compatible with $\widetilde{\omega}$ and the market utility function under mild restrictions on the meeting technology.

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

[32] For some mechanism $\widetilde{w}$ there may not exist a compatible queue because $\widetilde{w}$ is either too attractive or too unattractive. If $\widetilde{w}$ is too unattractive, we can set $\widetilde{\boldsymbol{q}}$ to be the zero vector. In Section 4, we show that sellers can not do better than posting a second-price auction with a reserve price, which implies that $\widetilde{w}$ will not be too attractive in the above sense.

Definition 1. *A directed search equilibrium is a tuple* $(\delta^s, \delta_1^b, \delta_2^b)$ *of strategies with the following properties:*

1. *Each* ω *in the support of* δ^s *maximizes* $\pi(\omega, \boldsymbol{q}(\omega))$*, where, depending on whether or not* ω *belongs to the support of* δ^s*,* $\boldsymbol{q}(\omega)$ *is given by equations* (17) *and* (19)*, respectively.*

2. *For each buyer type* x_k*, if* $\delta_k^b(D) > 0$*, every* ω *in the support of* δ_k^b *maximizes* $V_k(\omega, \boldsymbol{q}(\omega))$*. If* $\delta_k^b(D) = 0$*, then for any mechanism* ω *in the support of* δ^s *the buyer value* $V_k(\omega, \boldsymbol{q}(\omega))$ *is non-positive.*

3. *Aggregating queues across sellers does not exceed the total measure of buyers of each type. That is,* $\int q_k(\omega) d\delta^s(w) \leq B_k$ *for each* $k \in \{1, 2\}$*.*

4.2 Efficiency

Equivalence. To prove constrained efficiency of equilibrium, we show that even if sellers can buy queues directly in a hypothetical competitive market, they cannot do better than in the decentralized environment.[33] In other words, the following two problems are equivalent for sellers.

1. *Sellers' Relaxed Problem*, in which there exists a hypothetical competitive market for queues, with the price for each buyer given by the market utility function. That is, sellers choose a queue (μ, λ) to maximize

$$\pi(\mu, \lambda) = m(\lambda) + (x_2 - 1)\,\phi(\mu, \lambda) - \mu U_2 - (\lambda - \mu) U_1, \tag{20}$$

   where the first two terms are total surplus (3) and the last two terms are the price of the queue.

[33] A similar result appears in Cai et al. (2017), so it is worth emphasizing that the credit belongs with the current paper: as they explicitly acknowledge in their article, Cai et al. (2017) borrow Proposition 4 directly from our paper, of which a first draft was written in 2016. The same applies to a number of other results, e.g., Lemma 3 here vs. Lemma 1 in Cai et al. (2017).

2. *Sellers' Constrained Problem*, in which sellers must post mechanisms to attract queues of buyers, as described in detail in Section 4.1. For any mechanism, the corresponding queue must be compatible with the market utility function, which means that it needs to satisfy equation (18). In this case, a seller's profit is again given by equation (20), assuming that sellers post efficient mechanisms, but now queue length and queue composition depend on the posted mechanism.

In the relaxed problem, a seller will "buy" queues of buyers with valuation x_k until their expected marginal contribution T_k to surplus is equal to their marginal cost U_k, where $k = 1, 2$. Hence, if sellers can post a mechanism which delivers buyers their marginal contribution to surplus, then buyers' payoffs are equal to their market utility and the queue is compatible with the mechanism and the market utility function, as defined by (18). The following proposition establishes that auctions with an entry fee or a reserve price can achieve this.[34]

Proposition 5. *Any solution (μ, λ) to the sellers' relaxed problem is compatible with an auction with an entry fee in the sellers' constrained problem, where the fee is given by*

$$t = -\frac{(x_2 - 1)\phi_\lambda(\mu, \lambda) + \phi_\lambda(\lambda, \lambda)}{1 - Q_0(\lambda)}. \tag{21}$$

It is also compatible with an auction with a reserve price in the sellers' constrained problem, where the reserve price is given by

$$r = -\frac{(x_2 - 1)\phi_\lambda(\mu, \lambda) + \phi_\lambda(\lambda, \lambda)}{Q_1(\lambda)}. \tag{22}$$

Hence, the directed search equilibrium is constrained efficient for any meeting technology.

Proof. See Appendix A.8. □

Uniqueness. As shown before, the planner's solution is unique under Assumption 1 and 2. By Proposition 5, the equilibrium mechanism is not unique: the planner's solution can be

[34] As mentioned before, when Assumption 1,and 2 do not hold, the planner's solution is not necessarily unique. The above proposition then implies that any optimal allocation can be decentralized (different optimal allocations are then associated with different equilibrium entry fees or reserve prices).

decentralized in multiple ways, either by auction with meeting fees or by auctions with reserve prices. In a submarket with low-type buyers only, a second-price auction with a reserve price is equivalent to price posting; so, price posting can also be an equilibrium mechanism.

When Assumption 1 or 2 does not hold, the planner's solution is not necessarily unique either. Proposition 5 then implies that any optimal allocation can be decentralized, each one associated with different equilibrium entry fees or reserve prices.

Auctions with Meeting Fees. The intuition behind the case with meeting fees is the following. Recall that a buyer's marginal contribution T_k consists of two parts: (i) a direct effect, representing the fact that the buyer may increase the maximum valuation among the group of buyers meeting the seller,[35] and (ii) an indirect effect, $(x_2 - 1)\phi_\lambda(\mu, \lambda) + \phi_\lambda(\lambda, \lambda)$, representing the externalities that the buyer may impose by making it easier or harder for the seller to meet other buyers. As is well-known, auctions (without reserve prices or fees) provide buyers with a payoff equal to their direct contribution.[36] Buyers' indirect effect on surplus is independent of their type and can therefore be priced by an appropriate entry fee. Since buyers pay the fee whenever they meet a seller, which happens with probability $1 - Q_0(\lambda)$, a meeting fee equal to (21) guarantees that their expected payoff from the mechanism equals exactly T_k, which yields the desired result.

Auctions with Reserve Prices. Perhaps surprisingly, an auction with an appropriate reserve price is also an efficient mechanism that can price all meeting externalities. After all, in contrast to meeting fees, reserve prices may prevent efficient trade. To see this, consider a seller who sets a reserve price $r \in (1, x_2)$. Low-type buyers have a zero trading probability at this seller, while their trading probability would be strictly positive at an auction by the same seller with a meeting fee. However, this difference between the two mechanisms only affects out-of-equilibrium behavior; in equilibrium, low-type buyers would visit neither seller.

[35] For low-type buyers, the direct effect is given by $\phi_\mu(\lambda, \lambda) = Q_1(\lambda)$, and for high-type buyers, the direct effect is given by $(x_2 - 1)\phi_\mu(\mu, \lambda) + \phi_\mu(\lambda, \lambda)$.

