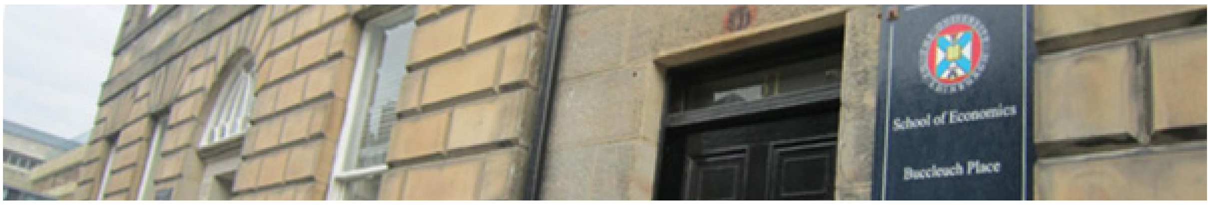

Edinburgh School of Economics

Discussion Paper Series

Number 242

Meeting Technologies and Optimal Trading Mechanisms in Competitive Search Markets

Benjamin Lester
(Federal Reserve Bank of Philadelphia)

Ludo Visschers
(The University of Edinburgh & Universidad Carlos III, Madrid)

Ronald Wolthoff
(University of Toronto)

Date
August 2014

Published by
School of Economics
University of Edinburgh
30 -31 Buccleuch Place
Edinburgh EH8 9JT
+44 (0)131 650 8361
http://edin.ac/16ja6A6

THE UNIVERSITY *of* EDINBURGH

Meeting Technologies and Optimal Trading Mechanisms in Competitive Search Markets*

Benjamin Lester†
Federal Reserve Bank of Philadelphia

Ludo Visschers‡
University of Edinburgh
& Universidad Carlos III, Madrid

Ronald Wolthoff
University of Toronto

August 2, 2014

Abstract

In a market in which sellers compete by posting mechanisms, we study how the properties of the *meeting technology* affect the mechanism that sellers select. In general, sellers have incentive to use mechanisms that are socially efficient. In our environment, sellers achieve this by posting an auction with a reserve price equal to their own valuation, along with a transfer that is paid by (or to) all buyers with whom the seller meets. However, we define a novel condition on meeting technologies, which we call "invariance," and show that the transfer is equal to zero *if and only if* the meeting technology satisfies this condition.

JEL classification: C78, D44, D83.
Keywords: search frictions, matching function, meeting technology, competing mechanisms.

*We would like to thank Philipp Kircher, Guido Menzio and Gábor Virág for helpful comments. All errors are our own.

†The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at www.philadelphiafed.org/research-and-data/publications/working-papers/.

‡Ludo Visschers gratefully acknowledges financial support from the Juan de la Cierva Grant; project grant ECO2010-20614 (Dirección general de investigación científica y técnica), and the Bank of Spain's Programa de Investigación de Excelencia.

1 Introduction

What trading mechanism should a seller use to sell a good? This is a classic question in economics and the answer, of course, depends on the details of the environment. We study this question in an environment with three key ingredients. First, there are a large number of sellers and a large number of buyers, each seller has one indivisible good, and sellers *compete* for buyers by posting the trading mechanism they will use to sell their good. Then, buyers observe the mechanisms that have been posted and choose the one that promises the maximal expected payoff, but the process by which buyers and sellers meet is *frictional*. In particular, when a buyer chooses to visit a seller, there may be uncertainty about whether he is actually able to meet with that seller, as well as the number of other buyers who also meet with that seller; we call the function that governs this process the "meeting technology." Finally, *after* buyers meet with a seller, they learn their idiosyncratic *private valuation* for the seller's good.

In this environment, we provide a complete characterization of the mechanism that sellers choose in equilibrium. More specifically, we show that sellers can do no better than a second-price auction with a reserve price equal to their own valuation and a fee (or subsidy) that is paid by (or to) each buyer that participates. We characterize this fee in closed form. Then, in the same spirit as Eeckhout & Kircher (2010), we study how this optimal mechanism depends on the properties of the meeting technology.

Our results suggest that conclusions drawn in the existing literature may depend quite heavily on the specific functional form that has typically been chosen for the meeting technology. In particular, previous studies of this basic environment (e.g., Peters (1997), Peters & Severinov (1997), Albrecht et al. (2012, 2013, 2014) and Kim & Kircher (2013)) have shared three features in common. First, the equilibrium in these papers is constrained efficient, in the sense that the solution to the planner's problem is implemented. Second, the trading mechanism that implements the constrained efficient allocation only requires transfers between the two agents that ultimately trade, so that agents who do not trade are not required to pay a fee nor are they offered a subsidy. Last, all of the papers cited above assume the same functional form for the meeting technology—what is often called the "urn-ball" meeting technology.

Taken together, the first two properties are somewhat surprising. After all, constrained efficiency places strict requirements on two distinct margins. First, from an ex ante point of view, it requires the proper allocation of buyers to sellers. Second, from an ex post point of view, it requires the proper allocation of the good once buyers arrive at a seller and learn their valuations. In order to implement an equilibrium that satisfies these two requirements, the payoffs received by buyers and sellers must inherit certain properties: to satisfy the first requirement, each buyer must receive an expected payoff exactly equal to his marginal contribution to the match surplus;[1] however, to satisfy the second requirement, payoffs must ensure that buyers truthfully report their valuation and that the good is allocated to the agent with the highest valuation. In principle, there is no reason to think that the payoffs satisfying the former requirement should coincide with the payoffs satisfying the latter.

We show that, in fact, these payoffs coincide—or, equivalently, that the meeting fee in the optimal mechanism is equal to zero—if, and only if, the meeting technology satisfies a novel condition that we call *invariance*. Loosely speaking, a meeting technology is invariant if the decision by one buyer to visit the mechanism a seller has posted does not interfere with the process by which that seller meets with other buyers. When invariance is violated, one buyer's search decision exerts an externality on the meeting prospects of the seller with other buyers, and this externality needs to be taxed (or subsidized).

The urn-ball process is one example of a meeting technology that satisfies invariance, which explains why the optimal mechanism in previous studies was so simple and why additional (rarely observed) side payments were not necessary. However, the urn-ball meeting technology also satisfies several other properties that previous studies have conjectured are important for determining the optimal trading mechanism. We close our analysis by comparing one of these properties—what Eeckhout & Kircher (2010) call "non-rivalry"—with our invariance property. We show that non-rivalry is a necessary (but not sufficient) condition for invariance.

2 Environment

Agents and Preferences. The economy is populated by a measure $\mu_B > 0$ of buyers and a measure $\mu_S > 0$ of sellers, with $\Lambda = \frac{\mu_B}{\mu_S}$.[2] Each seller possesses one, indivisible good and each buyer has unit demand for this good. All agents are risk neutral.

Sellers value their own good at y, which is common knowledge, while buyers have to visit a seller in order to learn their valuation. In particular, after a buyer meets a seller, he learns his private valuation x for that seller's good, which is drawn from a twice continuously-differentiable distribution $F(x)$ with support $[\underline{x}, \overline{x}] \subset \mathbb{R}_+$. An individual buyer's valuations are i.i.d. across sellers, as are the valuations of each buyer at an individual seller. We assume that $y \in [\underline{x}, \overline{x})$.[3]

[1] For detailed discussions of the relationship between the division of surplus and efficiency in search environments, see, e.g., Mortensen (1982), Hosios (1990) and Moen (1997).

[2] Though we focus our attention on an environment with an exogenous measure of sellers, our results do not change if we endogenize the measure of sellers through a free entry condition. We prove this formally in the Appendix.

[3] Since we do not impose much structure on $F(x)$, this is a fairly weak assumption, as the probability that a buyer's valuation x is smaller than y can be driven to zero without any loss of generality.

