

Search Frictions, Competing Mechanisms and Optimal Market Segmentation*

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

March 7, 2017

Abstract

In a market in which sellers compete for heterogeneous buyers by posting mechanisms, we analyze how the properties of the meeting technology affect the allocation of buyers to sellers. We show that a separate submarket for each type of buyer is the efficient outcome if and only if meetings are bilateral. In contrast, a single market with all agents is optimal if and only if the meeting technology satisfies a novel condition, which we call “joint concavity.” Both outcomes can be decentralized by sellers posting auctions combined with a fee that is paid by (or to) all buyers with whom the seller meets. Finally, we compare joint concavity to two other properties of meeting technologies, invariance and non-rivalry, and explain the differences.

JEL classification: C78, D44, D83.

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

*The idea for this paper grew out of earlier fruitful collaboration with James Albrecht, Ben Lester, Ludo Visschers and Susan Vroman. We thank the editor, the associate editor, three referees, Philipp Kircher, Guido Menzio, Michael Peters, Gábor Virág, and various seminar and conference participants for valuable comments. Ronald Wolthoff gratefully acknowledges financial support from the Connaught Fund at the University of Toronto.

[†]Tongji University. E-mail: xiaoming@tongji.edu.cn.

[‡]Vrije Universiteit Amsterdam and Tinbergen Institute. E-mail: p.a.gautier@vu.nl.

[§]University of Toronto. E-mail: ronald.p.wolthoff@gmail.com.

1 Introduction

Search theory offers a framework to study how buyers and sellers interact in a market without a Walrasian auctioneer. While past work in this area has already greatly improved our understanding of the trading process in such markets, one aspect remains relatively unexplored: the process that governs the initial meeting between agents, i.e. the *meeting technology*. After all, the vast majority of papers in this literature simply assumes either that each meeting is one-to-one (*bilateral meetings*, as in Moen, 1997) or that each meeting is n -to-1, where n follows a Poisson distribution (*urn-ball meetings*, as in Burdett et al., 2001).

Although rarely motivated, the choice between these two specifications often matters for outcomes. In particular, this is the case if agents who differ in characteristics affecting match surplus can direct their search to the contract (or mechanism) that maximizes their expected payoff. In such an environment, papers using urn-ball meetings tend to find that each equilibrium contract attracts multiple types of searchers (see e.g. Peters, 1997; Shi, 2002; Shimer, 2005; Albrecht et al., 2014; Auster and Gottardi, 2016). In contrast, papers using bilateral meetings tend to find that equilibrium exhibits perfect separation of the different types of searchers across different contracts (see e.g. Moen, 1997; Shi, 2001; Eeckhout and Kircher, 2010a; Guerrieri et al., 2010; Menzio and Shi, 2011).

This difference in market segmentation, highlighted in a unified environment by Eeckhout and Kircher (2010b), is important for a number of reasons. For example, it affects the mechanisms that we observe in equilibrium: separation into homogenous submarkets can generally be implemented with posted prices, while any form of pooling requires menus of prices and selection rules (if types are observable) or auctions (if they are not) to give priority to the highest types. Moreover, the difference is relevant from a technical point of view: as first shown by Shi (2009), perfect separation of types is analytically and computationally convenient in a dynamic context, because it makes the equilibrium block-recursive.

Moving beyond bilateral and urn-ball is important, because they are not necessarily the most adequate description of meeting technologies in real-life markets; in many cases, it might be more reasonable to suppose that a seller can meet and learn the type of multiple but not all buyers who try to match with him. This raises the question how robust the outcomes associated with these two particular technologies are. In other words, how informative are easily observable equilibrium outcomes, like trading mechanisms and market segmentation, on the frictions in the meeting process? In this paper, we answer this question for a static environment in which homogeneous sellers post mechanisms to attract buyers who differ in their private valuation, as in Peters and Severinov (1997), Eeckhout and Kircher (2010b) and Albrecht et al. (2014). To do so, we allow for arbitrary meeting technologies and derive necessary and sufficient conditions on the meeting technology under which it is optimal to have (i) perfect separation, i.e. a separate market for each type of buyer, or (ii) perfect pooling, i.e. a single market with all agents, for

all distributions of buyer types.

After describing the environment in detail in section 2, we start our analysis in section 3 by considering the trade-off of a social planner between the desire to spread high-type buyers as much as possible and the risk of them being crowded out by low-type buyers. Throughout, we extensively make use of a one-to-one transformation of the probability-generating function of a meeting technology. This transformation—the probability that a seller meets at least one buyer with an arbitrary label or characteristic—is developed in a companion paper Cai et al. (2016) and greatly simplifies the analysis.

Our first result concerns the optimality of perfect separation under any type distribution. We find that bilateral meetings are not only sufficient for this outcome, but also necessary. That is, if one moves away from bilateral meetings by allowing a seller to meet multiple buyers (potentially with arbitrary small probability), then there exist distributions of buyer valuations for which perfect separation is no longer efficient. Intuitively, separation does not exploit the efficiency gains that arise from sellers ranking multiple buyers: with homogenous submarkets, any meetings beyond the first are meaningless since they always yield the seller a clone of the buyer that he has already met. We establish that the set of distributions for which this necessity condition holds includes all distributions with a support that contains an interval.

Although the necessity of bilateral meetings for perfect separation is a new result in the literature, it is perhaps not very surprising. Most of our attention therefore goes out to the optimality of a single market. We show that this is the efficient outcome under any type distribution if and only if the meeting technology satisfies a novel condition which we call “joint concavity.” Joint concavity is satisfied by the urn-ball meeting technology, which explains why pooling is the efficient outcome in e.g. Peters and Severinov (1997), but we also describe a number of other meeting technologies that exhibit this property. Loosely speaking, joint concavity guarantees that social surplus can be increased by merging any two submarkets, irrespective of their composition. When this property is violated, there naturally exist type distributions for which pooling is not optimal. In contrast, we establish that joint concavity is not necessary for pooling if a type distribution has discrete support.

In the second half of section 3, we describe—based on Cai et al. (2016)—how both the separating and the pooling outcome can be decentralized by each seller posting a second-price auction, combined with a meeting fee to be paid by (or to) each buyer meeting him. Intuitively, in a large market, sellers take buyers’ equilibrium payoffs as given, making sellers the residual claimant on any extra surplus that they create and providing them with an incentive to post efficient mechanisms. Auctions guarantee that the good is allocated efficiently despite the presence of private information, while the meeting fees price any positive or negative externalities in the meeting process, providing all agents with a payoff equal to their social contribution. Interestingly, this mechanism reduces to price posting when meetings are bilateral, since a buyer meeting a seller will then trade with certainty by paying the meeting fee and bidding the

seller’s valuation.

We conclude our analysis by comparing our findings to existing results in section 4. In particular, we discuss how joint concavity relates to two other properties of meeting technologies described in the literature: (i) invariance as introduced by Lester et al. (2015), and (ii) non-rivalry as introduced by Eeckhout and Kircher (2010b). We show that invariance is a sufficient (but not a necessary) condition for joint concavity, while non-rivalry is a necessary (but not a sufficient) condition, and we explain why this is the case. Finally, the appendix contains all proofs, while additional results and discussion can be found in the online appendix.

Related Literature. While most of the literature assumes either bilateral or urn-ball meetings, a small number of papers explores arbitrary meeting technologies, as we do here.¹ The first to do so were Eeckhout and Kircher (2010b), whose formalization of meeting technologies has been used by subsequent work in this area. We analyze a similar environment, but contribute relative to their paper by identifying conditions for pooling or separating equilibria that are both necessary and sufficient.

Lester et al. (2015) use arbitrary meeting technologies in a model in which buyers’ valuations are realized after they arrive at sellers. Their paper has therefore no implications for market segmentation and focusses instead on the question when meeting fees arise as part of the equilibrium mechanism. They show that this is the case if and only if the meeting technology is not invariant. We discuss the relation between joint concavity and invariance in section 4.

Cai et al. (2016) make a methodological contribution by introducing the alternative representation of meeting technologies that we also use here. They show that this representation keeps the analysis of competing mechanisms with arbitrary meeting technologies tractable, even when heterogeneity is two-sided. They further establish efficiency of the market equilibrium—a result we employ in section 3.2—and derive conditions on the meeting technology such that sellers’ beliefs about the buyers that they will attract are uniquely determined by the market utility condition.² However, they do not characterize the optimal degree of market segmentation; note that a single market never arises with two-sided heterogeneity, as documented before by e.g. Shi (2001), Shi (2002), Shimer (2005), Eeckhout and Kircher (2010a) and Albrecht et al. (2014) for bilateral or urn-ball meetings.

Finally, Cai (2016) uses the alternative representation to extend the efficiency result of Hosios (1990) to an environment in which workers can meet multiple firms, e.g., through on-the-job search. However, he considers random search and wage bargaining, as in Gautier et al. (2010), and is therefore silent on market segmentation as well.