[36] This is easiest to see in a second-price auction. Suppose that the highest and the second highest value are x_2 and x_1. Then, the payoff for the highest value buyer is $x_2 - x_1$, which is also his contribution to surplus. Other bidders receive zero and their contributions to the surplus of the auction are also zero. Extension of this result to other auction formats follows from revenue equivalence.

High-type buyers are only affected by the reserve price when they are the only bidder, which happens with probability $Q_1(\lambda)$. A reserve price equal to (22) therefore guarantees that buyers' expected payoff again equals T_k.

Meeting Fees vs. Reserve Prices. Although the meeting fee is a useful instrument from a theoretical point of view, one could argue that it may be difficult to implement in practice. For example, if the meeting fee is positive, fake sellers with no intent to sell could open phantom auctions to collect the meeting fees from interested buyers, which would then discourage buyers from visiting sellers who charge fees in the first place.[37] Those concerns do not apply to auctions with reserve prices. The optimal reserve price has the same sign and plays a similar role as the optimal meeting fee, but is easier to implement because all buyers who do not win, pay (or receive) nothing. If, however, some buyers have valuations below the sellers' reservation value and the meeting externalities are positive, then auctions with negative reserve prices are not efficient, while auctions with entry subsidies *and* a reserve price equal to the sellers' valuation remain efficient.

Finally, consider Lester et al. (2015), where buyers are ex ante identical and learn their valuation only upon meeting the seller. In their framework, just as in ours, pricing negative meeting externalities would require a positive reserve price. However, unlike in our setup, a positive reserve price in their model would actually prevent mutually beneficial trade in equilibrium: as buyers' valuations are only revealed ex post, the highest buyer valuation is between the seller's own valuation and his reserve price with positive probability. This inefficiency prevents sellers from adopting reserve prices in equilibrium, instead they always opt for meeting fees. In Albrecht et al. (2014), who restrict attention to urn-ball meetings ($\phi_\lambda = 0$), this inefficiency does not arise, because sellers always choose to set their reserve price equal to their valuation.

Restricted Mechanism Space. We have shown that despite the potential presence of spillovers in the meeting process, business stealing externalities and agency costs, the com-

[37] Similarly, negative meeting fees are subsidies that could attract fake buyers with no intent to purchase.

peting mechanisms problem reduces to one where sellers can buy queues in a competitive market. This result, of course, requires a sufficiently large mechanism space. If it is not possible for sellers to either commit to a reserve price above their valuation or charge fees, the decentralized equilibrium will only be efficient for invariant meeting technologies (i.e., $\phi_\lambda = 0$). If $\phi_\lambda < 0$ (resp. > 0), buyers impose negative (resp. positive) externalities on other meetings and will receive more (resp. less) than their marginal social contribution.[38]

5 N Buyer Types

5.1 Surplus

In this section, we consider the case with N buyer values: $0 < x_1 < \cdots < x_N$. The measure of x_k buyers is B_k for $k = 1, \ldots, N$. The rest of the model remains the same, including the planner's problem and the definition of the decentralized equilibrium. Let q_k denote the number of x_k buyers per seller in a submarket. To use our alternative representation of meeting technologies, we apply a change of notation and define μ_k as the queue length of buyers with value x_k or higher, i.e., $\mu_k = q_k + \cdots + q_N$ for $k = 1, \ldots, N$. The queue in the submarket can then be represented by $\boldsymbol{\mu} \equiv (\mu_1, \ldots, \mu_N)$, where μ_1 is the total queue length. Thus $\phi(\mu_k, \mu_1)$ is the probability that a seller meets at least one buyer with value x_k or higher. We further adopt the convention $x_0 \equiv 0$ to simplify notation.

The following Lemma extends Lemma 1 and 2 to the case of N buyer values. It turns out that this general case does not add much complexity. The interpretation of equations (23) to (26) closely resembles the corresponding interpretation in the two-type case. Here, we only discuss equation (24) as an example and omit the others.

In equation (24), the first term of $T_k(\boldsymbol{\mu})$ reflects the direct contribution to surplus of a buyer with valuation x_k when this buyer has the highest value in an n-to-1 meeting; this contribution equals the difference between the highest and the second-highest buyer values. The second term of $T_k(\boldsymbol{\mu})$ represents the externalities that the buyer may impose on other

[38] With free entry of sellers, the buyer-seller ratio would be too high (resp. too low) in this case.

buyers and the seller. It does not depend on k, because the meeting function treats all buyers symmetrically. Specifically, if a buyer makes it easier for the other buyers to meet the seller ($\phi_\lambda \geq 0$), he increases total surplus through a positive meeting externality, even if he does not have the highest value. A similar logic applies for negative meeting externalities ($\phi_\lambda \leq 0$).

Lemma 3. *Consider a submarket with a measure 1 of sellers and a queue* $\boldsymbol{\mu} \equiv (\mu_1, \ldots, \mu_N)$ *of buyers. Total surplus in the submarket then equals*

$$S(\boldsymbol{\mu}) = \sum_{j=1}^{N} (x_j - x_{j-1}) \phi(\mu_j, \mu_1) \tag{23}$$

The marginal contribution to surplus of a buyer with valuation x_k *equals*

$$T_k(\boldsymbol{\mu}) = \sum_{j=1}^{k} (x_j - x_{j-1}) \phi_\mu(\mu_j, \mu_1) + \sum_{j=1}^{N} (x_j - x_{j-1}) \phi_\lambda(\mu_j, \mu_1). \tag{24}$$

A seller's marginal contribution to surplus equals

$$R(\boldsymbol{\mu}) = \sum_{j=1}^{N} (x_j - x_{j-1}) \left[\phi(\mu_j, \mu_1) - \mu_j \phi_\mu(\mu_j, \mu_1) - \mu_1 \phi_\lambda(\mu_j, \mu_1) \right]. \tag{25}$$

The Hessian matrix of the surplus function $S(\boldsymbol{\mu})$ *is negative definite if and only if*

$$-m''(\mu_1) x_1 - \sum_{k=2}^{N} (x_k - x_{k-1}) \left(\phi_{\lambda\lambda}(\mu_k, \mu_1) - \frac{\phi_{\mu\lambda}(\mu_k, \mu_1)^2}{\phi_{\mu\mu}(\mu_k, \mu_1)} \right) > 0. \tag{26}$$

Proof. See Appendix D.1. □

Equilibrium. We can again show that in equilibrium, sellers can do no better than posting second-price auctions with an entry fee or a reserve price, and their relaxed problem and constrained problem are equivalent so that the decentralized equilibrium is constrained efficient. For the proofs and derivations, we refer to Appendix D.2 and D.3.