Mechanisms. In order to attract buyers, each seller posts and commits to a direct mechanism. A mechanism specifies an extensive form game that determines for each buyer i a probability of trade and an expected payment as a function of: (i) the total number n of buyers to meet with the seller; (ii) the valuation x_i that buyer i reports; and (iii) the valuations x_{-i} reported by the $n-1$ other buyers. Formally, a mechanism is summarized by a probability of trade $\phi(x_i, x_{-i}, n)$ and a transfer $\tau(x_i, x_{-i}, n)$ for all $n \in \mathbb{N}_1 \equiv \{1, 2, 3, \ldots\}$ and $i \in \{1, 2, ...n\}$.

Meeting Technology. After observing all posted mechanisms, each buyer chooses the mechanism at which he will attempt to trade. However, the meeting process between buyers and sellers is frictional, so that some buyers may not meet a seller and some sellers may not meet with any buyers.[4] As is often done in the literature, we restrict attention to meeting technologies that exhibit constant returns to scale. More specifically, suppose a measure σ of sellers post the same mechanism and a measure β of buyers choose this mechanism. Then, letting $\lambda = \beta/\sigma$ denote the market tightness or *queue length* at that mechanism, a seller meets exactly $n \in \mathbb{N}_0 \equiv \{0, 1, 2, \ldots\}$ buyers with probability $P_n(\lambda)$.[5] We assume that $P_n(\lambda)$ is twice continuously-differentiable. Since the number of meetings cannot exceed the number of buyers who attempt to trade at a particular mechanism, the probability distribution must satisfy $\sum_{n=0}^{\infty} nP_n(\lambda) \leq \lambda$.

We denote the probability for a buyer to meet a seller with exactly $n-1$ other buyers by $Q_n(\lambda)$. This probability is related to $P_n(\lambda)$ through the consistency requirement

$$nP_n(\lambda) = \lambda Q_n(\lambda), \tag{1}$$

for all $n \in \mathbb{N}_1$. Intuitively, given market tightness $\lambda = \beta/\sigma$ at some mechanism, the measure of buyers at those sellers who end up with precisely n buyers, $\sigma n P_n(\lambda)$, must equal the measure of buyers who face $n-1$ buyers at their seller, $\beta Q_n(\lambda)$. Finally, we define $Q_0(\lambda) = 1 - \sum_{n=1}^{\infty} Q_n(\lambda)$ as the probability that the buyer fails to meet with a seller.

Probability-Generating Function. It will prove useful to define the function

$$m(z; \lambda) \equiv E[z^n|\lambda] = \sum_{n=0}^{\infty} P_n(\lambda) z^n \tag{2}$$

[4]Note that our language follows that of Eeckhout & Kircher (2010), who make the distinction between meetings and matches. All buyers who participate in a seller's trading mechanism are said to "meet" with the seller, whereas the buyer who acquires the seller's single, indivisible good is said to "match" with the seller.

[5]We follow most of the literature by assuming that the process through which buyers and sellers meet is exogenous. Only a small number of papers derive the meeting technology as the endogenous outcome of agents' optimal decisions. See Lester et al. (2013) for an example.

for $z \in [0,1]$ with, by convention, $m(0;\lambda) = P_0(\lambda)$. This function is the probability-generating function of n and provides an alternative representation of $P_n(\lambda)$, as $P_n(\lambda) = \frac{1}{n!}\frac{\partial^n}{\partial z^n} m(0;\lambda)$. As we will see below, introducing the notation and tools of probability-generating functions allows us to derive a number of results that would be more difficult to obtain using standard techniques from the directed search literature. In particular, it will be helpful to note that, setting $z = F(x)$, $m(z;\lambda)$ can be interpreted as the probability that the maximum valuation at a seller with a queue length λ is no greater than x.[6]

Let m_z and m_λ denote the partial derivative of $m(z;\lambda)$ with respect to the first and second argument, respectively. Given the properties of a probability-generating function, the expected number of meetings per unit mass of sellers can then be expressed as

$$\sum_{n=0}^{\infty} nP_n(\lambda) = m_z(1;\lambda). \tag{3}$$

Note that $m_z(\cdot) > 0$, with $m(\cdot)$ ranging from $m(0;\lambda) = P_0(\lambda)$ to $m(1;\lambda) = 1$. In addition, we impose that $m_\lambda(z;\lambda) < 0$ and $m_{\lambda\lambda}(z;\lambda) \equiv \frac{\partial^2}{\partial\lambda^2} m(z;\lambda) > 0$ for all $z \in [0,1)$ and $\lambda \in (0,\infty)$. Again letting $z = F(x)$, these assumptions imply that a marginal increase in the queue length always reduces the probability that the maximum valuation at a seller is below $x = F^{-1}(z)$, but that the effect is smaller when the queue length is longer.[7]

3 The Planner's Problem and Market Equilibrium

In this section, we first characterize the solution to the planner's problem, and then characterize the decentralized market equilibrium.

3.1 The Planner's Problem

Consider the problem of a benevolent social planner whose objective is to maximize net social surplus, subject to the frictions of the physical environment (e.g., the meeting frictions). The planner's decision rule can be broken down into two components. First, the planner has to assign queue lengths to each seller, subject to the constraint that the sum of these queue lengths across all sellers cannot exceed the total measure of available buyers. Second, the planner has to specify a trading rule for agents to follow after buyers arrive at sellers. We solve these stages in reverse order. As the derivation is standard for each component, we keep the exposition brief.

[6] If no buyers arrive, the maximum valuation is, by convention, $-\infty$.

[7] Monotonicity of $m(\cdot)$ in λ is implied if a change in λ causes a first-order stochastic dominance shift in $P_n(\lambda)$. Convexity is also common in the literature; see, e.g., Eeckhout & Kircher (2010) for a similar assumption.

Trading Rule. Consider a seller who is visited by n buyers. Clearly, the surplus at this seller is maximized by instructing the seller to trade with the buyer who has the highest valuation x if this valuation exceeds y, and keep the good himself otherwise. Therefore, the expected net surplus at a seller with queue length λ is

$$S(\lambda) = \sum_{n=0}^{\infty} P_n(\lambda) \int_y^{\overline{x}} (x-y)\, dF^n(x) = \overline{x} - y - \int_y^{\overline{x}} m(F(x);\lambda)\, dx. \tag{4}$$

Allocation of Buyers. Now consider the allocation of buyers to sellers. At this stage, the objective of the planner is to choose queue lengths at each seller, i.e., λ_j for $j \in [0, \mu_S]$, to maximize $\int_0^{\mu_S} S(\lambda_j)\, dj$ subject to the constraint that $\int_0^{\mu_S} \lambda_j dj = \mu_B$. As $S''(\lambda) = -m_{\lambda\lambda}(F(x);\lambda) < 0$, it follows immediately that the planner maximizes total surplus by assigning equal queue lengths across all sellers, so that $\lambda_j = \Lambda = \frac{\mu_B}{\mu_S}$ for all j and total surplus equals $\int_0^{\mu_S} S(\Lambda)\, dj = \mu_S S(\Lambda)$. The following proposition summarizes.

Proposition 1. *A unique solution to the planner's problem exists. In this solution, each seller is assigned a queue length Λ. After buyers arrive and learn their valuation, the planner assigns the good to the agent with the highest valuation.*

3.2 Market Equilibrium

To characterize the relevant properties of the mechanism that sellers select in equilibrium, we proceed in two steps. First, we restrict attention to second-price auctions with a reserve price r and a meeting fee t that is paid by (or to) each buyer before learning his type. Within this restricted set of mechanisms, we define an equilibrium and characterize the profit-maximizing values of r and t. Then, we establish that sellers can do no better than this second-price auction, even with access to the set of all direct mechanisms. More specifically, we show that any direct mechanism chosen by a seller in equilibrium will be payoff-equivalent to the second-price auction with r and t chosen optimally.

Equilibrium with Second-Price Auctions and Meeting Fees. As a first step, we calculate the expected revenue of a seller who sets r and t and attracts a queue length λ, along with the expected payoff of each buyer that visits this seller.