¹Alternatively, some papers propose specific (classes of) meeting technologies that are neither bilateral nor urn-ball (e.g. Fraja and Sákovics, 2001; Lester and Wolthoff, 2014; Wolthoff, 2017; Lester et al., 2017).

²In this paper, we use the standard assumption that sellers expect the most favorable queue in case of a multiplicity. See section 3.2 for a detailed discussion.

2 Environment

Agents and Preferences. We consider a static goods market with incomplete information, sharing key features with McAfee (1993) and Peters and Severinov (1997). The market is populated by a measure 1 of sellers, indexed by $j \in [0, 1]$, and a measure $\Lambda > 0$ of buyers. Both types of agents are risk-neutral. 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. A buyer’s valuation is private information and is an independent draw from a distribution $F(x)$ with $0 \leq x \leq 1$.³ We impose no additional structure on $F(x)$, although we will sometimes pretend that buyers have either a low valuation x_1 or a high valuation x_2 when describing the intuition behind our results.⁴ Buyers observe their valuation before making any decisions. An agent’s payoff is the sum of (i) his monetary transfers and (ii) his valuation if he possesses the good at the end of the period (and zero otherwise).

Search. In order to attract buyers, each seller posts and commits to a direct mechanism (hereafter: “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 that 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.⁵

All identical mechanisms are treated symmetrically by buyers and are therefore said to form a *submarket*. After observing all submarkets, each buyer chooses the one in which he wishes to attempt to match.⁶ As standard in the literature (see e.g. Shimer, 2005), we capture the anonymity of the large market by assuming that: i) sellers can condition their strategies on the actions of buyers but not on their identities, and ii) identical buyers must use identical mixed strategies in equilibrium. Consequently, agents’ search decisions can be summarized by three endogenous variables for each submarket: the measure s of sellers, the measure b of buyers, and the distribution $G(x)$ of valuations among these buyers.

Meeting Technology. Within a submarket, meetings between buyers and sellers are governed by a frictional process, the *meeting technology*, which we model as in Eeckhout and Kircher (2010b). That is, the probability $P_n(\lambda)$ that a seller meets $n \in \mathbb{N}_0 = \{0, 1, 2, \dots\}$ buyers only

³The assumption that all buyers have a (weakly) higher valuation than the seller is standard as well as innocuous. Buyers with lower valuations would simply never trade.

⁴Neither proposition 1 nor proposition 2 below is driven by the requirement that they should hold for all $F(x)$, i.e. they remain the same under the *weaker* requirement that they should hold for all $F(x)$ with only two points of support.

⁵See the online appendix for a precise definition. In line with most of the literature (e.g. Peters, 1997; Eeckhout and Kircher, 2010b; Lester et al., 2015, 2017), we abstract from mechanisms that condition on other mechanisms in the market. Epstein and Peters (1999) and Peters (2001) provide a detailed discussion.

⁶The assumption that a buyer can meet only one seller per period is standard in the directed search literature. See Albrecht et al. (2006), Galenianos and Kircher (2009), Kircher (2009), Wolthoff (2017) and Gautier and Holzner (2017) for papers that relax this assumption.

depends on the *queue length* $\lambda = b/s$ in the submarket, is twice-continuously differentiable, and satisfies $\sum_{n=0}^{\infty} nP_n(\lambda) \leq \lambda$, since the number of meetings cannot exceed the number of buyers in the submarket. Further, since valuations are private information, a seller who cannot meet all the buyers that try to visit him lacks the relevant knowledge to choose the most promising ones and meets a randomly selected subset instead.⁷ Formally, conditional on a seller meeting n buyers, the types of these buyers are n independent draws from the submarket’s valuation distribution $G(x)$. That is, if a fraction $\mu/\lambda \in [0, 1]$ of the buyers in the submarket are labeled “blue,” then the probability for a seller to meet i blue buyers and $n - i$ other buyers equals

$$P_n(\lambda) \binom{n}{i} \left(\frac{\mu}{\lambda}\right)^i \left(1 - \frac{\mu}{\lambda}\right)^{n-i}.$$

Alternative Representation. Cai et al. (2016) show that the analysis of arbitrary meeting technologies is often greatly simplified by using an alternative representation of $P_n(\lambda)$. This alternative representation is the probability $\phi(\mu, \lambda)$ that a seller with a queue μ of blue buyers and a queue $\lambda - \mu$ of other buyers meets at least one blue buyer. We follow this approach here. Given the assumption regarding the type-independent meetings, $\phi(\mu, \lambda)$ equals

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

To simplify notation, we will often omit the arguments of ϕ and use subscripts to indicate its partial derivatives.⁸

Examples of Meeting Technologies. For future reference, it will be useful to formally define a few examples of meeting technologies that satisfy all our assumptions.

1. *Urn-Ball.* First explored by Butters (1977) and Hall (1977), this technology specifies that—within a submarket—each buyer is randomly allocated to one of the sellers. As a result, the number of buyers that meet a particular seller follows a Poisson distribution with mean equal to the queue λ . That is, $P_n(\lambda) = e^{-\lambda} \frac{\lambda^n}{n!}$ and $\phi(\mu, \lambda) = 1 - e^{-\mu}$.
2. *Bilateral.* With this technology, each seller meets at most one buyer, i.e. $P_0(\lambda) + P_1(\lambda) = 1$ or $\phi(\mu, \lambda) = P_1(\lambda) \frac{\mu}{\lambda}$, with $P_0(\lambda)$ strictly convex. A potential micro-foundation consists of randomly pairing agents and keeping only pairs that consist of one buyer and one seller, yielding $P_1(\lambda) = \frac{\lambda}{1+\lambda}$.⁹

⁷The allocation of buyers to submarkets and the selection of a buyer for trade will of course not be random but depend on buyers’ types.

⁸See Cai et al. (2016) for a proof that the relation between $\phi(\mu, \lambda)$ and $P_n(\lambda)$ is one-to-one.

⁹This micro-foundation can be found in the money search literature (see e.g. Kiyotaki and Wright, 1993). Some papers in the labor search literature provide an alternative, consisting of an urn-ball process augmented with the constraint that each seller can only contact one random buyer among the ones that wish to meet him,

3. *Pairwise Urn-Ball.* This technology, described by Lester et al. (2015), is a variation on the urn-ball technology. Buyers first form pairs, after which each pair is randomly assigned to a seller in the submarket. That is, $P_n(\lambda) = 0$ for $n \in \{1, 3, 5, \dots\}$ and $P_n(\lambda) = e^{-\lambda/2} \frac{(\lambda/2)^{n/2}}{(n/2)!}$ for $n \in \{0, 2, 4, \dots\}$, which implies $\phi(\mu, \lambda) = 1 - e^{-\mu(1-\frac{1}{2}\frac{\mu}{\lambda})}$.
4. *Multi-Platform.* This technology consists of two platforms or rounds. In the first round, all b buyers and a fraction $0 < \alpha < 1$ of the s sellers in a submarket attempt to meet according to the random-pairing bilateral technology described above. The $\frac{b}{b+\alpha s} b = \frac{\lambda}{\lambda+\alpha} b$ buyers who fail to meet a seller then participate in the second round, in which they meet the remaining $(1 - \alpha) s$ sellers according to an urn-ball process. That is,

$$P_n(\lambda) = \begin{cases} \alpha \frac{\alpha}{\lambda+\alpha} + (1 - \alpha) e^{-\xi} & \text{if } n = 0 \\ \alpha \frac{\lambda}{\lambda+\alpha} + (1 - \alpha) \xi e^{-\xi} & \text{if } n = 1 \\ (1 - \alpha) \frac{\xi^n e^{-\xi}}{n!} & \text{if } n \in \{2, 3, \dots\}, \end{cases}$$

where $\xi = \frac{\lambda^2}{(1-\alpha)(\lambda+\alpha)}$ is the queue length in the second round. This yields $\phi(\mu, \lambda) = \alpha \frac{\mu}{\lambda+\alpha} + (1 - \alpha) \left(1 - e^{-\frac{\lambda\mu}{(1-\alpha)(\lambda+\alpha)}}\right)$.¹⁰

3 Planner's Problem and Market Equilibrium

We start by analyzing the problem of a planner whose objective is to maximize social surplus, subject to the frictions generated by the meeting technology. To keep the exposition as simple as possible, we initially assume that the planner knows buyers' valuations, allowing him to provide different types of buyers with different instructions. This naturally raises the question whether the planner's solution would be different without that knowledge. We establish that this is not case using the result of Cai et al. (2016) that—for any meeting technology—the solution can be decentralized with a particular incentive-compatible mechanism, which provides a way for a planner who does not know buyers' valuations to implement the same solution.

such that $P_1(\lambda) = 1 - e^{-\lambda}$ (see e.g. Albrecht et al., 2006; Galenianos and Kircher, 2009; Gautier and Wolthoff, 2009; Gautier et al., 2016).