5.2 Queues Across Submarkets

A larger number of buyer types increases the complexity of the planner's problem. Although the result in Proposition 1 generalizes in a straightforward way—i.e., with N types of buyers, no more than $N+1$ submarkets are required—a full characterization of these submarkets quickly becomes intractable.[39] Nevertheless, we provide a partial characterization by showing that we can compare queue compositions between any two submarkets in terms of first-order stochastic dominance under Assumption 1 and one additional mild assumption.

This additional assumption concerns the probability $\phi_\mu(\lambda z, \lambda)$ that a high-type buyer increases surplus directly—i.e., faces no competition from other high-type buyers. We assume that this probability decreases if we add more buyers to the queue, holding the fraction of high-type buyers constant at z.

Assumption 3. *$\phi_\mu(\lambda z, \lambda)$ is strictly decreasing in λ for $0 < z \leq 1$.*

As mentioned after equation (2), $\phi_\mu(\lambda, \lambda) = Q_1(\lambda)$ and $\phi_\mu(0, \lambda) = 1 - Q_0(\lambda)$. Thus, Assumption 3 implies that (i) in submarkets with longer queues, it is strictly less likely that a buyer turns out to be the only one present, and, by continuity, (ii) buyers are weakly less likely to meet a seller if the queue length in the submarket increases, which could be interpreted as a form of congestion.

First-Order Stochastic Dominance. Consider two arbitrary submarkets, indexed by $\ell \in \{a, b\}$, that attract a queue $\boldsymbol{\mu}^\ell$ of buyers. The following proposition compares queue compositions between the submarkets in terms of first-order stochastic dominance.

Proposition 6. *Consider two submarkets a and b with respective queues $\boldsymbol{\mu}^a$ and $\boldsymbol{\mu}^b$, satisfying*

[39] Unless, of course, the meeting technology is bilateral or jointly concave, which lead to perfect separation and perfect pooling, respectively. We conjecture that the following result, which is similar to Proposition 2, continues to hold: At the social optimum, there will be one submarket for all x_N buyers. Note that this conjecture has sharp predictions. If we take the submarket for x_N buyers out, then by the same logic, in the remaining submarkets there will be exactly one which contains all the remaining x_{N-1} buyers. Repeating this logic, implies then that there will be N submarkets where the highest buyer type in the ℓ-th submarket is $x_{N+1-\ell}$ and some submarkets can be idle. However, we were unable to prove this conjecture.

$\mu_1^a > \mu_1^b$. *If Assumption 1 and 3 hold, then for any k,*

$$\frac{\mu_k^b}{\mu_1^b} \geq \frac{\mu_k^a}{\mu_1^a}. \tag{27}$$

Proof. See Appendix D.4. □

This result is quite remarkable. It shows that under two weak assumptions on the meeting technology, the buyer value distribution of a short queue always first-order stochastically dominates that of a long queue. A simple consequence of this result is that the shorter queue always has a weakly higher upper and a weakly higher lower bound.

To understand the above proposition, assume that buyers with valuations x_{k-1} and x_k both visit submarkets a and b with positive probability. Since at the planner's solution, the marginal contribution of buyers with valuations x_{k-1} and x_k must be the same across the two submarkets, by equation (24) we have $\phi_\mu(\mu_k^a, \mu_1^a) = \phi_\mu(\mu_k^b, \mu_1^b)$. If $\mu_1^a > \mu_1^b$, by Assumption 3, queue a must have a lower proportion of buyers with values weakly greater than x_k.

Proposition 6 offers some testable implications that do not require characterization of the entire model. Also, since the assumptions are rather weak, they apply to almost all meeting technologies that are currently used in the literature. Consider for example two identical goods that are offered on eBay where the queue lengths and the buyer value distributions differ. Our theory derives a sharp prediction on the relation between the queue length and the buyer value distribution.

Proposition 6 is useful beyond the specific environment that we consider here. To see this, suppose that we add an epsilon degree of seller heterogeneity to the model. The equilibrium allocation of buyers and sellers will then change marginally. Without Proposition 6, we cannot order the resulting buyer value distributions of different types of sellers, making an analysis of sorting in terms of first-order stochastic dominance impossible.

6 Conclusion

In this paper, we develop a framework to study how technological innovations that make it easier for sellers to meet multiple buyers (e.g., the internet) may impact equilibrium trading mechanisms and allocations. In particular, we analyze an environment in which sellers compete for heterogeneous buyers by posting trading mechanisms and meetings are governed by a meeting technology from a wide class. We show how the equilibrium mechanism nests posted prices when sellers can meet at most one buyer and standard auctions when they are unconstrained in the number of buyers they can meet. Concerning market segmentation, when low-valuation buyers reduce the probability that sellers and high-valuation buyers meet, sellers will discourage the low-valuation buyers from visiting. This can lead to complete or partial market segmentation, depending on the dispersion of valuations and the degree of congestion in the meeting process. All high-valuation buyers are always in one segment, either with or without a subset of low-valuation buyers.

We also introduce a new function ϕ which makes the analysis of general meeting technologies tractable and allows us to generalize the competing mechanism literature. Using this function, we show that in a large economy, despite the presence of private information and possible search externalities, the directed search equilibrium is equivalent to a competitive equilibrium (where the commodities are buyer types and the prices are the market utilities). A seller can attract a desired queue by posting an auction with entry fee or by charging an appropriate reserve price which establishes the equivalence between the two equilibria.

Appendix A Additional Results and Omitted Proofs

A.1 Recoverability of P_n

Proposition 7. *If $\phi(\mu, \lambda)$ is generated by some $\{P_n(\lambda) : n = 0, 1, 2, \dots\}$, then we can recover $P_n(\lambda)$ from $\phi(\mu, \lambda)$ by*

$$P_n(\lambda) = \frac{(-\lambda)^n}{n!} \frac{\partial^n}{\partial \mu^n} (1 - \phi(\mu, \lambda)) \Big|_{\mu=\lambda}. \tag{A.1}$$

Proof. When $n = 0$, equation (A.1) is simply $P_0(\lambda) = 1 - \phi(\lambda, \lambda)$. When $n \geq 1$, by equation (1),

$$\frac{\partial^n}{\partial \mu^n} (1 - \phi(\mu, \lambda)) = \sum_{k=n}^{\infty} P_k(\lambda) n! \left(-\frac{1}{\lambda}\right)^n \left(1 - \frac{\mu}{\lambda}\right)^{k-n}.$$