Lemma 1. *A seller who posts a second-price auction with reserve price r and meeting fee t who* *attracts a queue λ obtains an expected payoff equal to*

$$\begin{aligned} R(r,t,\lambda) &= \overline{x} - (r-y)\,m(F(r);\lambda) - \int_r^{\overline{x}} m(F(x);\lambda)\,dx \\ &\quad - \int_r^{\overline{x}} (1-F(x))\,m_z(F(x);\lambda)\,dx + t m_z(1;\lambda). \end{aligned}$$

The expected payoff for each buyer in the seller's queue equals

$$U(r,t,\lambda) = \frac{1}{\lambda}\int_r^{\overline{x}} (1-F(x))\,m_z(F(x);\lambda)\,dx - \frac{t}{\lambda} m_z(1;\lambda). \tag{5}$$

A seller's objective is to maximize revenue, taking into account that his choice of r and t affects the queue λ that he attracts. In particular, optimal search behavior implies that buyers must be indifferent between all sellers who receive a strictly positive queue. Formally, let $\overline{U}$ denote the highest level of utility that buyers can obtain in this market, or what is often called the "market utility." Then, the equilibrium relationship between r, t, and λ is determined by $U(r,t,\lambda) = \overline{U}$.

Hence, an equilibrium in this environment is formally a distribution $G(r,t,\lambda)$ of reserve prices, meeting fees, and queue lengths across sellers, and a market utility $\overline{U}$, such that (i) given $\overline{U}$, each triple (r,t,λ) in the support of G maximizes revenue $R(r,t,\lambda)$ subject to the constraint $U(r,t,\lambda) = \overline{U}$; and (ii) aggregating queue lengths across all sellers yields the total measure of buyers, μ_B. However, one can prove that there exists a unique equilibrium in which all sellers post the same reserve price r^* and meeting fee t^*, and attract the same queue $\lambda^* = \Lambda$. Moreover, it turns out that the optimal reserve price is equal to the sellers' valuation ($r^* = y$), so that the good is allocated to the agent (one of the buyers or the seller) with the highest valuation. Hence, the decentralized equilibrium is constrained efficient, i.e., it creates the same amount of surplus as the planner's solution. The following proposition summarizes and derives t^* in closed form.

Proposition 2. *A unique equilibrium exists and it is constrained efficient. In equilibrium, the reservation price equals $r^* = y$, the meeting fee equals*

$$t^* = \frac{\Lambda}{m_z(1;\Lambda)} \int_y^{\overline{x}} \left[\frac{1}{\Lambda}(1-F(x))\,m_z(F(x);\Lambda) + m_\lambda(F(x);\Lambda)\right] dx, \tag{6}$$

and the queue length equals $\lambda^* = \Lambda$ at all sellers, with buyers receiving market utility $\overline{U} = U(r^*, t^*, \Lambda)$.

Intuitively, the first term in brackets in (6) is the buyer's private expected payoff from participating in an auction with reserve price $r = y$ when the queue length is Λ, which is evident from (5). The second term in brackets is the (negative value of) this buyer's marginal contribution to the social surplus generated by the auction, which is evident from (4). Hence, t^* is simply the difference between the two, adjusted to account for the fact that not every buyer actually meets the seller and pays this fee.

Equilibrium with General Mechanisms. Proposition 2 describes the equilibrium outcome when sellers have access to a restricted set of mechanisms and shows that the equilibrium coincides with the solution to the planner's problem. In Proposition 3 below, we establish that sellers would not benefit from having access to a larger set of mechanisms; both the expected payoffs and the allocations will be identical in any other equilibrium. The reasoning is simple: since buyers' equilibrium payoffs are equal to the market utility $\overline{U}$, an individual seller is the residual claimant on any additional surplus he creates. Hence, he has incentive to select an efficient mechanism in order to maximize this surplus. The mechanism that we consider—a second-price auction with a reserve price r and a meeting fee t—is sufficiently flexible that it allows the seller to maximize the "size of the pie" without placing restrictions on how the pie is divided. Therefore, access to additional mechanisms is largely irrelevant.[8]

Proposition 3. *A mechanism* $\{\phi(x_i, x_{-i}, n), \tau(x_i, x_{-i}, n)\}$ *is an equilibrium if, and only if, it is payoff-equivalent to the equilibrium characterized in Proposition 2.*

4 Meeting Technologies and the Optimal Mechanism

In the previous section, we characterized the equilibrium mechanism that arises in a fairly general environment. An important feature of the optimal mechanism is a meeting fee, t^*, that allows sellers to manipulate expected revenue without distorting the efficient allocation. Somewhat surprisingly, this meeting fee does not play a role in the optimal mechanism that has been identified in the previous literature. Rather, these papers find that the optimal mechanism is simply an auction with reserve price $r^* = y$.

We now reconcile our results with the existing literature by showing that previous studies had focused on a special class of meeting functions. In particular, we introduce a novel property of meeting technologies that we call *invariance* and show that this property is necessary and sufficient for $t^* = 0$. We use several examples to illustrate the relationship between invariance and the necessity of meeting fees. We then conclude by reviewing another property of meeting technologies that has been identified in the literature and discuss its relationship to invariance.

[8] One may also wonder whether there exists an alternative equilibrium in which sellers use a mechanism that only requires the buyer who trades to pay a transfer, and leaves all other buyers with a payoff of zero. In the Appendix, we show that the answer is "no."

4.1 Invariance

We now introduce a property of meeting technologies that we call "invariance," which requires that the probability of a seller meeting with n members of any subset of buyers in the queue is unaffected by the presence (or lack thereof) of buyers in the queue who are outside of this subset, for all $n \in \mathbb{N}_0$. To formalize this concept, suppose that each buyer is assigned a certain label with probability $\gamma \in [0,1]$, independent of the realizations of this label for other buyers. Then, the probability that a seller meets exactly n buyers with this label equals[9]

$$\widetilde{P}_n(\lambda,\gamma) \equiv \sum_{N=n}^{\infty} P_N(\lambda) \begin{pmatrix} N \\ n \end{pmatrix} \gamma^n (1-\gamma)^{N-n}.$$

Definition 1. *A meeting technology is* invariant *if, and only if,* $\widetilde{P}_n(\lambda,\gamma) = P_n(\gamma\lambda)$ *for all* $\gamma \in (0,1)$*,* $\lambda \in (0,\infty)$*, and all* $n \in \mathbb{N}_0$.

To understand the intuition, consider a sub-market in which the ratio of buyers to sellers is λ, and suppose that a fraction γ of buyers are red and the remainder are green.[10] Invariance, then, says that the ratio of red buyers to sellers must be a sufficient statistic for the distribution that governs the number of red buyers who arrive at a seller. In particular, changing the measure of green buyers should not change the probability that n red buyers arrive at a seller, for any $n \in \mathbb{N}_0$.

In other words, a marginal buyer choosing to visit the mechanism a seller has posted does not interfere with the process by which that seller meets with *other* buyers. When the invariance condition is violated, the seller either wants to tax or subsidize buyers for choosing his mechanism, depending on whether they exert a negative or positive externality on his prospects for meeting other buyers, respectively. The proposition below formalizes this argument: we establish that meeting fees are set to zero in equilibrium if, and only if, the meeting technology is invariant, and we derive a simple condition for verifying invariance for arbitrary technologies.

Proposition 4. *Consider a meeting technology* $P_n(\lambda)$*. The following statements are equivalent:*

1. *Meeting fees are not used in equilibrium, i.e.,* $t^* = 0$ *for any distribution* $F(x)$ *with support* $[\underline{x}, \overline{x}] \subset [0,\infty)$*,* $y \in [\underline{x}, \overline{x})$ *and* $\Lambda \in (0,\infty)$*;*

2. *The meeting technology is invariant;*

[9] This transformation, introduced by Rényi (1956), is known as "binomial thinning." See Harremoës et al. (2010).