¹⁰While this technology may seem more involved than the other examples, the two-round structure actually resembles the meeting process in various real-life markets: (i) buyers who cannot find a product at the local bazaar may subsequently submit a bid at an online auction site; (ii) workers who have trouble finding a job are often put in touch with firms by a public employment agency; and (iii) singles who fail to meet someone in a bar may subscribe to a dating website. Admittedly, the analogy is not perfect because terms of trade may differ across platforms in reality, which is ruled out here by the definition of a submarket.

3.1 Planner’s Problem

The problem of a planner who knows buyers’ valuations consists of two parts. First, the planner has to allocate buyers and sellers to submarkets, creating a queue length and a distribution of buyer types at each seller. Second, the planner has to specify how trade will take place after buyers arrive at sellers. We solve these stages in reverse order.

Trading Rule. Once a number of buyers $n \in \mathbb{N}_1 \equiv \{1, 2, 3, \dots\}$ has arrived at a seller, surplus is clearly maximized by allocating the good to the buyer with the highest valuation. Cai et al. (2016) show that the expected surplus generated by this trading rule can be written as the integral of ϕ .

Lemma 1 (Cai et al., 2016). *The surplus created by a seller with a queue λ of buyers whose types are distributed according to the distribution $G(x)$ equals*

$$S(\lambda, G) = \int_0^1 \phi(\lambda(1 - G(x)), \lambda) dx. \quad (2)$$

Proof. See Cai et al. (2016) or online appendix B.1. □

Allocation of Buyers. Now consider the allocation of buyers to sellers. For each seller $j \in [0, 1]$, the planner chooses—with a slight abuse of notation—a queue length $\lambda(j)$ and a distribution of buyer types $G(j, x)$ to maximize total surplus $\int_0^1 S(\lambda(j), G(j, x)) dj$. Of course, the planner cannot allocate more buyers of a certain type than are available. Formally, $\int_0^1 \lambda(j) \nu(j, B) dj \leq \Lambda \nu_F(B)$ for any Borel-measurable set B , where ν_F is the measure associated with F and $\nu(j, \cdot)$ is the measure associated with $G(j, \cdot)$.

Perfect Separation. We first establish that bilateral meetings are a necessary and sufficient condition for the optimality of perfect separation. Recall that bilateral meeting technologies are defined by the property that $P_0(\lambda) + P_1(\lambda) = 1$ for any λ , with $P_1(\lambda)$ strictly concave.

Proposition 1. *Bilateral meetings are a necessary and sufficient condition for the planner to create a separate submarket for each type of buyer under any $F(x)$ and Λ .*

Proof. See appendix A.1. □

As mentioned before, sufficiency of bilateral meetings for perfect separation is a well-known result in the literature: a separate submarket for each active buyer type avoids the high degree of crowding-out that arises if high-type and low-type buyers visit the same submarket and sellers meet one of both at random. Note that a planner may keep the lowest types out of the market altogether if a marginal seller can generate more surplus in a different submarket.

Necessity is however—to the best of our knowledge—a new result. To understand the intuition, suppose that a seller can meet two or more buyers with positive probability. With perfect

separation, any meetings beyond the first are irrelevant—as a seller will always meet a clone of the first buyer—and the gain in surplus relative to a bilateral technology is zero. Letting one high-type and one low-type buyer swap submarket, however, can increase total surplus. To see this, note that the swap increases surplus in the low-type submarket by $x_2 - x_1$ as long as the selected buyer meets a seller. In contrast, surplus in the high-type submarket may decrease by $x_2 - x_1$, but only if the selected buyer meets a seller *in a bilateral fashion*. Which effect is larger depends on the queue lengths, but the positive effect necessarily dominates if the queue lengths are similar enough across the two submarkets. Necessity therefore follows from distributions in which buyers’ types are sufficiently close. The following corollary builds on that insight to establish the necessity of bilateral meetings for distributions of which the support contains an interval.

Corollary 1. *If the support of a type distribution F contains an interval, bilateral meetings are necessary for a separate submarket for each type of buyer to be efficient for any Λ .*

Proof. See appendix A.2. □

In contrast, if types are wide apart, a meeting technology may give rise to perfect separation even though it is not strictly bilateral (see Eeckhout and Kircher, 2010b, for a detailed discussion of this case).

Perfect Pooling. To state our main result regarding the optimality of a single market, we define a novel property of meeting technologies, which we call “joint concavity.”

Definition 1. *A meeting technology exhibits joint concavity if and only if $\phi(\mu, \lambda)$ is concave in (μ, λ) , i.e.*

$$\phi_{\mu\mu}\phi_{\lambda\lambda} \geq \phi_{\mu\lambda}^2, \tag{3}$$

for all $0 \leq \mu \leq \lambda < \infty$.¹¹

The following proposition then establishes that joint concavity is closely related to the optimality of a single market.

Proposition 2. *Joint concavity is a necessary and sufficient condition for the planner to send all agents to the same market under any $F(x)$ and Λ .*

Proof. See appendix A.3. □

The intuition for this result is straightforward.¹² Joint concavity is sufficient because it implies that by merging two submarkets, the probability $\phi(\lambda(1 - G(x)), \lambda)$ of a seller meeting at

¹¹Condition (3) is necessary and sufficient for concavity, since $\phi_{\mu\mu} \leq 0$ for all meeting technologies.

¹²Here the advantage of using ϕ becomes apparent; the equivalent condition in terms of P_n , which we derive in the online appendix, is far less simple and intuitive.

least one buyer with a valuation higher than x will increase for all x . Since the surplus created by a submarket is the integral of ϕ over x , pooling all agents into one submarket is optimal. In contrast, necessity follows from distributions that have two points of support, with one of both being zero, i.e. $x_1 = 0$ and $x_2 > 0$. For these distributions, only the number of meetings with high-type buyers matters for surplus and (2) reduces to $x_2\phi(\lambda(1 - G(x_1)), \lambda)$. Hence, if joint concavity of ϕ fails, there exist measures of low-type and high-type buyers such that either partially or completely separating them into multiple submarkets is optimal.

A natural question is now whether—analogueous to corollary 1—the necessity of joint concavity can be established for a given F . It is not clear that this is the case. Instead, we provide a related (negative) result in the following proposition.

Proposition 3. *For any F with discrete support, there exists a meeting technology which does not satisfy joint concavity but for which perfect pooling is optimal for any Λ .*

Proof. See appendix A.4. □

A particularly intuitive example of this result is the case in which F is degenerate. Pooling is then optimal if and only if $P_0(\lambda)$ is convex, which holds even for a bilateral technology. For F with $N \in \{2, 3, \dots\}$ points of support, the proof of the proposition exploits the fact that calculation of surplus from pooling only requires evaluation of ϕ in N points. This makes it possible to perturb an urn-ball technology, for which pooling is known to be optimal, to construct a new meeting technology that creates the same amount of surplus as urn-ball under pooling, but less surplus than urn-ball for any other allocation. We show that this new meeting technology violates joint concavity, which yields the desired result.¹³

3.2 Market Equilibrium

Equilibrium Definition. Next, we define the market equilibrium.¹⁴ To do so, let $R(m, \lambda, G)$ denote the expected payoff of a seller who posts a mechanism m and attracts a queue of buyers (λ, G) . Further, let $U(x, m, \lambda, G)$ denote the expected payoff of a buyer with valuation x who visits this seller. Each seller aims to maximize his revenue R , but must take into account that his queue (λ, G) is endogenously determined and depends on the mechanism that he posts. In line with the literature, we follow the market utility approach. That is, given a tuple $(m(j), \lambda(j), G(j, \cdot))$ for each seller $j \in [0, 1]$, let $\bar{U}(x)$ denote the highest expected payoff that a buyer with valuation x can obtain, i.e. $\bar{U}(x) = \max_{j \in [0, 1]} U(x; m(j), \lambda(j), G(j, \cdot))$. A seller posting a mechanism m then expects a queue satisfying

$$U(x, m, \lambda, G) \leq \bar{U}(x), \text{ with equality for each } x \text{ in the support of } G. \quad (4)$$

¹³In the online appendix, we present a numerical example and discuss why this procedure fails when the support of F contains an interval.

¹⁴We provide a brief discussion here. See the online appendix for additional details.

For many meeting technologies, (4) uniquely determines the seller’s queue. In case of multiplicity, we follow McAfee (1993), Eeckhout and Kircher (2010b) and Auster and Gottardi (2016) by assuming that sellers are optimistic and expect the solution that maximizes their revenue.¹⁵ Finally, sellers expect a queue that gives them a non-positive payoff if no solution to (4) exists. An equilibrium can then be defined as follows.