Evaluating the above equation at $\mu = \lambda$ yields equation (A.1). □

A.2 Proof of Lemma 1

When a seller meets $n \geq 1$ buyers, the surplus x from the meeting equals $x_1 = 1$ with probability $\left(1 - \frac{\mu}{\lambda}\right)^n$ and equals x_2 with the complementary probability $1 - \left(1 - \frac{\mu}{\lambda}\right)^n$. Expected surplus per seller in the submarket therefore equals

$$\begin{aligned}
S &= \sum_{n=1}^{\infty} P_n(\lambda) \left[\left(1 - \frac{\mu}{\lambda}\right)^n + x_2 \left(1 - \left(1 - \frac{\mu}{\lambda}\right)^n\right)\right] \\
&= \sum_{n=1}^{\infty} P_n(\lambda) + \sum_{n=1}^{\infty} P_n(\lambda)(x_2 - 1)\left(1 - \left(1 - \frac{\mu}{\lambda}\right)^n\right) \\
&= m(\lambda) + (x_2 - 1)\left(\sum_{n=1}^{\infty} P_n(\lambda) - \sum_{n=1}^{\infty} P_n(\lambda)\left(1 - \frac{\mu}{\lambda}\right)^n\right) \\
&= m(\lambda) + (x_2 - 1)\left(1 - \sum_{n=0}^{\infty} P_n(\lambda)\left(1 - \frac{\mu}{\lambda}\right)^n\right) \\
&= m(\lambda) + (x_2 - 1)\,\phi(\mu, \lambda),
\end{aligned}$$

where we use the definition of $m(\lambda)$ for the third equality, we add a term $P_0(\lambda)$ to both summations in the large parenthesis and start the summation from $n = 0$ for the fourth equality, and use the definition of ϕ from equation (1) for the last equality.

Adding ε more low-type buyers per seller will increase the total queue length λ by ε but will not change μ, the queue length of high-type buyers. Thus, the marginal contribution to surplus by low-type sellers is the partial derivative $S_\lambda(\mu, \lambda)$, which is given by (4). Similarly, adding ε more high-type buyers per seller will not only increase the total queue length λ by ε but also increase μ by the same amount. Thus the marginal contribution to surplus by high-type sellers is the partial derivative $S_\mu(\mu, \lambda) + S_\lambda(\mu, \lambda)$, which is given by (5). Finally, after adding ε more sellers, surplus will become $(1+\varepsilon)\, S\left(\frac{\mu}{1+\varepsilon}, \frac{\lambda}{1+\varepsilon}\right)$. Taking the derivative with respect to ε yields equation (6). $\square$

A.3 Proof of Proposition 1

Recall that the social planner's problem is given by (7). Below, we rewrite (7) slightly by introducing a new function $\widehat{S}$, total surplus per agent, which has two advantages: i) the domain of $\widehat{S}$ is compact, and ii) the accounting constraints for buyers and sellers hold with equalities so that we can apply directly the Fenchel-Bunt Theorem.

Suppose that the planner creates $\widetilde{L}$ submarkets, which may include an inactive one. In submarket ℓ, the measure of sellers is $\widetilde{\alpha}^\ell$ and the measure of buyers with value x_j is $\widetilde{B}_j^\ell$ for $j = 1, 2$. Therefore, $\sum_{\ell=1}^{\widetilde{L}} \widetilde{\alpha}^\ell = 1$ and $\sum_{\ell=1}^{\widetilde{l}} \widetilde{B}_j^\ell = B_j$ for $j = 1, 2$. Define $\widetilde{z}_1^\ell = (\widetilde{B}_1^\ell + \widetilde{B}_2^\ell)/(\alpha^\ell + \widetilde{B}_1^\ell + \widetilde{B}_2^\ell)$ and $\widetilde{z}_2^\ell = \widetilde{B}_2^\ell/(\alpha^\ell + \widetilde{B}_1^\ell + \widetilde{B}_2^\ell)$, i.e., $\widetilde{z}_1^\ell$ is the fraction of buyers and $\widetilde{z}_2^\ell$ is the fraction of x_2 buyers in a submarket ℓ.

Since total surplus in each submarket exhibits constant returns to scale with respect to the number of sellers and the number of high-type and low-type buyers, we can normalize the total number of buyers and sellers in each submarket (active or inactive) to 1, and define the surplus per agent (both buyers and sellers) in submarket ℓ as $\widehat{S}(\widetilde{z}_1^\ell, \widetilde{z}_2^\ell)$. When $\widetilde{\alpha}^\ell > 0$, it

is given by

$$\widehat{S}(z_1, z_2) = \frac{1}{\alpha^\ell + \widetilde{B}_1^\ell + \widetilde{B}_2^\ell} \cdot \widetilde{\alpha}^\ell S\left(\frac{\widetilde{B}_2^\ell}{\widetilde{\alpha}^\ell}, \frac{\widetilde{B}_2^\ell + \widetilde{B}_1^\ell}{\widetilde{\alpha}^\ell}\right).$$

and when $\widetilde{\alpha}^\ell = 0$, it is simply zero. The function $\widehat{S}$ is well defined even in a submarket with buyers only ($\widetilde{z}_1^\ell = 1$) and its domain is compact. The total surplus generated from all submarkets is

$$\sum_{\ell=1}^{\widetilde{L}} (\alpha^\ell + \widetilde{B}_1^\ell + \widetilde{B}_2^\ell) \widehat{S}(\widetilde{z}_1^\ell, \widetilde{z}_2^\ell).$$

Therefore, as in equation (7), total surplus is a convex combination of the individual submarkets' surpluses, which are represented by $\widehat{S}$ here. The planner's solution is thus the supreme of all such convex combinations. Because the function $\widehat{S}$ is continuous and its domain is compact, the graph of $\widehat{S}$ is compact, which implies that the convex hull of the graph is also compact (see, for example, Theorem 17.2 of Rockafellar, 1970). Thus the supreme can be reached as a maximum. Furthermore, note that the domain of $\widehat{S}$ is the set $\{(z_1, z_2) \,|\, 0 \leq z_2 \leq z_1 \leq 1\}$, which is connected. By the Fenchel-Bunt Theorem (see Theorem 18 (ii) of Eggleston, 1958), which is an extension of Caratheodory's theorem, it suffices to create 3 submarkets. □

A.4 Proof of Lemma 2

If $P_0(\lambda) + P_1(\lambda) = 1$, then total surplus per seller is given by $m(\lambda)\,(1 + (x_2 - x_1)\,\mu/\lambda)$. We can obtain the same surplus by splitting the submarket in two: one contains only high-type buyers and the other contains only low-type buyers, where the queue lengths in both submarkets equal λ. We can further increase total surplus by moving sellers from the submarket with only low-type buyers to the other submarket with only high-type buyers.