[10] Note that labeling buyers "red" and "green" is simply a convenient device to split the queue into subsets. Indeed, different labels can be arbitrarily assigned to ex ante homogeneous buyers, which will be the case in our analysis, or they can be assigned to buyers that are ex ante heterogeneous, as we discuss in Section 4.2. For the purpose of defining invariance, it does not matter. It does, however, slightly change the intuition in some of our explanations below. In what follows, we find it easiest to convey the intuition by treating a buyer's color as being realized ex ante, but the discussion could easily be adapted to treat the alternative case as well.

3. *The meeting technology can be written as*

$$P_n(\lambda) = \frac{(-\lambda)^n M^{(n)}(\lambda)}{n!}, \tag{7}$$

for some function $M : [0,\infty) \to [0,1]$ *that satisfies* $m(z;\lambda) = M(\lambda(1-z))$ *and preserves all of our assumptions on* $P_n(\lambda)$ *and* $m(z;\lambda)$, *with* $M^{(n)}(\cdot)$ *denoting the* n^{th}*-order derivative of* M.[11]

4.2 Discussion

Examples of Meeting Technologies. In order to facilitate a discussion of the invariance condition introduced above, it will be helpful to introduce a few examples of meeting technologies.

1. *Urn-Ball.* Popular in the directed search literature, this specification arises when a measure β of buyers randomize evenly across a measure σ of sellers who have posted the same mechanism.[12] As a result, the number of buyers that visit a particular seller is determined by a Poisson distribution with mean equal to $\lambda = \beta/\sigma$. That is, $P_n(\lambda) = e^{-\lambda}\frac{\lambda^n}{n!}$.

2. *Geometric.* The geometric meeting technology specifies that $P_n(\lambda) = \frac{\lambda^n}{(1+\lambda)^{n+1}}$. Again, letting $\lambda = \beta/\sigma$, this meeting technology can be interpreted as the outcome of a process in which a mass β of buyers and a mass σ of sellers are randomly positioned on a circle, and buyers walk to the nearest seller to their right.

3. *Bilateral.* This specification, often used in the literature that uses search-theoretic models to study monetary theory, has the following interpretation.[13] Suppose there is a mass β and σ of buyers and sellers, respectively, and all agents are randomly matched in pairs. Then, given $\lambda = \beta/\sigma$, the probability distribution over the number of *buyers* that a seller meets is: $P_0(\lambda) = \frac{1}{1+\lambda}$, $P_1(\lambda) = \frac{\lambda}{1+\lambda}$, and $P_n(\lambda) = 0$ for $n \in \{2,3,\ldots\}$.

4. *Pairwise Urn-Ball.* If a measure β of buyers are first grouped in pairs and then allocated across a measure σ of sellers according to the urn-ball technology described above then, letting $\lambda = \beta/\sigma$, the meeting technology can be described by $P_n(\lambda) = 0$ for $n \in \{1,3,5,\ldots\}$ and $P_n(\lambda) = e^{-\lambda/2}\frac{(\lambda/2)^{n/2}}{(n/2)!}$ for $n \in \{0,2,4,\ldots\}$.

There are several things to note from these examples. First, the urn-ball technology satisfies the invariance condition: when all buyers randomize across a group of sellers, the presence of more

[11] For the sake of completeness, we describe exactly which properties of M are necessary to preserve our assumptions on P_n and m in the Appendix.

[12] This specification was first used by Butters (1977) and Hall (1977). See Burdett et al. (2001) for an explicit derivation of the micro-foundations of this meeting technology in a directed search model.

[13] See, e.g., Kiyotaki & Wright (1993) and Trejos & Wright (1995). This technology has also been used more recently to study trading dynamics in over-the-counter financial markets; see, e.g., Duffie et al. (2007).

green buyers has no effect on the number of red buyers to arrive at a particular seller. Since the existing literature has focused primarily on the case of urn-ball matching, it should then come as no surprise that a simple second-price auction with $r^* = y$ was identified as the optimal mechanism.

Second, since the geometric technology satisfies (7) for $M(\lambda) = \frac{1}{1+\lambda}$, this meeting technology also satisfies invariance: since a red buyer meets with the first seller to his right, it is irrelevant (by construction) how many green buyers he passes on his way.[14] This confirms that the set of invariant meeting technologies is not limited to derivatives of the urn-ball process, but accepts other reasonable meeting processes as well.

Third, note that the bilateral meeting technology is *not* invariant: since sellers only meet with one randomly selected buyer under this meeting technology, an increase in the measure of green buyers will decrease the probability that a seller meets one red buyer (and increase the probability he meets zero red buyers). Hence, in a competitive environment with this meeting technology, sellers would optimally employ a meeting fee to tax this congestion externality appropriately.

Lastly, note that the pairwise urn-ball technology also violates invariance. Intuitively, when the measure of green buyers is insignificant, most red buyers are matched in pairs with other red buyers and it is highly likely that a seller is visited by an even number of red buyers. As the measure of green buyers gets large, however, it becomes increasingly likely that red buyers are paired with green buyers, which implies a greater probability that a seller is visited by an odd number of red buyers. Hence, this process is not invariant. However, in contrast to the bilateral meeting technology, buyers exert a positive externality on one another under the pairwise urn-ball meeting process, and thus sellers optimally offer a meeting *subsidy* in equilibrium (i.e., $t^* < 0$).

Non-Rivalry. We close our analysis by comparing the invariance condition to another property of meeting technologies that has been identified in the literature: what has often been called "non-rivalry." In addition to distinguishing our results from the existing literature, this comparison also sheds light on other environments in which the invariance condition is likely to be an important determinant of the optimal mechanism for price (or wage) determination.

In words, Eeckhout & Kircher (2010) describe a (purely) non-rival meeting technology as one in which "the meeting probability for a buyer is not affected by the presence of other buyers in the market."[15] That is, $1 - Q_0(\lambda)$ is independent of λ. Using the definition of $Q_0(\lambda)$, it is easy to verify that this condition is equivalent to the requirement that there exists a $\gamma \in [0, 1]$, such that $\sum_{n=0}^{\infty} nP_n(\lambda) = \gamma\lambda$ for all λ.[16]

[14] More generally, the family of negative-binomial distributions satisfies equation (7) for $M(\lambda) = \left(\frac{\theta}{\theta+\lambda}\right)^{\theta}$ and $\theta \in \mathbb{N}_1$. This family of distributions is therefore invariant as well.

[15] For a recent example of a paper that cites the importance of non-rivalry, see Albrecht et al. (2014).

[16] Note that $\sum_{n=1}^{\infty} Q_n(\lambda) = \frac{1}{\lambda} \sum_{n=0}^{\infty} nP_n(\lambda)$.

Clearly the urn-ball technology described earlier in the text is non-rival because each buyer meets a seller with probability one. The bilateral meeting technology, however, violates non-rivalry. Since each seller can meet at most one buyer, an increase in the ratio of buyers to sellers implies that each buyer is more likely to be crowded out. As a result, the probability that each buyer meets a seller, $1 - Q_0(\lambda) = \frac{1-e^{-\lambda}}{\lambda}$, is decreasing in λ. Given these observations, a natural question is how non-rivalry relates to invariance, which we address in the following lemma.[17]

Lemma 2. *Invariance implies non-rivalry, but non-rivalry does not imply invariance.*

To understand the first statement in Lemma 2, it is helpful to recall our previous example, where buyers are either red or green. If the meeting technology does not satisfy non-rivalry, then increasing the measure of green buyers, while holding constant the measure of sellers and red buyers, will decrease the probability that each red buyer meets a seller. This implies that increasing the measure of green buyers will decrease the probability that a seller meets with $n \geq 1$ red buyers, which violates the invariance condition. The second statement in the lemma follows from considering the pairwise urn-ball technology. This technology is non-rival since each buyer meets a seller with probability one, regardless of the ratio of buyers to sellers. However, as we explained above, it is not invariant.