Definition 2. *A directed search equilibrium is a mechanism $m(j)$ and a queue $(\lambda(j), G(j, \cdot))$ for each seller $j \in [0, 1]$, and a market utility $\bar{U}(x)$ for each type of buyer x , such that ...*

1. *each $(m(j), \lambda(j), G(j, \cdot))$ maximizes $R(m, \lambda, G)$ subject to equation (4);*
2. *aggregating queues across sellers does not exceed the total measure of buyers of each type;*
3. *incentive compatibility is satisfied, so buyers report their valuations truthfully.*

Decentralization. The following proposition, which follows from Cai et al. (2016), states the main result regarding the market equilibrium: for any meeting technology, the planner’s solution can be decentralized by having each seller post a second-price auction combined with a meeting fee, to be paid by each buyer meeting the seller.¹⁶

Proposition 4 (Cai et al., 2016). *For any meeting technology, the planner’s solution $\{\lambda(j), G(j, x)\}$ can be decentralized as a directed search equilibrium in which seller $j \in [0, 1]$ posts a second-price auction and a meeting fee equal to*

$$\tau(j) = - \frac{\int_0^1 \phi_\lambda(\lambda(j)(1 - G(j, x)), \lambda(j)) dx}{\phi_\mu(0, \lambda(j))}. \quad (5)$$

Proof. See Cai et al. (2016) or online appendix B.2. □

The intuition for this result is similar to the intuition for efficiency in many other directed search models. Since sellers take buyers’ equilibrium payoffs as given, they are the residual claimant on any surplus that they create. This provides them with an incentive to post mechanisms that decentralize the planner’s solution, which requires efficiency along two margins: (i) the allocation of buyers to sellers, and (ii) the allocation of the good given a queue of buyers.

The second-price auction fulfills the second requirement and provides each buyer with a payoff equal to the extra surplus that he creates when he has the highest valuation. To satisfy the first requirement however, each buyer must receive an expected payoff exactly equal to his marginal contribution to social surplus, which includes the externality that he may impose during the meeting process (e.g. by preventing a buyer with a higher valuation from meeting

¹⁵We discuss this issue in detail and derive conditions for uniqueness in Cai et al. (2016).

¹⁶The meeting fee can be negative, turning it into a subsidy paid to each buyer. The equilibrium is of course not unique. Because of risk neutrality, the seller could for example condition the meeting fee on the number of buyers that shows up. However, all equilibria give rise to the same *expected* payoffs. See Peters and Severinov (1997), Albrecht et al. (2014), Lester et al. (2015) for detailed discussions of efficiency in related models.

the seller). Because this externality is type-independent, it can be priced by the meeting fee (5), which equals (the negative of) the spillovers that a buyer imposes on other buyers (the numerator) conditional on the event that he meets a seller (the probability of which is given by the denominator).

Posted Prices versus Auctions. It is worth highlighting that the equilibrium mechanism nests two popular trading mechanisms as special cases. As we discuss in more detail in the next section, technologies that exhibit joint concavity give rise to meeting fees that are non-positive. For a subset of those technologies, the equilibrium meeting fee is exactly zero, reducing the equilibrium mechanism to a *standard auction* or, equivalently, a menu of prices. In contrast, when meetings are bilateral, the second-price auction does not generate any revenue and the meeting fees, which are then strictly positive, act as a single *posted price*. Note that these are the only cases in which payments exclusively take place between the agents that ultimately trade; for all other meeting technologies, payments to or by buyers who do not trade are a necessary feature of the efficient equilibrium (see Lester et al., 2015).

4 Categorization of Meeting Technologies

Bilateral meetings are well understood, but joint concavity is a novel condition and warrants discussion. To better understand this condition, we compare it in this section to two other properties of meeting technologies described in the literature, *invariance* and *non-rivalry*. We show that invariance is a sufficient (but not a necessary) condition, while non-rivalry is a necessary (but not a sufficient) condition. Figure 1 summarizes this discussion.

Invariance. Introduced by Lester et al. (2015), an invariant technology is one in which the queue of blue buyers μ at a seller is a sufficient statistic for the distribution of the number of meetings between blue buyers and that seller. Formally,

$$\sum_{N=n}^{\infty} P_N(\lambda) \binom{N}{n} \left(\frac{\mu}{\lambda}\right)^n \left(1 - \frac{\mu}{\lambda}\right)^{N-n} = P_n(\mu), \quad (6)$$

for all $0 \leq \mu \leq \lambda < \infty$ and $n \in \mathbb{N}_0$. Perhaps the best-known example of an invariant technology is the urn-ball technology.¹⁷ In lemma 2, we establish that if (6) holds for $n = 0$, then it holds for all n . That is, invariance can alternatively be defined as the condition that the probability that a seller meets *at least one* of the μ blue buyers is independent of the number of other buyers visiting the same submarket.

Lemma 2. *A meeting technology is invariant if and only if $\phi_\lambda(\mu, \lambda) = 0$ for all $0 \leq \mu \leq \lambda < \infty$.*

¹⁷Recall that urn-ball implies $\phi(\mu, \lambda) = 1 - e^{-\mu}$. As discussed in Lester et al. (2015), a second example of an invariant technology is the geometric distribution $P_n(\lambda) = \frac{\lambda^n}{(1+\lambda)^{n+1}}$, which yields $\phi(\mu, \lambda) = \frac{\mu}{1+\mu}$.

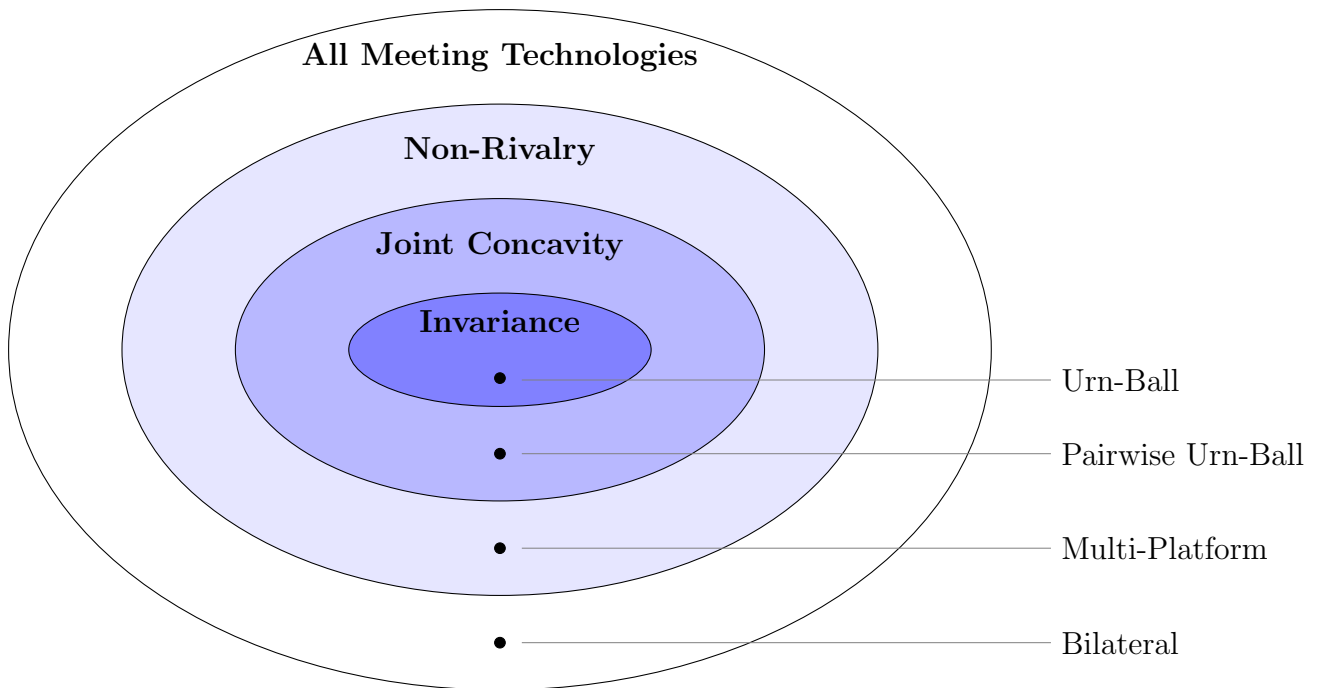


Figure 1: Venn Diagram of Meeting Technologies

Proof. See appendix A.5. □

Using this lemma, it is easy to establish that invariance is a sufficient condition for joint concavity: if $\phi_\lambda = 0$, then $\phi_{\mu\lambda} = \phi_{\lambda\lambda} = 0$ and joint concavity is (weakly) satisfied. In words, invariance implies that meetings between blue buyers and sellers are unaffected by the presence of other buyers in the submarket. Joint concavity therefore reduces to concavity in the measure of blue buyers, which is always satisfied.