The Hessian matrix of the surplus function is

$$\begin{pmatrix} S_{\mu\mu}(\mu,\lambda) & S_{\mu\lambda}(\mu,\lambda) \\ S_{\mu\lambda}(\mu,\lambda), & S_{\lambda\lambda}(\mu,\lambda) \end{pmatrix} = \begin{pmatrix} (x_2-x_1)\phi_{\mu\mu}(\mu,\lambda), & (x_2-x_1)\phi_{\mu\lambda}(\mu,\lambda) \\ (x_2-x_1)\phi_{\mu\lambda}(\mu,\lambda), & x_1 m''(\lambda)+(x_2-x_1)\phi_{\lambda\lambda}(\mu,\lambda) \end{pmatrix}$$

Since $\phi_{\mu\mu} < 0$, by Sylvester's criterion the Hessian matrix is negative semidefinite if and only if its determinant is positive. That is,

$$(x_2-x_1)\phi_{\mu\mu}(\mu,\lambda)\left(x_1 m''(\lambda)+(x_2-x_1)\phi_{\lambda\lambda}(\mu,\lambda)\right)-(x_2-x_1)^2\phi_{\mu\lambda}(\mu,\lambda)^2 \geq 0$$

Dividing both sides by $(x_2-x_1)\,\phi_{\mu\mu}(\mu,\lambda)$ gives (8). □

A.5 Verifying Assumptions 1 and 2

Consider first the case of bilateral technologies, where $\phi(\mu,\lambda) = 1 - P_0(\lambda) - P_1(\lambda)(1-\frac{\mu}{\lambda}) = \mu m(\lambda)/\lambda$ (recall $m(\lambda) \equiv 1 - P_0(\lambda)$). Thus $\phi(\lambda z,\lambda) = z m(\lambda)$, which is strictly concave in λ for any $z \in (0,1]$. Next, since $P_0(\lambda) + P_1(\lambda) = 1$ for all λ, the requirement of Assumption 2 becomes void and is thus always satisfied.

Next, consider technologies exhibiting joint concavity. For $\gamma \in (0,1)$, $\gamma\phi(\lambda_1 z,\lambda_1) + (1-\gamma)\,\phi(\lambda_2 z,\lambda_2) \leq \phi\left(\gamma\lambda_1 z + (1-\gamma)\lambda_2 z, \gamma\lambda_1 + (1-\gamma)\lambda_2\right)$ since $\phi(\mu,\lambda)$ is concave in (μ,λ). Thus for a given z, $\phi(\lambda z,\lambda)$ is concave in λ, i.e., Assumption 1 is satisfied. Since $\phi_{\mu\mu}\phi_{\lambda\lambda} - \phi_{\mu\lambda}^2 \geq 0$ for all μ and λ, $H(\lambda z,\lambda)$ is always weakly negative and the requirement of Assumption 2 becomes void and is thus always satisfied.

For the stochastic capacity example, we show that $\phi(\mu,\lambda) = \mu/(1+\sigma\mu+(1-\sigma)\lambda)$ in Appendix B.2. Hence, $\phi(\lambda z,\lambda) = \lambda z/(1+\lambda(1-\sigma+\sigma z))$, which is strictly concave in λ when $z > 0$. For Assumption 2, note that $H(\lambda z,\lambda)$ is always positive:

$$H(\lambda z,\lambda) = \frac{(\lambda+1)^3(1-\sigma)^2}{4\sigma(1+(1-\sigma)\lambda)(1+\lambda(1-\sigma+\sigma z))} > 0.$$

From the above equation, $\partial H(\lambda z, \lambda)/\partial z < 0$, and

$$\frac{\partial H(\lambda z, \lambda)}{\partial \lambda} = \frac{(1+\lambda)^2(1-\sigma)^2\left(\lambda^2(1-\sigma)(1-\sigma(1-z)) + 2\lambda\left(1-\sigma^2(1-z)\right) + 1 + \sigma(2-z)\right)}{4\sigma(1+(1-\sigma)\lambda)^2(1+\lambda(1-\sigma+\sigma z))^2},$$

which is strictly positive. Hence, the level curve of $H(\lambda z, \lambda)$ is upward sloping. On the other hand,

$$\frac{\partial \phi_\mu(\lambda z, \lambda)}{\partial \lambda} = -\frac{1-\sigma+2\sigma z + \lambda(1-\sigma)(1-\sigma(1-z))}{(1+\lambda(1-\sigma+\sigma z))^3} < 0,$$

which implies that the level curve of $\phi_\mu(\lambda z, \lambda)$ is downward sloping. Thus Assumption 2 is trivially satisfied.

Finally, we present an example of a meeting technology which satisfies Assumption 1 but not Assumption 2. This example is a mechanical combination of a bilateral and an invariant meeting technology. There are two locations within a submarket. With probability α, all buyers and sellers are in location 1 and meetings are dictated by a geometric technology; with complementary probability, all buyers and sellers are in location 2 and meetings are dictated by a telephone-line technology. In this case, $\phi(\mu, \lambda)$ is a weighted average, i.e., $\phi(\mu, \lambda) = (1-\alpha)\,\mu/(1+\lambda) + \alpha\mu/(1+\mu)$.

Given this functional form, it is straightforward to check that $\phi(\lambda z, \lambda)$ is strictly concave in λ when $z > 0$ so that Assumption 1 is satisfied. For Assumption 2, note that

$$H(\lambda z, \lambda) = \frac{(1-\alpha)}{4\alpha(\lambda+1)}\left[\lambda^3(1-\alpha)z^3 + \lambda^2\left(3z^2(1-\alpha) + 4\alpha z\right) + \lambda(3+\alpha)z + 1 - \alpha\right] > 0.$$

Direct computation yields that

$$\lim_{\lambda\to 0}\frac{\partial H(\lambda z, \lambda)}{\partial \lambda} = \frac{(1-\alpha)(-1+\alpha+(\alpha+3)z)}{4\alpha}$$

which is strictly negative when $z \to 0$. Hence, Assumption 2 is violated. □

A.6 Collection of Technical Lemmas

Below, we collect several technical lemmas which will be useful to characterize the planner's solution established in Proposition 2.

An alternative way of understanding Assumption 1 is the following. Holding fixed the number of low-type and high-type buyers in a submarket, adding one more seller decreases the queue length but keeps the fraction of high-type buyers constant. Assumption 1 then implies that the total surplus in this submarket is always concave in the number of sellers.