Intuitively, non-rivalry is a statement about whether a buyer's decision to visit a sub-market (or mechanism) affects the probability that another buyer will meet a seller in that sub-market, while invariance is a statement about whether an individual buyer's decision to visit a sub-market affects the *distribution* of buyers at each seller in that sub-market. Given this distinction, we conclude that whether or not a meeting technology is invariant is likely to be a crucial determinant of the optimal trading mechanism in any environment where the payoff of a seller (or firm) depends on the distribution of buyers (or workers) that arrive. In this paper, we provided perhaps the simplest environment where the distribution of buyers is relevant: when buyers draw independent valuations after choosing a seller, a key object of interest for the seller is the maximum valuation among all buyers—an order statistic—which clearly depends on whether there are two or three or four buyers.

However, by this logic, invariance is also likely to be an important property of the meeting technology in models of directed search in which buyers can attempt to meet multiple sellers, as in Albrecht et al. (2006) and Galenianos & Kircher (2009). In these models, ex post heterogeneity arises because some buyers have received two price quotes and others have received only one; a seller's trading probability is thus $1 - \widetilde{P}_0(\lambda, \gamma)$, where γ is the fraction of buyers that don't trade with a different seller. As such, our results suggest that meeting fees will likely be needed to price

[17] In the Appendix, we compare invariance to a third property of meeting technologies that the urn-ball technology also satisfies, which we call "independence." Put succinctly, this property requires that the arrival of a buyer at a particular seller does not change the distribution of other buyers at that seller. We show that independence is neither sufficient nor neccessary for invariance.

this externality if the meeting technology is not invariant.[18]

Importantly, the role of the invariance condition in determining the optimal trading mechanism is not limited to environments with ex post heterogeneity, but is also likely to be important in environments with ex ante heterogeneity as well. For example, in a directed search model with ex ante heterogeneous buyers, if invariance is violated—e.g., if one buyer's search decision affects the distribution of other buyers that the seller meets—then our results suggest that standard methods of price determination may need to be supplemented with fees or subsidies.[19] We leave a detailed investigation of the role of invariance in models with ex ante heterogeneity for future research.

5 Conclusion

We consider an environment in which sellers compete by posting mechanisms and buyers direct their search toward the mechanism offering the maximal expected payoff. We characterize the trading mechanism that sellers choose in equilibrium, given a fairly general class of meeting technologies. We show that, in general, sellers can do no better than posting an auction with no reserve price, but with a meeting fee (or subsidy) that is paid by (or to) all buyers who participate in the auction. Then, we completely characterize a subset of meeting technologies for which the meeting fee is set to zero, and show that the meeting technology that previous studies have used was contained in this subset. We call the meeting technologies in this subset *invariant.*

Appendix

Proofs

Proof of Proposition 1. The proof of this result is provided in the main text.

[18] This is consistent with the results in this literature that wage posting (without meeting fees) decentralizes the planner's solution when the meeting technology is urn-ball (Kircher, 2009), but not when firms are constrained in the number of applications that they can process (Galenianos & Kircher, 2009; Wolthoff, 2012).

[19] Indeed, with ex ante heterogeneity, the meeting technology affects both the optimal mechanism *and* whether different types of buyers pool or separate in equilibrium: an urn-ball technology leads to perfect pooling, while a bilateral technology leads to perfect separation. Peters & Severinov (1997), Virág (2010), Albrecht et al. (2012, 2014) and Peters (2013) study the urn-ball case, while Eeckhout & Kircher (2010) analyze both urn-ball and bilateral meeting technologies. Shi (2001) derives related results in a labor market setting with bilateral meetings, while Shimer (2005) and Peters (2010) study a similar problem with the urn-ball meeting technology. Given our results, a natural conjecture is that the pooling equilibrium without meeting fees extends to all invariant meeting technologies, while meeting technologies with congestion externalities give rise to an equilibrium with meeting fees and at least some separation of types.

Proof of Lemma 1. From the description of the auction in the main text, it follows immediately that a buyer's expected payoff equals

$$U(r,t,\lambda)=\sum_{n=1}^{\infty}Q_n(\lambda)\left[\int_r^{\overline{x}}\int_{\underline{x}}^{x}(x-\max\{\widetilde{x},r\})\,dF^{n-1}(\widetilde{x})\,dF(x)-t\right].$$

The inner integral simplifies to $\int_r^x F^{n-1}(\widetilde{x})\,d\widetilde{x}$ after integration by parts. Changing the order of integration and substituting equations (1) and (2) then yields equation (5).

To calculate the payoff of a seller, consider the expected valuation $V(r,\lambda)$ of the agent that holds the good at the end of the period, given a queue length λ and reserve price r. As the maximum valuation among n buyers is distributed according to $F^n(x)$, $V(r,\lambda)$ equals

$$\begin{aligned}V(r,\lambda)&=P_0(\lambda)\,y+\sum_{n=1}^{\infty}P_n(\lambda)\left[F^n(r)\,y+\int_r^{\overline{x}}x\,dF^n(x)\right]\\&=\overline{x}-(r-y)\,m(F(r);\lambda)-\int_r^{\overline{x}}m(F(x);\lambda)\,dx,\end{aligned}$$

where the second line follows after integration by parts and substitution of (2). As the payoff of a seller must satisfy $R(r,t,\lambda)=V(r,\lambda)-\lambda U(r,t,\lambda)$, the desired expression follows.

Proof of Proposition 2. The Lagrangian of the seller's maximization problem equals

$$L(r,t,\lambda,\zeta)=R(r,t,\lambda)+\zeta\left(U(r,t,\lambda)-\overline{U}\right),$$

where ζ denotes the multiplier on the market utility constraint. The first-order condition with respect to t equals

$$\left(1-\frac{\zeta^*}{\lambda^*}\right)m_z(1;\lambda^*)=0.$$

As $m_z(1;\lambda)>0$, this implies that $\zeta^*=\lambda^*$ in any equilibrium. Substituting this into the first-order condition with respect to r gives $(y-r^*)\,m_z(F(r^*);\lambda^*)\,F'(r^*)=0$. Hence, $r^*=y$ in any equilibrium. In combination with these results, the first-order condition with respect to λ subsequently implies that the optimal meeting fee must satisfy

$$t^*=\frac{\lambda^*}{m_z(1;\lambda^*)}\int_y^{\overline{x}}\left[\frac{1}{\lambda^*}(1-F(x))\,m_z(F(x);\lambda^*)+m_\lambda(F(x);\lambda^*)\right]dx.$$

In combination with the first-order condition with respect to ζ, this implies that the optimal queue length λ^* is determined by

$$-\int_{y}^{\overline{x}} m_{\lambda}\left(F\left(x\right);\lambda^*\right)dx = \overline{U}. \tag{8}$$

Since $m\left(z;\lambda\right)$ is convex in λ, a unique solution exists for any $\overline{U}$ small enough. In equilibrium, the market utility has to be such that $\lambda^* = \Lambda$, or otherwise the market clearing constraint $\int_0^{\mu_S} \lambda^* dj = \mu_B$ is violated. It immediately follows that $\overline{U} = -\int_y^{\overline{x}} m_{\lambda}\left(F\left(x\right);\Lambda\right)dx$ and that the equilibrium is efficient.