In contrast, invariance is not necessary for joint concavity; the condition for joint concavity can hold without ϕ_λ , $\phi_{\mu\lambda}$ or $\phi_{\lambda\lambda}$ being zero. We prove this using the pairwise urn-ball technology: although this technology exhibits joint concavity, it is not invariant, as explained by Lester et al. (2015). Intuitively, when there are very few other buyers in the submarket, most buyer pairs consist of two blue types, making it likely that a seller will meet an even number of blue buyers. Adding additional other buyers to this submarket increases the probability that a buyer pair will consist of one blue and one other type, and that a seller will meet an odd number of blue buyers. This makes it more likely that a seller will meet at least one blue buyer, i.e. $\phi_\lambda > 0$, violating invariance. The following proposition formalizes this.

Proposition 5. *Invariance implies joint concavity, but joint concavity does not imply invariance.*

Proof. See appendix A.6. □

Although pairwise urn-ball is a convenient technology for proving this proposition, it is not

necessarily realistic. In the online appendix, we therefore discuss a second example of a meeting technology which satisfies joint concavity while violating invariance. We call this example the *spatial meeting technology*, as it describes how buyers in a submarket spread themselves across spatially dispersed sellers within each submarket: buyers stay relatively concentrated at a subset of the sellers (e.g. those that are near) if λ is small, but spread themselves more evenly across all sellers in the submarket if λ increases, yielding again $\phi_\lambda > 0$.¹⁸

Non-Rivalry. Eeckhout and Kircher (2010b) define a (purely) non-rival technology as one in which a buyer’s probability to meet one of the sellers is not affected by the presence of other buyers in the market. We first establish that their definition is equivalent to $\phi_{\mu\lambda}(0, \lambda) = 0$ for all $0 \leq \lambda < \infty$.

Lemma 3. *A meeting technology is non-rival if and only if $\phi_{\mu\lambda}(0, \lambda) = 0$ for all $0 \leq \lambda < \infty$.*

Proof. See appendix A.7. □

To understand this expression, recall that $\phi(\mu, \lambda)$ represents the probability that a seller meets at least one blue buyer, which is clearly zero if $\mu = 0$. The partial derivative $\phi_\mu(0, \lambda)$ captures how this changes if a single buyer (or more precisely, an arbitrarily small measure of buyers) in the queue becomes blue and must therefore equal the probability that this blue buyer succeeds in meeting the seller. Since meetings are type-independent, the same expression applies to all λ buyers in the queue, irrespective of how many of them are blue. Non-rivalry then says that this meeting probability should be independent of λ .

It is easy to verify that the above examples of technologies that exhibit joint concavity, i.e. urn-ball and pairwise urn-ball, both satisfy non-rivalry. This is not a coincidence. As the following proposition establishes, all technologies that exhibit joint concavity are non-rival. However, not all non-rival technologies exhibit joint concavity.

Proposition 6. *Joint concavity implies non-rivalry, but non-rivalry does not imply joint concavity.*

Proof. See appendix A.8. □

To understand why non-rivalry is a necessary condition for joint concavity, consider a submarket with an arbitrarily small measure of blue buyers, i.e. $\mu \rightarrow 0$. A seller’s probability of meeting at least one blue buyer in this submarket then goes to zero, $\phi(0, \lambda) = 0$. Since this is the case irrespective of the queue λ of other buyers, the derivatives of all orders of $\phi(0, \lambda)$ with respect to λ are zero as well. By (3), joint concavity then requires $\phi_{\mu\lambda}(0, \lambda) = 0$, which is exactly the condition for non-rivalry.

¹⁸We further prove in the online appendix that $\phi_\lambda \geq 0$ is a necessary but not a sufficient condition for joint concavity.

More generally, however, joint concavity requires that condition (3) is satisfied for any $0 \leq \mu \leq \lambda$ and not just for $\mu = 0$. Hence, non-rivalry is not sufficient. The multi-platform technology is a good example. This technology is non-rival, because every buyer meets a seller with probability 1. However, this technology clearly does not satisfy joint concavity for sufficiently large α , as it converges to a bilateral technology for $\alpha \rightarrow 1$. The proof of proposition 6 formalizes this and establishes that joint concavity is in fact violated for any $\alpha > 0$. Hence, non-rivalry does not imply joint concavity.

Surprisingly, this result (joint with proposition 3) contradicts proposition 5 in Eeckhout and Kircher (2010b) which states that non-rivalry is a sufficient condition for a single market. Inspection of the proof of their proposition reveals that the discrepancy arises because the proof implicitly assumes invariance rather than non-rivalry when treating the trading probability for high-type buyers as independent of the queue of low-type buyers.

5 Conclusion

We study an environment in which sellers post mechanisms to compete for buyers who differ in their private valuation for a good. Buyers can direct their search to the mechanism that maximizes their expected payoff, but may experience frictions in meeting a particular seller. We derive necessary and sufficient conditions on the technology that governs these meetings under which either a separate submarket for each type of buyer or a single market with all agents are optimal. We find that perfect separation is the efficient equilibrium outcome if and only if meetings are bilateral, while perfect pooling is optimal if and only if the meeting technology satisfies a novel property, which we call “joint concavity.”

Although we focus on a specific setting, some of our assumptions can be relaxed at no cost. For example, our results carry over to a setting in which buyer types are publicly observable rather than private information; this change simply means that auctions are no longer necessary and can be replaced by a menu of prices and an ex post selection rule, as in Shi (2006). Other assumptions are likely more crucial. For example, the nature of the equilibrium may change if the heterogeneous rather than the homogeneous side of the market posts, as in e.g. Delacroix and Shi (2013) and Albrecht et al. (2016). We leave a detailed analysis of arbitrary meeting technologies in such an environment for future research.¹⁹

A Proofs

A.1 Proof of Proposition 1

Part 1 (bilateral meetings imply perfect separation). Eeckhout and Kircher (2010b)

¹⁹See also Delacroix and Shi (2016), who study whether buyers or sellers should post.

demonstrate sufficiency for two types of buyers; we extend their result to arbitrary distributions. We do so in two steps: (i) buyers with different valuations must belong to different submarkets, and (ii) buyers with the same valuation must belong to the same submarket.

For (i), suppose that there exists a submarket with a measure s of sellers and a queue (λ, G) . Because of lemma 1 and the fact that meetings are bilateral, the surplus created in this submarket equals

$$sS(\lambda, G) = sP_1(\lambda) \int_0^1 (1 - G(x)) dx. \quad (7)$$

Now suppose the planner would decompose this submarket into a separate submarket for each type of buyer x , allocating the s sellers according to a distribution $H(x)$. Let $\hat{\lambda}(x) = \lambda \frac{dG(x)}{dH(x)}$ denote the queue length in submarket x . A seller in this submarket then creates a surplus $P_1(\hat{\lambda}(x))x$, such that surplus across all submarkets equals

$$s \int_0^1 P_1(\hat{\lambda}(x)) x dH(x). \quad (8)$$

Clearly, if the planner chooses $H(x) = G(x)$, then $\hat{\lambda}(x) = \lambda$ for all x and surpluses (7) and (8) are equal to each other. In that case, the marginal value of a seller in submarket x equals $(P_1(\lambda) - \lambda P_1'(\lambda))x$. This value is increasing in x , which means that the allocation of sellers is suboptimal and surplus (8) can be increased by sending some sellers to different submarkets. Since $P_1(\lambda)$ is concave, $P_1(\lambda) - \lambda P_1'(\lambda)$ is increasing in λ , and the planner can increase surplus by allocating relatively more sellers to the submarkets in which buyers have high valuations.

For (ii), suppose that there are two submarkets for buyers with valuation x . Let s_i and λ_i respectively denote the measure of sellers and the queue length in submarket $i \in \{1, 2\}$. By the strict concavity of $P_1(\lambda)$, merging the two submarkets then increases surplus:

$$x s_1 P_1(\lambda_1) + x s_2 P_1(\lambda_2) < x (s_1 + s_2) P_1\left(\frac{s_1 \lambda_1 + s_2 \lambda_2}{s_1 + s_2}\right).$$

Part 2 (perfect separation implies bilateral meetings). We prove this result in two steps. We first prove that (i) if $P_0(\lambda)$ is not convex, then the planner may create multiple submarkets, even when all buyers are homogeneous, and (ii) if $P_0(\lambda)$ is convex and, for some $\Pi > 0$, $P_0(\Pi) + P_1(\Pi) < 1$, then there exist two-point distributions of buyer types such that perfect separation is not optimal.²⁰

For (i), if $P_0(\lambda)$ is not convex, then by definition there exist s , λ_1 , and λ_2 such that

$$sP_0(\lambda_1) + (1 - s)P_0(\lambda_2) < P_0(s\lambda_1 + (1 - s)\lambda_2). \quad (9)$$

²⁰Of course, if $P_0(\Pi) + P_1(\Pi) < 1$, then—by continuity—there exists a small neighborhood of Π for which $P_0 + P_1 < 1$.