Lemma 4. *Consider a submarket where the measure of sellers, low-type buyers and high-type buyers are α, b_1, and b_2 respectively. Under Assumption 1, total surplus $\alpha S\left(\frac{b_2}{\alpha}, \frac{b_1+b_2}{\alpha}\right)$ is strictly concave in α.*

Proof. First define $\widetilde{b} = b_1 + b_2$ and $z = b_2/\widetilde{b}$. Surplus generated from the submarket is $\alpha\left(m(\frac{\widetilde{b}}{\alpha}) + (x_2-1)\phi(\frac{\widetilde{b}z}{\alpha}, \frac{\widetilde{b}}{\alpha})\right)$. Then, for any $\gamma \in (0,1)$,

$$\gamma\alpha_1\left(m\left(\frac{\widetilde{b}}{\alpha_1}\right) + (x_2-1)\,\phi\left(\frac{\widetilde{b}z}{\alpha_1}, \frac{\widetilde{b}}{\alpha_1}\right)\right) + (1-\gamma)\,\alpha_2\left(m\left(\frac{\widetilde{b}}{\alpha_2}\right) + (x_2-1)\phi\left(\frac{\widetilde{b}z}{\alpha_2}, \frac{\widetilde{b}}{\alpha_2}\right)\right)$$
$$> (\gamma\alpha_1 + (1-\gamma)\alpha_2)\left(m\left(\frac{\widetilde{b}}{\gamma\alpha_1 + (1-\gamma)\alpha_2}\right) + (x_2-1)\phi\left(\frac{\widetilde{b}z}{\gamma\alpha_1 + (1-\gamma)\alpha_2}, \frac{\widetilde{b}}{\gamma\alpha_1 + (1-\gamma)\alpha_2}\right)\right),$$

where we use the fact that $m(\lambda)$ is strictly concave and $\phi(\lambda z, \lambda)$ is concave in λ. □

Our next lemma shows how $T_1(\lambda z, \lambda)$ and $T_2(\lambda z, \lambda)$ (the marginal contribution to surplus by a buyer) vary along a level curve of $R(\lambda z, \lambda)$ and how this depends on the sign of the determinant of the Hessian matrix of the surplus function.

Lemma 5. *For any given R^*, let $\lambda(z)$ be implicitly determined by the level curve $R(\lambda z, \lambda) = R^*$. Then, $T_2(\lambda(z)z, \lambda(z))$ as a function of z is strictly decreasing at some point z_0 if the determinant of the Hessian of the surplus function in equation (3) is strictly positive for the queue $(z_0\lambda(z_0), \lambda(z_0))$ (and decreasing if the determinant is strictly negative). The reverse result holds for $T_1(\lambda(z)z, \lambda(z))$.*

Proof. Since $T_1(\lambda z, \lambda)$ and $R(\lambda z, \lambda)$ are given by Equation (4) and (6), respectively, we have

$$\begin{aligned}\frac{dT_1(\lambda z, \lambda)}{dz}\Big|_{R(\lambda z,\lambda)=R^*} &= \frac{\partial T_1(\lambda z, \lambda)}{\partial z} + \frac{\partial T_1(\lambda z, \lambda)}{\partial \lambda}\left(-\frac{\partial R(\lambda z, \lambda)/\partial z}{\partial R(\lambda z, \lambda)/\partial \lambda}\right) \\ &= \frac{z(x_2-1)\lambda\phi_{\mu\mu}}{m''(\lambda) + (x_2-1)\left(z^2\phi_{\mu\mu} + 2z\phi_{\mu\lambda} + \phi_{\lambda\lambda}\right)}\left(-m''(\lambda) - (x_2-1)\left(\phi_{\lambda\lambda} - \frac{\phi_{\mu\lambda}^2}{\phi_{\mu\mu}}\right)\right),\end{aligned}$$

where in the second line we have suppressed the arguments $(\lambda z, \lambda)$ from the relevant functions. The denominator in the first term on the right-hand side is negative because of Assumption 1, and the second term corresponds to the sign of the determinant of the Hessian of the surplus function, as demonstrated in equation (8).

The results regarding $T_2(\lambda z, \lambda)$ follow from the equation

$$\frac{dT_2(\lambda z, \lambda)}{dz}\Big|_{R(\lambda z,\lambda)=R^*} = -\frac{1-z}{z}\frac{dT_1(\lambda z, \lambda)}{dz}\Big|_{R(\lambda z,\lambda)=R^*}.$$

This equation uses the fact that the surplus function $S(\mu, \lambda)$ exhibits constant returns to scale. To understand it, consider a marginal change in the queue and let $\Delta T_k = T_k(\mu + \Delta\mu, \lambda + \Delta\lambda) - T_k(\mu, \lambda)$ for $k = 1, 2$. Since sellers' marginal contribution to surplus remains constant, we have $S(\mu, \lambda) - \mu T_2(\mu, \lambda) - (\lambda - \mu)T_1(\mu, \lambda) = S(\mu + \Delta\mu, \lambda + \Delta\lambda) - (\mu + \Delta\mu)T_2(\mu + \Delta\mu, \lambda + \Delta\lambda) - (\lambda + \Delta\lambda - \mu - \Delta\mu)T_1(\mu + \Delta\mu, \lambda + \Delta\lambda)$. Recall that $T_1 = S_\lambda$ and $T_2 = S_\mu + S_\lambda$. A first-order Taylor approximation yields $\mu\Delta T_2 = (\lambda - \mu)\Delta T_1$, which is exactly the above equation. Alternatively, it can be proved by direct computation. □

Our last lemma shows the single-crossing result, which complements Lemma 5 and is critical to the result that at the social optimum, there is one submarket for all x_2 buyers.

Lemma 6. *Under Assumption 2, each level curve of $R(\lambda z, \lambda)$ intersects with the curve $H(\lambda z, \lambda) = 1/(x_2 - 1)$ at most once and from the left in the z-λ plane.*

Proof. Suppose that a level curve of $R(\lambda z, \lambda)$ intersects with the curve $H(\lambda z, \lambda) = 1/(x_2 - 1)$ at point $(\lambda z, \lambda)$ (with a slight abuse of notation). Then, x_2 is given by $1 + 1/H(\lambda z, \lambda)$. By

direct computation we have,

$$\left.\frac{\partial R(\lambda z,\lambda)/\partial z}{\partial R(\lambda z,\lambda)/\partial \lambda}\right|_{x_2=1+1/H(\lambda z,\lambda)} = \frac{\lambda\phi_{\mu\mu}(\lambda z,\lambda)}{z\phi_{\mu\mu}(\lambda z,\lambda)+\phi_{\mu\lambda}(\lambda z,\lambda)} = \frac{\partial\phi_{\mu}(\lambda z,\lambda)/\partial z}{\partial\phi_{\mu}(\lambda z,\lambda)/\partial\lambda},$$

where the left-hand side of the above equation denotes the slope of the level curve of $R(\lambda z,\lambda)$ at point $(\lambda z,\lambda)$ in the z-λ plane. Assumption 2 then implies Lemma 6. Finally, note that

$$\left.\frac{\partial T_2(\lambda z,\lambda)/\partial z}{\partial T_2(\lambda z,\lambda)/\partial \lambda}\right|_{x_2=1+1/H(\lambda z,\lambda)} = \left.\frac{\partial T_1(\lambda z,\lambda)/\partial z}{\partial T_1(\lambda z,\lambda)/\partial \lambda}\right|_{x_2=1+1/H(\lambda z,\lambda)} = \frac{\partial\phi_{\mu}(\lambda z,\lambda)/\partial z}{\partial\phi_{\mu}(\lambda z,\lambda)/\partial\lambda}.$$

□

A.7 Proof of Proposition 2

Before moving to the main part of the proof, we need the following simple mathematical fact. Suppose that $f(x,y)$ is an arbitrary function and is strictly concave in y. Furthermore, y as a function of x is implicitly defined by $f_2(x,y(x))=0$ (subscripts of f indicate partial derivatives). Then $f(x,y(x))$ is locally concave in x if and only if $f(x,y)$ is locally concave in (x,y).