Proof of Proposition 3. This proof closely resembles the corresponding one in Lester et al. (2013) and suppresses the arguments of ϕ and τ to keep notation concise. Suppose there exists an equilibrium in which one or more sellers post a particular mechanism $\{\phi_0, \tau_0\}$ and attract a queue length λ_0, which yields them a payoff $\mathbf{R}\left(\phi_0, \tau_0, \lambda_0\right)$, yields buyers a market utility $\overline{U} = \mathbf{U}\left(\phi_0, \tau_0, \lambda_0\right)$, and creates a surplus $\mathbf{S}\left(\phi_0, \lambda_0\right) = \mathbf{R}\left(\phi_0, \tau_0, \lambda_0\right) + \lambda_0 \mathbf{U}\left(\phi_0, \tau_0, \lambda_0\right) - y$.[20] A deviant seller who posts a mechanism $\{\phi, \tau\}$ and attracts a queue λ, determined by the market utility condition $\mathbf{U}\left(\phi, \tau, \lambda\right) = \overline{U}$, obtains a payoff

$$\mathbf{R}\left(\phi, \tau, \lambda\right) = \mathbf{S}\left(\phi, \lambda\right) - \lambda\overline{U} + y.$$

This payoff is maximized when the deviant chooses trading probabilities ϕ which correspond with the planner's solution ϕ^*, and τ such that

$$\frac{\partial}{\partial\lambda}\mathbf{S}\left(\phi^*, \lambda\right) \equiv \frac{d}{d\lambda}S\left(\lambda\right) = \overline{U}. \tag{9}$$

Note that $\{\phi_0, \tau_0\}$ can be part of an equilibrium if, and only if, it is a solution to the deviant's maximization problem, i.e., $\phi_0 = \phi^*$ and λ_0 solves (9). Strict concavity of $S\left(\lambda\right)$ implies that a unique solution exists to (9) for each level of market utility. Hence, all sellers must attract the same queue length and this queue length must equal $\lambda = \Lambda$, as any other value is inconsistent with the aggregate buyer-seller ratio. It is easy to verify that $\frac{d}{d\lambda}S\left(\lambda\right)\big|_{\lambda=\Lambda} = U\left(r^*, t^*, \Lambda\right)$, which implies that $\{\phi_0, \tau_0\}$ is an equilibrium strategy if and only if it is pay-off equivalent to the equilibrium characterized in Proposition 2.

Proof of Proposition 4. First, we show that invariance implies a zero meeting fee (i.e., statement 2 implies statement 1). Note that $\widetilde{P}_n\left(\lambda, \gamma\right)$ is a compound distribution describing the sum of $N \sim P_n\left(\lambda\right)$ independent Bernoulli variables. The probability-generating function of a Bernoulli

[20] Naturally, the amount of surplus created by a seller does not depend on the payments.

variable is $1-\gamma+\gamma z$, so the probability-generating function of $\widetilde{P}_n(\lambda,\gamma)$ equals $m(1-\gamma+\gamma z;\lambda)$. The invariance condition is then equivalent to

$$m(1-\gamma+\gamma z;\lambda)=m(z;\gamma\lambda). \tag{10}$$

Taking the derivative of this condition with respect to γ and evaluating the result in $\gamma=1$ yields

$$-(1-z)\,m_z(z;\lambda)=\lambda m_\lambda(z;\lambda), \tag{11}$$

for all z and λ. If we substitute this expression into (6) for $z=F(x)$ and $\lambda=\Lambda$, then $t^*=0$ follows immediately.

Second, we show that a meeting technology that does not need a meeting fee for all $F(x)$, y and Λ must satisfy equation (7) (i.e., statement 1 implies statement 3). Note that since y affects the lower bound of integration in (6), expression (11) is not only a sufficient condition, but it is also a necessary condition. This expression is a first-order partial differential equation with solution

$$m(z;\lambda)=M(\lambda(1-z)), \tag{12}$$

for arbitrary differentiable M, although our assumptions about $P_n(\lambda)$ and $m(z;\lambda)$ require a number of additional restrictions on M, which we characterize in the Appendix. Equation (7) directly follows from (12).

To complete the proof, it suffices—by the transitivity of implication—to show that a meeting technology which satisfies equation (7) is invariant (i.e. statement 3 implies statement 2). This is immediate using probability-generating functions, as (12) satisfies (10).

Proof of Lemma 2. The necessity of non-rivalry follows from combining equation (3) and equation (12), which reveals that invariant meeting technologies satisfy $\sum_{n=0}^{\infty} nP_n(\lambda)=-M'(0)\,\lambda$. We establish in the online Appendix that $M'(0)\in[-1,0]$. Hence, meeting technologies that are invariant are also non-rival.

To prove that non-rivalry is not sufficient, we show that the pairwise urn-ball technology is a counterexample. For this technology, $1-Q_0(\lambda)=1$ is independent of λ. Hence, pairwise urn-ball is non-rival. At the same time, pairwise urn-ball violates invariance, because $\widetilde{P}_1(\lambda,\gamma)>0$ for all $\gamma>0$, while $P_1(\gamma\lambda)=0$. Hence, non-rivalry does not imply invariance.

Free Entry

Although Proposition 2 establishes that the equilibrium is efficient for any arbitrary ratio of buyers to sellers, a traditional concern in much of the search literature is whether efficiency is also achieved when this ratio is determined endogenously. To address this issue, suppose that sellers can freely enter the market by paying a cost k, as in standard search models (see, e.g., Pissarides, 1985). The planner will then choose the measure of sellers μ_S to maximize net social surplus,

$$\max_{\mu_S \in (0,\infty)} \mu_S \left[S\left(\frac{\mu_B}{\mu_S}\right) - k \right], \tag{13}$$

while the equilibrium measure of sellers follows from the indifference condition

$$R\left(r^*, t^*, \frac{\mu_B}{\mu_S}\right) - k = y. \tag{14}$$

The following lemma states that the market tightness achieved in equilibrium indeed coincides with the solution to the planner's problem.

Lemma A1. *The market equilibrium with free entry is constrained efficient.*

This extends some of the results of of Albrecht et al. (2012, 2014) to arbitrary meeting technologies. Clearly, our results regarding the importance of meeting fees and (invariance of) the meeting technology are robust to endogenizing the measure of sellers, as they hold for arbitrary values of the buyer-seller ratio.

Proof. The proof resembles Lester et al. (2013) and Albrecht et al. (2014). The equilibrium measure of sellers is determined by the free entry condition (14). Since $S(\lambda) = R(r^*, t^*, \lambda) + \lambda U(r^*, t^*, \lambda) - y$, this condition is equivalent to

$$S\left(\frac{\mu_B}{\mu_S}\right) - k = \frac{\mu_B}{\mu_S} U\left(r^*, t^*, \frac{\mu_B}{\mu_S}\right). \tag{15}$$

The efficient measure of sellers follows from the first-order condition of net social surplus, described in (13), with respect to μ_S. After simplification, this yields

$$S\left(\frac{\mu_B}{\mu_S}\right) - k = \frac{\mu_B}{\mu_S} S'\left(\frac{\mu_B}{\mu_S}\right). \tag{16}$$

Since $S'\left(\frac{\mu_B}{\mu_S}\right) = U\left(r^*, t^*, \frac{\mu_B}{\mu_S}\right)$, it follows that the solution to (15) solves (16). ■

Necessity of Meeting Fees

Proposition 3 establishes that all equilibria must be payoff-equivalent to the one with second-price auctions and meeting fees. In this proposition, payoff-equivalence concerns the expected payoffs and not the realized payoffs, as it is straightforward to change the latter while keeping the former the same, e.g. by changing the auction format. This observation may raise the question whether one can construct an equilibrium mechanism in which the meeting fees/subsidies are replaced by an extra payment by/to the trading buyer, such that buyers who do not trade receive a zero payoff, even when the meeting technology is not invariant. Lemma A2 establishes that this is not feasible. Before presenting the formal argue, we sketch the intuition.

We know that any mechanism that is chosen in equilibrium has to implement the planner's solution, and it has to be payoff equivalent to the mechanism we describe. This puts a lot of structure on the mechanism: it pins down the probability that a buyer is awarded the good, and it pins down the expected transfer he pays. Suppose we divide the game into three stages: at stage 1, buyers choose a seller; at stage 2, buyers arrive and are potentially asked to pay a meeting fee (or receive a subsidy); and at stage 3, buyers learn their type and report it, whereupon the good is allocated and additional transfers occur. Our mechanism specifies that all buyers pay t at stage 2, and then a second price auction occurs at stage 3. The question can then be phrased: is there an alternative mechanism with *no* transfers at stage 2, and only a transfer between the seller and the buyer who receives the good at stage 3?