Now, let the market be populated by a measure 1 of sellers and a measure $s\lambda_1 + (1-s)\lambda_2$ of buyers with valuation 1. A single market with all agents generates a surplus of $1 - P_0(s\lambda_1 + (1-s)\lambda_2)$. However, the surplus generated by two submarkets—one with s sellers and $s\lambda_1$ buyers and the other one with the remaining buyers and sellers—equals $s(1 - P_0(\lambda_1)) + (1-s)(1 - P_0(\lambda_2))$, which is higher because of (9).

For (ii), let the market be populated by a measure 1 of sellers, a measure b_1 of buyers with valuation x_1 , and a measure b_2 of buyers with valuation x_2 , satisfying $b_1 + b_2 = \Pi$ and $x_2 > x_1$. We will prove the claim by contradiction. Suppose the planner fully separates the two types of buyers and optimally allocates s_i sellers to the submarket for valuation x_i , where $s_1 + s_2 = 1$. Define queue lengths $\lambda_i = \frac{b_i}{s_i}$.

Let now a measure ε of buyers with valuation x_1 and an equally large measure of buyers with valuation x_2 swap submarket, such that—in both submarkets—the queue lengths stay the same, but the composition of types becomes marginally more diverse. Again by lemma 1, social surplus of this new allocation equals

$$\mathcal{S}(\varepsilon) = s_2 \left[(x_2 - x_1) \phi \left(\frac{b_2 - \varepsilon}{s_2}, \lambda_2 \right) + x_1 \phi(\lambda_2, \lambda_2) \right] + s_1 \left[(x_2 - x_1) \phi \left(\frac{\varepsilon}{s_1}, \lambda_1 \right) + x_1 \phi(\lambda_1, \lambda_1) \right].$$

Clearly, $\varepsilon = 0$ corresponds to perfect separation. For this to be the optimal allocation, it must be the case that $\mathcal{S}(\varepsilon) \leq \mathcal{S}(0)$ for all $\varepsilon > 0$. Hence, a necessary condition is that $\mathcal{S}'(0) \leq 0$. Differentiating $\mathcal{S}(\varepsilon)$ and evaluating at $\varepsilon = 0$ gives

$$\mathcal{S}'(0) = (x_2 - x_1) (\phi_\mu(0, \lambda_1) - \phi_\mu(\lambda_2, \lambda_2)). \quad (10)$$

Next, we show that if x_1 and x_2 are sufficiently close, then equation (10) is strictly positive, implying that full separation is not optimal. First, note that $\phi_\mu(\mu, \lambda) = \sum_{n=1}^{\infty} \frac{nP_n(\lambda)}{\lambda} \left(1 - \frac{\mu}{\lambda}\right)^{n-1}$, which implies that

$$\phi_\mu(0, \Pi) - \phi_\mu(\Pi, \Pi) = \sum_{n=2}^{\infty} \frac{nP_n(\Pi)}{\Pi} \geq \sum_{n=2}^{\infty} \frac{2P_n(\Pi)}{\Pi} = \frac{2(1 - P_0(\Pi) - P_1(\Pi))}{\Pi} > 0.$$

Hence, by the continuity of $P_n(\lambda)$, there exists a $\Delta\Pi > 0$ such that for $\Pi - \Delta\Pi < \lambda_2 < \lambda_1 < \Pi + \Delta\Pi$, equation (10) is strictly positive. For $x_1 \rightarrow x_2$, we are back in the homogeneous case. In this limit, the convexity of $P(\lambda)$ implies that a single market with all agents is the desirable outcome, i.e. $\lambda_1 \rightarrow \Pi$ and $\lambda_2 \rightarrow \Pi$, such that equation (10) converges to

$$\frac{\mathcal{S}'(0)}{x_2 - x_1} = \phi_\mu(0, \Pi) - \phi_\mu(\Pi, \Pi) > 0. \quad (11)$$

By continuity, there then exists a $\Delta x > 0$ such that $x_2 - x_1 < \Delta x$ implies that $\mathcal{S}'(0) > 0$. We can therefore conclude that when x_1 is sufficiently close to x_2 , perfect separation is not optimal

and social surplus can be increased by slightly mixing the submarkets. \square

A.2 Proof of Corollary 1

The proof of this result resembles part 2 of the proof of proposition 1. As shown there, the requirement of only one submarket for each buyer type conditional on full separation implies that $P_0(\lambda)$ is convex. To establish the necessity of $P_0(\lambda) + P_1(\lambda) = 1$ for all λ , assume again that this condition is violated for some $\Pi > 0$. As established in the proof of proposition 1, this implies that there exists a $\Delta\Pi > 0$ such that $\phi_\mu(0, \lambda_1) > \phi_\mu(\lambda_2, \lambda_2)$ for $\Pi - \Delta\Pi < \lambda_2 < \lambda_1 < \Pi + \Delta\Pi$.

By assumption, the support of the type distribution F contains an interval $[x_0, x_2]$. Let $\lambda(x)$ denote the optimal queue length in submarket x , conditional on full separation, and select any point x_1 between x_0 and x_2 . As we will show below, there always exists an aggregate buyer measure Λ such that the optimal queue length in submarket x_1 is Π , i.e. $\lambda(x_1) = \Pi$. By continuity, we can then select Δx_1 small enough so that the optimal queue length around x_1 satisfies $\Pi - \Delta\Pi < \lambda(x_1 + \Delta x_1) < \lambda(x_1 - \Delta x_1) < \Pi + \Delta\Pi$. From the proof of proposition 1, it is then easy to see that swapping workers between submarket $x_1 - \Delta x$ and $x_1 + \Delta x$, where $\Delta x < \Delta x_1$, increases output. Hence, full separation is not optimal in this case.

The final step is then to prove the claim that there exists a buyer measure Λ such that $\lambda(x_1) = \Pi$, conditional on full separation. The desired Λ can be found in a constructive way. In submarket x , surplus per seller equals $x(1 - P_0(\lambda(x)))$ and buyers' marginal contribution to output is $-xP'_0(\lambda(x))$. By Euler's theorem, sellers' marginal contribution to surplus in this submarket then equals $x(1 - P_0(\lambda(x)) + \lambda(x)P'_0(\lambda(x)))$. Since sellers' marginal contributions to surplus are the same in submarket x and x_1 at the optimum, we have

$$x(1 - P_0(\lambda(x)) + \lambda(x)P'_0(\lambda(x))) = x_1(1 - P_0(\Pi) + \Pi P'_0(\Pi)). \quad (12)$$

Given x_1 and Π , the above equation provides an implicit solution for $\lambda(x)$.²¹ The desired buyer measure Λ then follows from

$$1 = \Lambda \int_0^1 \frac{1}{\lambda(x)} dF(x),$$

since the total seller measure is 1. \square

A.3 Proof of Proposition 2

Part 1 (joint concavity implies perfect pooling). To prove this result, suppose that there are two submarkets, indexed by $i \in \{1, 2\}$, consisting of $s_i > 0$ sellers who each have a queue

²¹Of course, if (12) does not admit a solution $\lambda(x)$, then we set $\lambda(x) = \infty$, i.e., we assign no sellers to that submarket because x is too small.

(λ_i, G_i) . By lemma 1, total surplus across the two submarkets is equal to

$$s_1 \int_0^1 \phi(\lambda_1(1 - G_1(x)), \lambda_1) dx + s_2 \int_0^1 \phi(\lambda_2(1 - G_2(x)), \lambda_2) dx. \quad (13)$$

We show a higher surplus can be generated by merging the two submarkets, creating one market with $s_0 = s_1 + s_2$ sellers, each with a queue $\lambda_0 = \frac{s_1\lambda_1 + s_2\lambda_2}{s_1 + s_2}$ of buyers whose valuations are distributed according to

$$G_0(x) = \frac{s_1\lambda_1 G_1(x) + s_2\lambda_2 G_2(x)}{s_1\lambda_1 + s_2\lambda_2}.$$

Again by lemma 1, this combined market will create a surplus $s_0 \int_0^1 \phi(\lambda_0(1 - G_0(x)), \lambda_0) dx$, which is larger than (13) because concavity of $\phi(\mu, \lambda)$ implies that

$$s_1\phi(\mu_1, \lambda_1) + s_2\phi(\mu_2, \lambda_2) \leq s_0\phi\left(\frac{s_1\mu_1 + s_2\mu_2}{s_1 + s_2}, \frac{s_1\lambda_1 + s_2\lambda_2}{s_1 + s_2}\right).$$

Hence, a single market is optimal for technologies that exhibit joint concavity.

Part 2 (perfect pooling implies joint concavity). We prove this result by showing that if ϕ is not concave, there exists a two-type distribution of buyers such that one market is not optimal. Note that if ϕ is not concave, then—by the definition of concavity—there exist values $\alpha, \mu_1, \mu_2, \lambda_1$ and λ_2 , such that

$$\alpha\phi(\mu_1, \lambda_1) + (1 - \alpha)\phi(\mu_2, \lambda_2) > \phi(\mu_0, \lambda_0), \quad (14)$$

where $\mu_0 = \alpha\mu_1 + (1 - \alpha)\mu_2$ and $\lambda_0 = \alpha\lambda_1 + (1 - \alpha)\lambda_2$.