To see this, differentiating with respect to x gives $y'(x) = -f_{12}(x,y(x))/f_{22}(x,y(x))$. Therefore, $\frac{d}{dx}f(x,y(x)) = f_1(x,y(x))$, and

$$\frac{d^2}{dx^2}f(x,y(x)) = f_{11}(x,y(x)) - \frac{f_{11}(x,y(x))^2}{f_{22}(x,y(x))},$$

which then proves our claim above.

Recall that the two-step problem of the planner is

$$\text{(A.2)} \qquad \max_{b_1}\max_{\alpha} \quad \widetilde{S}(b_1,\alpha) \equiv \alpha S\left(\frac{B_2}{\alpha},\frac{B_2+b_1}{\alpha}\right) + (1-\alpha)S\left(0,\frac{B_1-b_1}{1-\alpha}\right).$$

Thus $\overline{S}(b_1) = \max_\alpha \widetilde{S}(b_1,\alpha)$. Note that by Lemma 4, $\widetilde{S}(b_1,\alpha)$ is always strictly concave in α. We define the first term on the right-hand side as $\widetilde{S}^a(b_1,\alpha)$ and the second term as $\widetilde{S}^b(b_1,\alpha)$.

Given b_1, the optimal $\alpha^*(b_1)$ is determined by the first-order condition: $\widetilde{S}_2(b_1, \alpha) = 0$ when $\alpha \in (0, 1)$ and $\widetilde{S}_2(b_1, \alpha) \geq 0$ when $\alpha = 1$ (subscripts of S, $\widetilde{S}$, $\widetilde{S}^a$, and $\widetilde{S}^b$ indicate partial derivatives), or equivalently equation (12).

We now prove Proposition 2 by showing that if $\overline{S}'(b_1) = 0$, then $\overline{S}''(b_1) < 0$, which rules out the scenario that $\overline{S}(b_1)$ decreases first and then increases. Hence, there are only three possibilities: $\overline{S}(b_1)$ is monotonically increasing, monotonically decreasing, or increasing first and then decreasing. In either case, the first-order condition is sufficient and implies a unique solution.

Suppose that $\overline{S}'(b_1) = 0$. We first consider the case the case $\alpha^*(b_1) < 1$ or equivalently $R(B_2, b_1) < 1$, where both submarkets contain sellers. Since the optimal $\alpha^*(b_1)$ is characterized by the first-order condition $\widetilde{S}_2(b_1, \alpha) = 0$, $\overline{S}(b_1)$ is locally concave if $\widetilde{S}(b_1, \alpha)$ is locally concave in (b_1, α). In the following, we will show both $\widetilde{S}^a(b_1, \alpha)$ and $\widetilde{S}^b(b_1, \alpha)$ are locally concave in (b_1, α). By Lemma 4, both $\widetilde{S}^a(b_1, \alpha)$ and $\widetilde{S}^b(b_1, \alpha)$ are strictly concave in α. Consider $\widetilde{S}^a(b_1, \alpha)$ first. Since both R and T_1 are the same between the two submarkets, by Assumption 2 and Lemma 5, $S(\mu, \lambda)$ must be strictly locally concave at point $(\frac{B_2}{\alpha}, \frac{B_2+b_1}{\alpha})$, which then implies that $\widetilde{S}^a(b_1, \alpha)$ is locally concave due to the surplus function being constant returns to scale. Since $\widetilde{S}^b(b_1, \alpha) = (1-\alpha)m(\frac{B_1-b_1}{1-\alpha})$, again due to constant returns to scale, $\widetilde{S}^b(b_1, \alpha)$ is always concave in (b_1, α). Therefore, we have $\overline{S}''(b_1) < 0$.

Next, we consider the case $\alpha^*(b_1) = 1$, which then implies $T_1(\frac{B_2}{\alpha^*(b_1)}, \frac{B_2+b_1}{\alpha^*(b_1)}) \geq 0$ and $R(\frac{B_2}{\alpha^*(b_1)}, \frac{B_2+b_1}{\alpha^*(b_1)}) \geq 1$. The following lemma shows that in this case surplus per seller is locally strictly concave.

Lemma 7. *Under Assumption 1 and 2, if at some point (μ_0, λ_0), $R(\mu_0, \lambda_0) \geq 1$ and $T_1(\mu_0, \lambda_0) \geq 0$, then the Hessian matrix of $S(\mu, \lambda)$ at point (μ_0, λ_0) is negative definitive.*

Proof. Step 1: For any given z, $\lim_{\lambda\to\infty} T_1(\lambda z, \lambda) = 0$. To see this, note that $\lim_{\lambda\to\infty} m'(\lambda) = 0$, by equation 4 we only need to show that $\lim_{\lambda\to\infty} \phi_\lambda(\lambda z, \lambda) = 0$. Because $\phi(\mu, \lambda)$ is always concave in μ, we have $\phi(\lambda z, \lambda) > \lambda z \phi_\mu(\lambda z, \lambda)$. For $z > 0$, this implies that $\lim_{\lambda\to\infty} \phi_\mu(\lambda z, \lambda) \leq \lim_{\lambda\to\infty} \phi(\lambda z, \lambda)/\lambda z = 0$. Next, Assumption 1 implies that $\lim_{\lambda\to\infty} z\phi_\mu(\lambda z, \lambda) + \phi_\lambda(\lambda z, \lambda) = 0$,

which then implies $\lim_{\lambda\to\infty}\phi_\lambda(\lambda z,\lambda)=0$ regardless of whether $z=0$ or $z>0$.