Loosely speaking, the reason the answer is "no" is that transfers at stage 3 have to respect both an individual rationality (or participation) constraint and an incentive compatibility (or truth-telling) constraint. If the seller did not charge t at stage 2, he would have to charge higher prices at stage 3; otherwise his revenue would be different, which violates payoff-equivalence. However, he cannot raise prices on only buyers who report high valuations, as this would violate incentive compatibility. Therefore, he would have to raise prices on *all* buyers, but this would violate the participation constraint of buyers with valuations very close to y. Hence, the seller *must* extract rents at stage 2; he *cannot* derive a pricing scheme that extracts the same rents but still respects the buyers' IR and IC constraints.

Lemma A2. *All equilibrium mechanisms must include a transfer, (in expectation) equal to t^*, which is paid by/to each of the buyers who arrive at a seller before they learn their valuations*

Proof. As a first step, consider a mechanism $\{\phi_n(x), \tau_n(x)\}$ that assigns a probability of trade $\phi_n(x)$ in exchange for a transfer $\tau_n(x)$ to an agent who reports being of type x when there are n buyers participating in the mechanism.[21] Once a buyer has learned his type, x, he reports a

[21] This representation of a mechanism constitutes a slight abuse of notation whereby, e.g., $\phi_n(x_1) = \int \ldots \int \phi(x_1, x_2, \ldots, x_n, n) dF(x_2) \ldots dF(x_n)$, where $\phi(x_1, \ldots, x_n, n)$ is defined in the main text.

valuation x' which maximizes

$$\phi_n(x')\,x - \tau_n(x')\,. \tag{17}$$

The incentive compatibility (or truth-telling) constraint then requires

$$\phi_n'(x)\,x - \tau_n'(x) = 0. \tag{18}$$

Hence,

$$\begin{aligned}\tau_n(x) &= \int_y^x \tau_n'(\widetilde{x})\,d\widetilde{x} + C_0 = \int_y^x \widetilde{x}d\phi_n(\widetilde{x}) + C_0 \\ &= \phi_n(x)\,x - y\phi_n(y) - \int_y^x \phi_n(\widetilde{x})\,d\widetilde{x} + C_0 \\ &= \phi_n(x)\,x - \int_y^x \phi_n(\widetilde{x})\,d\widetilde{x} + C_1, \end{aligned} \tag{19}$$

for constants C_0 and C_1, where the second equality follows from integration by parts. Combining (17) and (19) yields

$$\int_y^x \phi_n(\widetilde{x})\,d\widetilde{x} - C_1. \tag{20}$$

Hence, in *any* incentive compatible mechanism, the expected payoff to a buyer with valuation x is completely determined by the probability of trade, $\phi_n(x)$, and a constant C_1.

Next, as a second step, we will show that both the probabilities of trade and the constant C_1 are uniquely determined in any equilibrium of our game. To see this, first note that any equilibrium mechanism must be constrained efficient (Proposition 2 in the text), which implies that the good must be allocated to the agent who values it most. Hence,

$$\phi_n(x) = \begin{cases} F^{n-1}(x) & \text{if } x \geq y \\ 0 & \text{otherwise.} \end{cases} \tag{21}$$

Substituting (21) into (20) and taking expectations over n and x, we find that the ex ante expected utility of a buyer is

$$U = \sum_{n=1}^{\infty} Q_n(\lambda)\left[\int_{\underline{x}}^{\overline{x}} \int_y^x F^{n-1}(\widetilde{x})\,d\widetilde{x}dF(x) - C_1\right]. \tag{22}$$

By Proposition 3, $U = \bar{U}(r^*, t^*, \Lambda)$ in any equilibrium, which implies that C_1 is also uniquely determined. Indeed, in our mechanism, $C_1 = t$, the meeting fee.

Finally, as a last step, consider an equilibrium of our game in which $C_1 = t > 0$. Moreover, suppose we decompose $C_1 = C_1^a + C_1^b$, where C_1^a denotes any fees that are paid before buyers learn their type, and C_1^b are transfers made after reporting x. The individual rationality (or participation) constraint of a buyer with valuation x is given by

$$\int_y^x \phi_n(\widetilde{x})\, d\widetilde{x} - C_1^b \geq 0. \tag{23}$$

Since this constraint must hold for all $x \in [y, \overline{x}]$, it follows that $C_1^b = 0$ in every equilibrium. ■

Restrictions on M

Our assumptions about $P_n(\lambda)$ and $m(z;\lambda)$ impose a number of restrictions on $M(\cdot)$.[22]

(i) m being a probability-generating function requires that M is an analytic function.[23]

(ii) $\frac{1}{n!}\frac{\partial^n}{\partial z^n} m(0;\lambda) \in [0,1]$ for all n and all λ requires that $\frac{(-\lambda)^n}{n!} M^{(n)}(\lambda) \in [0,1]$ for all n and λ;

(iii) $m(1;\lambda) = 1$ for all λ requires that $M(0) = 1$;

(iv) $m_z(1;\lambda) \in [0,\lambda]$ for all λ requires that $M'(0) \in [-1,0]$;

(v) $m_\lambda(z;\lambda) < 0$ for all z and all λ requires that $M'(\lambda) < 0$ for all λ;

(vi) $m_{\lambda\lambda}(z;\lambda) > 0$ for all z and all λ requires that $M''(\lambda) > 0$ for all λ.

Jointly, these restrictions completely characterize the set of feasible $M(\cdot)$. However, a tighter characterization is possible. For example, restriction (ii) can be tightened as follows

(ii′) M satisfies $\frac{(-\lambda)^n}{n!} M^{(n)}(\lambda) \in (0,1)$ for all n and λ.

Proof. We provide a proof by contradiction. Suppose that there exist a $\widehat{n} \in \mathbb{N}_0$ and a $\widehat{\lambda} > 0$ such that $\frac{(-\widehat{\lambda})^{\widehat{n}}}{\widehat{n}!} M^{(\widehat{n})}\left(\widehat{\lambda}\right) = 0$. As $\widehat{\lambda} > 0$, this implies $M^{(\widehat{n})}\left(\widehat{\lambda}\right) = 0$. Because $M^{(\widehat{n})}(\lambda)$ is continuously differentiable but cannot cross zero, it must be the case that its first derivative is zero in $\lambda = \widehat{\lambda}$ as well, i.e. $M^{(\widehat{n}+1)}\left(\widehat{\lambda}\right) = 0$. By induction, it then follows that all higher derivatives must equal zero in this point, $M^{(n)}\left(\widehat{\lambda}\right) = 0$ for all $n \in \{\widehat{n}, \widehat{n}+1, \widehat{n}+2, \ldots\}$.

With all of its derivatives being zero in $\widehat{\lambda}$, the analytic function $M^{(\widehat{n})}(\lambda)$ must be zero in a neighborhood around $\widehat{\lambda}$. Again by induction, we therefore obtain that $M^{(n)}(\lambda) = 0$ for all

[22] To keep notation concise, we write "all λ" instead of "all $\lambda \in (0,\infty)$", "all n instead of "all $n \in \mathbb{N}_0$", and "all z" instead of "all $z \in [0,1)$".

[23] Probability-generating functions are analytic functions (Sachkov, 1997).

$n \in \{\widehat{n}, \widehat{n}+1, \widehat{n}+2, \ldots\}$ and all λ. In that case,

$$\sum_{i=0}^{\widehat{n}-1} P_i(\lambda) = \sum_{i=0}^{\widehat{n}-1} \frac{(-\lambda)^i}{i!} M^{(i)}(\lambda) = 1,$$

which is a differential equation with solution $M(\lambda) = 1 + \sum_{i=1}^{\widehat{n}-1} c_i \lambda^i$, for some coefficients c_i. Because $M'(\lambda) < 0$ rules out the possibility that $M(\lambda)$ is a constant, one obtains that $\lim_{\lambda\to\infty} |M(\lambda)| = +\infty$.