Consider now a market in which buyers' valuations are either x_1 or x_2 , with $0 = x_1 < x_2$. Set the measure of high-type buyers equal to μ_0 and the measure of low-type buyers equal to $\lambda_0 - \mu_0$, while maintaining the assumption that the measure of sellers equals 1. Then by lemma 1, the social surplus of creating a single market is $\mathcal{S}_1 = x_2\phi(\mu_0, \lambda_0)$.

Now, decompose the single market into two submarkets A and B , with seller measures α and $1 - \alpha$, total queue lengths λ_1 and λ_2 , and high-type queue lengths μ_1 and μ_2 , respectively.²² The social surplus per seller for the two submarkets is $S_2^A = x_2\phi(\mu_1, \lambda_1)$ and $S_2^B = x_2\phi(\mu_2, \lambda_2)$, while total surplus across the two submarkets equals $\mathcal{S}_2 = \alpha S_2^A + (1 - \alpha)S_2^B$.

The two submarkets create more surplus than the single market, i.e. $\mathcal{S}_2 > \mathcal{S}_1$, if and only if

$$x_2(\alpha\phi(\mu_1, \lambda_1) + (1 - \alpha)\phi(\mu_2, \lambda_2)) > x_2\phi(\mu_0, \lambda_0),$$

which holds because it is exactly equation (14). Hence, joint concavity is a necessary condition

²²This is possible because $\mu_0 = \alpha\mu_1 + (1 - \alpha)\mu_2$ and $\lambda_0 = \alpha\lambda_1 + (1 - \alpha)\lambda_2$.

for a single market. \square

A.4 Proof of Proposition 3

First, consider the case in which F is degenerate. It is easy to see that pooling is then optimal if and only if $P_0(\lambda)$ is convex, which is the case even for a bilateral meeting technology.

Next, consider the case in which F has $N \in \{2, 3, \dots\}$ points of support, denoted by $x_1 < x_2 < \dots < x_N$. Let $0 < \gamma_n < 1$ denote the probability that x equals x_n and define $\zeta_n = \gamma_n + \gamma_{n+1} + \dots + \gamma_N$, for $n = 1, \dots, N$. Below, we perturb the urn-ball meeting technology, i.e. $P_n(\lambda) = e^{-\lambda} \lambda^n / n!$ or $\phi(\mu, \lambda) = 1 - e^{-\mu}$, to create a new meeting technology, $\tilde{P}_n(\lambda)$ or $\tilde{\phi}(\mu, \lambda)$, which violates joint concavity but satisfies the following property:

$$\tilde{\phi}(\mu, \lambda) \leq \phi(\mu, \lambda) \quad \text{with equality if and only if} \quad \frac{\mu}{\lambda} \in \{0, \zeta_1, \dots, \zeta_N\}. \quad (15)$$

To see why this proves the proposition, suppose that buyers and sellers in the market are distributed across $M > 1$ submarkets, such that—subject to aggregate consistency—submarket $i \in \{1, \dots, M\}$ consists of s_i sellers, each with queue λ_i and a type distribution described by $(\zeta_{i,1}, \dots, \zeta_{i,N})$. Total surplus under \tilde{P}_n then satisfies

$$\begin{aligned} \sum_{i=1}^M s_i \sum_{n=1}^N (x_n - x_{n-1}) \tilde{\phi}(\lambda_i \zeta_{i,n}, \lambda_i) &\leq \sum_{i=1}^M s_i \sum_{n=1}^N (x_n - x_{n-1}) \phi(\lambda_i \zeta_{i,n}, \lambda_i) \\ &\leq \sum_{n=1}^N (x_n - x_{n-1}) \phi(\Lambda \zeta_n, \Lambda) \\ &= \sum_{n=1}^N (x_n - x_{n-1}) \tilde{\phi}(\Lambda \zeta_n, \Lambda), \end{aligned}$$

where the four expressions follow from lemma 1, equation (15), proposition 2, and again equation (15), respectively. Hence, under \tilde{P}_n , a higher surplus can be generated by merging the M submarkets and the desired result follows.

To construct $\tilde{P}_n(\lambda)$, let a_0, \dots, a_{2N+1} denote the coefficients that solve

$$\sum_{n=0}^{2N+1} a_n x^n = (1-x) \prod_{n=1}^N (1-x-\zeta_n)^2, \quad (16)$$

for any x . We then define $\tilde{P}_n(\lambda)$ as

$$\tilde{P}_n(\lambda) = \begin{cases} P_n(\lambda) + \varepsilon(\lambda) a_n & \text{for } n = 0, 1, \dots, 2N+1 \\ P_n(\lambda) & \text{for } n > 2N+1, \end{cases} \quad (17)$$

where $\varepsilon(\lambda) > 0$ is a smooth function of λ , small enough to guarantee that $\tilde{P}_n(\lambda) \geq 0$ for any n .

For \tilde{P}_n to be a meeting technology, two additional conditions need to be satisfied: i) $\sum_{n=0}^{\infty} \tilde{P}_n(\lambda) = 1$, and ii) $\sum_{n=1}^{\infty} n\tilde{P}_n(\lambda)/\lambda \leq 1$. Substituting (17) reveals that these conditions are equivalent to $\sum_{n=0}^{2N+1} a_n = 0$ and $\sum_{n=1}^{2N+1} na_n \leq 0$, respectively. The former condition is satisfied because $x = 1$ is a root of the polynomial (16). To see that the latter condition is satisfied, evaluate the derivative of (16) with respect to x in $x = 1$ to get $\sum_{n=1}^{2N+1} na_n = -\Pi_{n=1}^N \zeta_n^2 < 0$.

To establish that \tilde{P}_n does not satisfy joint concavity, note that $\tilde{\phi}(0, \lambda) = 1 - \sum_{n=0}^{\infty} \tilde{P}_n(\lambda) = 0$ for all $0 \leq \lambda < \infty$. Hence, $\tilde{\phi}_{\lambda\lambda}(0, \lambda) = 0$. Note further that

$$\tilde{\phi}_{\mu\lambda}(0, \lambda) = \frac{d}{d\lambda} \sum_{n=1}^{\infty} \frac{n\tilde{P}_n(\lambda)}{\lambda} = \frac{d}{d\lambda} \sum_{n=1}^{\infty} \frac{nP_n(\lambda)}{\lambda} + \sum_{n=1}^{2N+1} na_n \frac{d\varepsilon(\lambda)}{d\lambda} \frac{1}{\lambda} = \sum_{n=1}^{2N+1} na_n \frac{d\varepsilon(\lambda)}{d\lambda} \frac{1}{\lambda}.$$

Since $\varepsilon(\lambda)$ needs to be chosen in such a way that $\tilde{P}_n(\lambda)$ is nonnegative for any λ , it is impossible that $\varepsilon(\lambda)$ is a linear function. Hence, $\tilde{\phi}_{\mu\lambda}(0, \lambda) \neq 0$. Consequently, the determinant of the Hessian evaluated in $\mu = 0$ equals $\tilde{\phi}_{\mu\mu}(0, \lambda)\tilde{\phi}_{\lambda\lambda}(0, \lambda) - \tilde{\phi}_{\mu\lambda}^2(0, \lambda) = -\tilde{\phi}_{\mu\lambda}^2(0, \lambda) < 0$, which means that joint concavity is violated.

Finally, note that

$$\phi(\mu, \lambda) - \tilde{\phi}(\mu, \lambda) = \sum_{n=0}^{\infty} \left(\tilde{P}_n(\lambda) - P_n(\lambda) \right) \left(1 - \frac{\mu}{\lambda} \right)^n = \varepsilon(\lambda) \sum_{n=0}^{2N+1} a_n \left(1 - \frac{\mu}{\lambda} \right)^n.$$

By equation (16), this expression is zero for $\frac{\mu}{\lambda} \in \{0, \zeta_1, \dots, \zeta_N\}$ and strictly positive otherwise. Hence, \tilde{P}_n satisfies equation (15).

A.5 Proof of Lemma 2

Part 1 (invariance implies $\phi_\lambda = 0$). Evaluating the definition of invariance (6) in $n = 0$ yields

$$\sum_{N=0}^{\infty} P_N(\lambda) \left(1 - \frac{\mu}{\lambda} \right)^N = P_0(\mu). \quad (18)$$

The left-hand side of this equation is $1 - \phi(\mu, \lambda)$ and the right-hand side is independent of λ . Hence, $\phi_\lambda(\mu, \lambda) = 0$ for all $0 \leq \mu \leq \lambda < \infty$.