Step 2: Since $R(\mu_0,\lambda_0)\geq 1$, $\lim_{\lambda\to\infty}R(0,\lambda)=1$, and $\lim_{\lambda\to\infty}R(\lambda,\lambda)=x_2$, there exists some z^* such that $\lim_{\lambda\to\infty}R(\lambda z^*,\lambda)=R(\mu_0,\lambda_0)$. Along the level curve λ-z where $R(\lambda z,\lambda)=R(\mu_0,\lambda_0)$, we have $T_1(\mu_0,\lambda_0)\geq 0$ and $\lim_{z\to z^*}T_1(\lambda z,\lambda)=0$. By Lemma 5, $T_1(\lambda z,\lambda)$ is first decreasing and increasing with z along the level curve of $R(\lambda z,\lambda)$. Therefore, $T_1(\lambda z,\lambda)$ crosses the x-axis at most twice, once from above and once from below. This implies that around point (μ_0,λ_0), $T_1(\lambda z,\lambda)$ is strictly decreasing in z along the level curve of $R(\lambda z,\lambda)$, which implies the Hessian matrix of $S(\mu,\lambda)$ is negative definite by Lemma 5. □

Again due to constant returns to scale, the above lemma implies that $\widetilde{S}^a(b_1,\alpha)$ is locally strictly concave at point $(b_1,1)$. Hence, if $\overline{S}'(b_1)=0$, then $\overline{S}''(b_1)<0$. □

A.8 Proof of Proposition 5

In the relaxed problem, sellers select a queue directly in a hypothetical competitive market. The expected payoff for a seller in this market is the difference between the surplus that he creates and the price of the queue. Suppose that a queue (μ,λ) solves sellers' relaxed problem. If queue (μ,λ) contains buyers of value x_k, then $T_k(\mu,\lambda)=U_k$, where $T_k(\mu,\lambda)$ is given by equations (4) and (5); if queue (μ,λ) does not contain buyers of value x_k, then $T_k(\mu,\lambda)\leq U_k$.

Note that when a seller posts a second-price auction with entry fee, t and attracts queue (μ,λ), then the expected payoff of low-type buyers from visiting this seller is $V_1=Q_1(\lambda)-(1-Q_0(\lambda))t=\phi_\mu(\lambda,\lambda)-(1-Q_0(\lambda))t$, and the expected payoff of a high-type buyer from visiting this seller is $V_2=V_1+(x_2-1)\phi_\mu(\mu,\lambda)$, which can be verified directly by considering two different scenarios: a high-type buyer faces no competition from any other buyer types, or he faces no competition from other high-type buyers but does compete with low-type buyers.[40] To summarize, the expected payoffs from a second-price auction with an entry fee

[40] Alternatively, we can use standard auction theory (see Myerson, 1981) and consider the integral of the trading probability, which in our case is $\phi_\mu(\mu,\lambda)$.

are

$$V_1 = \phi_\mu(\lambda, \lambda) - (1 - Q_0(\lambda))t,$$
$$V_2 = (x_2 - 1)\phi_\mu(\mu, \lambda) + \phi_\mu(\lambda, \lambda) - (1 - Q_0(\lambda))t.$$

An important observation is that if we set t according to equation (21) in the above equation, then $V_k = T_k(\mu, \lambda)$ for $k = 1, 2$. Thus, buyers' expected payoffs from the auction equal their marginal contribution to surplus, which implies that the solution (μ, λ) to a seller's relaxed problem is also compatible with a second-price auction with entry fee t in the sellers' constrained problem, where compatibility is defined by equation (18).

The reserve price case is similar except for one difference. When $r < x_1$, then things are exactly the same as the case with an entry fee and we have $V_k = T_k(\mu, \lambda)$, where $k = 1, 2$. When $r \in (1, x_2)$, which happens only when there are no low-type buyers ($\mu = \lambda$), then $V_1 = 0 \leq U_1$ and $V_2 = T_2(\mu, \lambda) = U_2$, in which case the queue is again compatible with a second-price auction with reserve price r in the sellers' constrained problem.

To prove efficiency, note that the sellers' relaxed problem boils down to a competitive market for buyer types. Therefore, the first welfare theorem applies and the equilibrium is efficient. Since the sellers' constrained problem is equivalent to the sellers' relaxed problem, the directed search equilibrium is also efficient. □

References

Agrawal, A., Horton, J., Lacetera, N., and Lyons, E. (2015). Digitization and the contract labor market: A research agenda. In Goldfarb, A., Greenstein, S. M., and Tucker, C. E., editors, *Economic Analysis of the Digital Economy*, pages 219–250. University of Chicago Press.

Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2012). A note on Peters and Severinov, "Competition among sellers who offer auctions instead of prices". *Journal of Economic Theory*, 147:389–392.

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

Albrecht, J. W. and Jovanovic, B. (1986). The efficiency of search under competition and monopsony. *Journal of Political Economy*, 94(6):1246–1257.

Auster, S. and Gottardi, P. (2017). Competing mechanisms in markets for lemons. mimeo.

Barro, R. J. and Romer, P. M. (1987). Ski-lift pricing, with applications to labor and other markets. *The American Economic Review*, 77(5):875–890.

Burdett, K., Shi, S., and Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, 109:1060–1085.

Butters, G. (1977). Equilibrium distributions of sales and advertising prices. *Review of Economic Studies*, 44(3):465–491.

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.

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

Eeckhout, J. and Kircher, P. (2010a). Sorting and decentralized price competition. *Econometrica*, 78:539–574.

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

Eggleston, H. G. (1958). *Convexity*. Cambridge University Press.

Einav, L., Farronato, C., Levin, J. D., and Sundaresan, N. (2017). Sales mechanisms in online markets: What happened to internet auctions? *Journal of Political Economy*, forthcoming.

Epstein, L. and Peters, M. (1999). A revelation principle for competing mechanisms. *Journal of Economic Theory*, 88(1):119–161.

Geertz, C. (1978). The bazaar economy: Information and search in peasant marketing. *The American Economic Review*, 68(2):28–32.

Guerrieri, V., Shimer, R., and Wright, R. (2010). Adverse selection in competitive search equilibrium. *Econometrica*, 78(6):1823–1862. mimeo.

Han, L. and Strange, W. C. (2014). Bidding wars for houses. *Real Estate Economics*, 42(1):1–32.

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.

Levin, D. and Smith, J. (1994). Equilibrium in auctions with entry. *The American Economic Review*, pages 585–599.

Lucking-Reiley, D. (2000). Auctions on the internet: What's being auctioned, amd how? *Journal of Industrial Economics*, 48(3):227–252.

McAfee, R. P. (1993). Mechanism design by competing sellers. *Econometrica*, 61(6):pp. 1281–1312.

Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119:468–510.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105:385–411.

Montgomery, J. D. (1991). Equilibrium wage dispersion and interindustry wage differentials. *Quarterly Journal of Economics*, 106:163–179.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Peters, M. (1997). A competitive distribution of auctions. *The Review of Economic Studies*, 64(1):pp. 97–123.

Peters, M. (2001). Common agency and the revelation principle. *Econometrica*, 69(5):1349–1372.

Peters, M. and Severinov, S. (1997). Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory*, 75:141–179.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5):996–1025.

Wolthoff, R. P. (2018). Applications and interviews: Firms' recruiting decisions in a frictional labor market. *Review of Economic Studies*, 85(2):1314–1351.