This contradicts the requirement that $M(\lambda) \in [0,1]$ for all λ. Hence, it must be true that $\frac{(-\lambda)^n}{n!} M^{(n)}(\lambda) > 0$ for all n and all λ. This immediately implies that $\frac{(-\lambda)^n}{n!} M^{(n)}(\lambda) < 1$, for all n and all λ, which completes the proof. ■

Note that restriction (ii′) implies restriction (v) and (vi), which are therefore redundant. Hence, the set of feasible M can be characterized by restrictions (i), (ii′), (iii), and (iv).

Independence

Another well-known property of the urn-ball meeting technology is that $P_n(\lambda)$ does not only represent the distribution of the total number of buyers at each seller, but also the distribution of the number of competitors a buyer faces when arriving at a seller. In other words, a buyer who meets with a seller has no effect on the distribution (and thus the expectation) of the number of other buyers at the same seller. We will say that a meeting technology that satisfies this property exhibits "independence".

Formally, independence means $Q_n(\lambda) = (1 - Q_0(\lambda)) P_{n-1}(\lambda)$ for all λ and $n \in \mathbb{N}_1$. This property can be shown to be satisfied if and only if the meeting technology is of the following form:[24]

$$P_n(\lambda) = e^{-(1-Q_0(\lambda))\lambda} \frac{[(1-Q_0(\lambda))\lambda]^n}{n!}. \tag{24}$$

The following lemma establishes that independence is neither a necessary nor a sufficient condition for invariance.

Lemma 3. *Invariance does not imply independence and independence does not imply invariance.*

To see why invariance does not imply independence, consider the geometric technology. As we established in the main text, this meeting technology is invariant. However, this technology is not independent, as the fact that an individual buyer meets with a seller changes the probability

[24] By the consistency requirement and induction, the condition implies that $P_n(\lambda) = \frac{1}{n!}\lambda^n [1 - Q_0(\lambda)]^n P_0(\lambda)$. Since $P_n(\lambda)$ is a probability distribution, we must have $1 = \sum_0^\infty \frac{1}{n!}\lambda^n [1 - Q_0(\lambda)]^n P_0(\lambda) = e^{\lambda[1-Q_0(\lambda)]} P_0(\lambda)$. Solving the second equation for $P_0(\lambda)$ and substituting the solution into the first equation yields (24). Conversely, equation (24) implies the independence condition immediately.

distribution over the number of other buyers who meet with the seller. In particular, when a buyer meets with a seller, he learns that the seller has only met with buyers thus far, which increases the expected number of other buyers that this seller will ultimately meet; that is, the conditional mean of the queue length is greater than the unconditional mean.

Finally, to understand why independence does not imply invariance, consider the following variation on the urn-ball technology, in which a longer queue length reduces each buyers' chances of meeting a seller. That is,[25]

$$P_n(\lambda) = e^{-\Phi(\lambda)\lambda} \frac{[\Phi(\lambda)\lambda]^n}{n!}, \tag{25}$$

where $\Phi(\lambda) : [0, \infty) \to [0, 1]$ satisfies $\Phi(0) = 1$ and $\Phi'(\lambda) < 0$. To understand this meeting technology, imagine that buyers and sellers in each sub-market begin on separate islands. The measure σ of sellers each send a boat to transport the measure β of buyers, so that each boat carries $\lambda = \beta/\sigma$ buyers to the sellers' island. However, each boat sinks with probability $1 - \Phi(\lambda)$, so that heavier boats are more likely to sink. Then, the buyers that arrive safely at the sellers' island randomly select a seller, as in the urn-ball specification.

Since the probability of each boat's safe passage depends on the ratio of buyers to sellers in the sub-market, this technology does not satisfy the requirements of non-rivalry, and hence is not invariant. However, once buyers arrive, the meeting process ensues according to the standard urn-ball technology, so that the arrival of an individual buyer has no effect on the distribution of other buyers to arrive. Hence, this meeting process satisfies independence.[26]

References

Albrecht, J. W., Gautier, P. A., & Vroman, S. B. (2006). Equilibrium directed search with multiple applications. *Review of Economic Studies*, 73, 869–891.

Albrecht, J. W., Gautier, P. A., & 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., & Vroman, S. B. (2013). Directed search in the housing market. mimeo.

Albrecht, J. W., Gautier, P. A., & Vroman, S. B. (2014). Efficient entry with competing auctions. *American Economic Review*, forthcoming.

[25] See Kaas (2010) for a related example of this class of meeting technologies.

[26] This highlights the difference between non-rivalry and independence. Using our analogy from above, non-rivalry is primarily a requirement on the probability that "a boat arrives safely" (i.e., it must be independent of λ), while independence is primarily a condition on the meeting process that occurs "after boats arrive on shore."

Burdett, K., Shi, S., & 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.

Duffie, D., Gârleanu, N., & Pedersen, L. H. (2007). Valuation in over-the-counter markets. *Review of Financial Studies*, 20(6), 1865–1900.

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

Galenianos, M. & Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 114, 445–471.

Hall, R. (1977). *Microeconomic Foundations of Macroeconomics*, chapter An Aspect of the Economic Role of Unemployment. Macmillan, London.

Harremoës, P., Johnson, O., & Kontoyiannis, I. (2010). Thinning, entropy, and the law of thin numbers. *IEEE Transactions on Information Theory*, 56(9), 4228–4244.

Hosios, A. K. (1990). On the efficiency of matching and related models of search and unemployment. *Review of Economic Studies*, 57, 279–298.

Kaas, L. (2010). Variable search intensity in an economy with coordination unemployment. *B.E. Journal of Macroeconomics*, 10(1), 1–33.

Kim, K. & Kircher, P. (2013). Efficient competition through cheap talk: Competing auctions and competitive search without ex-ante commitment. CEPR Discussion Paper DP9785.

Kircher, P. (2009). Efficiency of simultaneous search. *Journal of Political Economy*, 117, 861–913.

Kiyotaki, N. & Wright, R. (1993). A search-theoretic approach to monetary economics. *American Economic Review*, 83(1), 63–77.

Lester, B., Visschers, L., & Wolthoff, R. P. (2013). Competing with asking prices. *Federal Reserve Bank of Philadelphia Working Paper 13-7*.

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

Mortensen, D. T. (1982). Efficiency of mating, racing, and related games. *American Economic Review*, 72, 968–979.

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

Peters, M. (2010). Noncontractible heterogeneity in directed search. *Econometrica*, 78(4), 1173–1200.

Peters, M. (2013). Competing mechanisms. In N. Vulcan, A. Roth, & Z. Neeman (Eds.), *The Handbook of Market Design*. Oxford University Press.

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

Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies and real wages. *American Economic Review*, 75, 676–690.

Rényi, A. (1956). A characterization of Poisson processes. *Magyar Tud. Akad. Mat. Kutaló Int. Közl.*, 1, 519–527.

Sachkov, V. (1997). *Probabilistic Methods in Discrete Mathematics*, volume 56 of *Encyclopedia of Mathematics and its Applications*. New York: Cambridge University Press.

Shi, S. (2001). Frictional assignment I: Efficiency. *Journal of Economic Theory*, 98(2), 232–260.

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

Trejos, A. & Wright, R. (1995). Search, bargaining, money, and prices. *Journal of Political Economy*, 103(1), 118–141.

Virág, G. (2010). Competing auctions: Finite markets and convergence. *Theoretical Economics*, 5, 241–274.

Wolthoff, R. P. (2012). Applications and interviews: A structural analysis of two-sided simultaneous search. FEEM Working Paper 637.