Part 2 ($\phi_\lambda = 0$ implies invariance). Note that $\phi_\lambda(\mu, \lambda) = 0$ for all $0 \leq \mu \leq \lambda < \infty$ implies that $\phi(\mu, \lambda) = \phi(\mu, \mu)$. By equation (1), $\phi(\mu, \mu) = 1 - P_0(\mu)$. Consequently, equation (18) must hold for all $0 \leq \mu \leq \lambda < \infty$. By standard results from analytic function theory (see e.g. Ahlfors, 1979, p.32), we can differentiate both sides of this equation n times with respect to μ , which yields

$$\sum_{N=n}^{\infty} \frac{N!}{(N-n)!} P_N(\lambda) \left(-\frac{1}{\lambda} \right)^n \left(1 - \frac{\mu}{\lambda} \right)^{N-n} = P_0^{(n)}(\mu), \quad (19)$$

for all $0 \leq \mu \leq \lambda < \infty$. For $\mu = \lambda$, this gives $P_0^{(n)}(\mu) = \frac{n!}{(-\mu)^n} P_n(\mu)$. Substitute this into the right hand side of equation (19) and rearrange the term $\frac{n!}{(-\mu)^n}$ to the left hand side gives (6). Hence, $\phi_\lambda = 0$ implies invariance. \square

A.6 Proof of Proposition 5

Part 1 (invariance implies joint concavity). This result follows immediately from lemma 2: $\phi_\lambda(\mu, \lambda) = 0$ for all $0 \leq \mu \leq \lambda < \infty$ implies that $\phi_{\lambda\lambda}(\mu, \lambda) = \phi_{\mu\lambda}(\mu, \lambda) = 0$ for all $0 \leq \mu \leq \lambda < \infty$, which in turn implies that equation (3) is satisfied.

Part 2 (joint concavity does not imply invariance). Consider the pairwise urn-ball technology, which satisfies $\phi(\mu, \lambda) = 1 - e^{-\mu(1 - \frac{1}{2}\frac{\mu}{\lambda})}$. Since $1 - e^{-y}$ is an increasing, concave function, a sufficient condition for $\phi(\mu, \lambda)$ to be concave is that the map $(\mu, \lambda) \rightarrow \mu(1 - \frac{1}{2}\frac{\mu}{\lambda})$ is concave (see e.g. Theorem 5.1 in Rockafellar, 1970, p.32). The Hessian of this map is indeed negative semi-definite. However, the technology is not invariant, as

$$\phi_\lambda(\mu, \lambda) = \frac{1}{2} \frac{\mu^2}{\lambda^2} e^{-\mu(1 - \frac{1}{2}\frac{\mu}{\lambda})} > 0.$$

Hence, joint concavity does not imply invariance. \square

A.7 Proof of Lemma 3

As shown by Lester et al. (2015), non-rivalry is satisfied if and only if $\frac{\partial}{\partial \lambda} \frac{1}{\lambda} \sum_{n=1}^{\infty} n P_n(\lambda) = 0$, for all $0 \leq \lambda < \infty$. The desired result then follows directly from observing that $\phi_\mu(0, \lambda) = \frac{1}{\lambda} \sum_{n=1}^{\infty} n P_n(\lambda)$. \square

A.8 Proof of Proposition 6

Part 1 (joint concavity implies non-rivalry). For any meeting technology, $\phi(0, \lambda) = 0$ for all $0 \leq \lambda < \infty$, i.e. the probability of meeting a blue buyer is zero if $\mu = 0$, irrespective of the value of λ . This implies that $\frac{\partial^n}{\partial \lambda^n} \phi(0, \lambda) = 0$ for all $n \in \mathbb{N}_1$.

Suppose now that a technology does not satisfy non-rivalry, i.e. $\phi_{\mu\lambda}(0, \lambda) \neq 0$. Then

$$\phi_{\mu\mu}(0, \lambda) \phi_{\lambda\lambda}(0, \lambda) - \phi_{\mu\lambda}^2(0, \lambda) < \phi_{\mu\mu}(0, \lambda) \phi_{\lambda\lambda}(0, \lambda) = 0,$$

i.e. ϕ does not exhibit joint concavity. Hence, joint concavity implies non-rivalry.

Part 2 (non-rivalry does not imply joint concavity). Consider the multi-platform technology. Starting from the expression for $\phi(\mu, \lambda)$ for this technology, one can derive

$$\phi_{\mu\lambda} = - \left(1 - e^{-\frac{\lambda\mu}{(1-\alpha)(\lambda+\alpha)}} - \frac{\lambda\mu}{(1-\alpha)(\lambda+\alpha)} e^{-\frac{\lambda\mu}{(1-\alpha)(\lambda+\alpha)}} \right) \frac{\alpha}{(\lambda+\alpha)^2} \leq 0,$$

which equals 0 (only) for $\mu = 0$. Hence, the multi-platform technology is non-rival.

Further, we get

$$\phi_\lambda = - \left(1 - e^{-\frac{\lambda\mu}{(1-\alpha)(\lambda+\alpha)}} \right) \frac{\mu\alpha}{(\lambda + \alpha)^2},$$

which is strictly negative for all $0 < \mu \leq \lambda < \infty$ and $\alpha > 0$. As we show in the online appendix, $\phi_\lambda \geq 0$ is a necessary condition for joint concavity. Hence, the multi-platform technology does not exhibit this property. \square

References

- Ahlfors, L. (1979). *Complex Analysis*. McGraw-Hill.
- Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2006). Equilibrium directed search with multiple applications. *Review of Economic Studies*, 73:869–891.
- 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., Gautier, P. A., and Vroman, S. B. (2016). Directed search in the housing market. *Review of Economic Dynamics*, 19:218–231.
- Auster, S. and Gottardi, P. (2016). Competing mechanisms in markets for lemons. mimeo.
- 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. (2016). Efficiency of wage bargaining with on-the-job search. mimeo.
- Cai, X., Gautier, P. A., and Wolthoff, R. P. (2016). Meetings and mechanisms. mimeo.
- Delacroix, A. and Shi, S. (2013). Pricing and signaling with frictions. *Journal of Economic Theory*, 148:1301–1332.
- Delacroix, A. and Shi, S. (2016). Should buyers or sellers organize trade in a frictional market? 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.
- Epstein, L. and Peters, M. (1999). A revelation principle for competing mechanisms. *Journal of Economic Theory*, 88(1):119–161.

- Fraja, G. D. and Sákovics, J. (2001). Walras retrouvé: Decentralized trading mechanisms and the competitive price. *Journal of Political Economy*, 109(4):pp. 842–863.
- Galenianos, M. and Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 114:445–471.
- Gautier, P. A. and Holzner, C. (2017). Simultaneous search and efficiency of entry and search intensity. *American Economic Journal: Microeconomics*, forthcoming.
- Gautier, P. A., Moraga-Gonzalez, J. L., and Wolthoff, R. P. (2016). Search costs and efficiency: Do unemployed workers search enough? *European Economic Review*, 84:123–139.
- Gautier, P. A., Teulings, C. N., and Van Vuuren, A. (2010). On-the-job search, mismatch and efficiency. *The Review of Economic Studies*, 77(1):245–272.
- Gautier, P. A. and Wolthoff, R. P. (2009). Simultaneous search with heterogeneous firms and ex post competition. *Labour Economics*, 16:311–319.
- Guerrieri, V., Shimer, R., and Wright, R. (2010). Adverse selection in competitive search equilibrium. *Econometrica*, 78(6):1823–1862. mimeo.
- Hall, R. (1977). *Microeconomic Foundations of Macroeconomics*, chapter An Aspect of the Economic Role of Unemployment. Macmillan, London.
- Hosios, A. J. (1990). On the efficiency of matching and related models of search and unemployment. *Review of Economic Studies*, 57:279–298.
- Kircher, P. (2009). Efficiency of simultaneous search. *Journal of Political Economy*, 117:861–913.
- Kiyotaki, N. and Wright, R. (1993). A search-theoretic approach to monetary economics. *American Economic Review*, 83(1):63–77.
- 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.
- Lester, B., Visschers, L., and Wolthoff, R. P. (2017). Competing with asking prices. *Theoretical Economics*, forthcoming.
- Lester, B. and Wolthoff, R. P. (2014). Interviews and the assignment of workers to jobs. mimeo.
- 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.
- 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.
- Shi, S. (2001). Frictional assignment I: Efficiency. *Journal of Economic Theory*, 98:232–260.
- Shi, S. (2002). A directed search model of inequality with heterogeneous skills and skill-biased technology. *Review of Economic Studies*, 69(2):467–491.
- Shi, S. (2006). Wage differentials, discrimination and efficiency. *European Economic Review*, 50:849–875.
- Shi, S. (2009). Directed search for equilibrium wage-tenure contracts. *Econometrica*, 77:561–584.
- 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. (2017). Applications and interviews: Firms’ recruiting decisions in a frictional labor market. *The Review of Economic Studies*, forthcoming.