Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 145 (2010) 1354–1385
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Sorting versus screening: Search frictions and competing mechanisms ☆

Jan Eeckhout [a,b,∗], Philipp Kircher [a,c,d]

[a] *Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA, United States*
[b] *ICREA-GSE-UPF, Barcelona, Spain*
[c] *University of Oxford, United Kingdom*
[d] *IZA, Germany*

Received 17 February 2009; final version received 31 December 2009; accepted 14 January 2010

Available online 22 January 2010

Abstract

In a market where sellers compete by posting trading mechanisms, we allow for a general search technology and show that its features crucially affect the equilibrium mechanism. Price posting prevails when meetings are rival, i.e., when a meeting by one buyer reduces another buyer's meeting probability. Under price posting buyers reveal their type by sorting ex-ante. Only if the meeting technology is sufficiently non-rival, price posting is not an equilibrium. Multiple buyer types then visit the same sellers who screen ex-post through auctions.

JEL classification: C78; D44; D83

Keywords: Competing mechanism design; Matching function; Meeting function; Sorting; Screening; Price posting; Auctions

☆ This paper used to circulate under the title "Sorting versus screening: Prices as optimal competitive sales mechanisms". We are grateful to Shouyong Shi, two anonymous referees and the associate editor for valuable comments. We also benefited from the feedback of several seminar audiences.

∗ Corresponding author at: Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA, United States.

E-mail addresses: eeckhout@ssc.upenn.edu (J. Eeckhout), kircher@econ.upenn.edu (P. Kircher).

0022-0531/$ – see front matter
doi:10.1016/j.jet.2010.01.005

1. Introduction

Prices are much more prevalent than auctions, yet common wisdom has it that auctions can achieve more than prices can. In this paper we argue that while this wisdom is true in the partial equilibrium setting of a monopolistic principal, in competitive markets with competing mechanisms this need not be true. In particular, we show that the prevalence of prices over more general auction-like mechanisms crucially depends on the features of the meeting technology. If buyers only rarely end up simultaneously bidding for the same good, sellers choose price posting. In contrast, when buyers do tend to simultaneously compete, for example in art or antique auctions, prices are dominated. Our findings highlight the role of the search process for mechanism design. The important insight here is that it is not necessarily the fine details of the mechanism space that determine the competitive sales mechanism, but rather the properties of the meeting process. We can thus characterize the prevalence of price posting as a function of the meeting technology.[1]

The role of the meeting technology can best be illustrated by considering two extreme versions that are commonly assumed. First, consider a *purely non-rival* meeting technology, as is often done in much of the directed search literature. Buyers simultaneously meet a given seller and they all contemporaneously compete for the good for sale. Each additional meeting by another buyer does not affect one's chances of meeting with the seller. Key here is the distinction between meeting and matching (or trade). Even if meeting is non-rival, the good itself is clearly rival: the more buyers meet, the lower the trading probability. As an example of a non-rival meeting technology, consider a seller of a piece of art who fixes a date and time when the good will be sold. Irrespective of how many other buyers turn up, the opportunity to enter the auction is invariant. Second, consider a *purely rival* meeting technology, as in much of the competitive search literature. At any given seller, there is always at most one buyer at the time. Another buyer's meeting clearly reduces one's own meeting probability. This is often the case in environments without recall where in any small time interval there is at most one meeting which must immediately end up in trade or separation. For example, a firm continuously hires and once a candidate turns up, a hiring decision is made.[2] There is of course a whole continuum in between these extreme meeting technologies. Suppose several workers simultaneously apply for a job, but the firm only considers say half of the applications (there could be many reasons: it is too costly, only those that have been referred by trusted friends and colleagues are considered, ...). This renders a meeting technology *partially rival*. We are not aware of work that considers the impact of variations in the meeting technology, and this work attempts to fill the gap.

The approach in most of the search literature is to assume a particular trading arrangement (typically price posting, but in other instances also competition in auctions) without questioning whether this particular mechanism would actually be chosen as an equilibrium outcome when a set of different mechanisms is available. In contrast, the competing mechanism design litera-

[1] Price posting is pervasive in many economic transactions. Even the internet auction house eBay derives 40% of its revenue from price posting. There could be many reasons why prices are pervasive, including low transaction costs (see for example Wang [35]). Our objective is to find out under which conditions price posting is an efficient trading mechanism in the presence of search frictions, and without assuming different transaction costs for other mechanisms.

[2] The purely rival meeting technology is maintained in work by Moen [18], Acemoğlu and Shimer [1], Mortensen and Wright [20], Moen and Rosen [19]. The purely non-rival meeting technology is assumed in such work as Peters [22–26], Peters and Severinov [28], Burdett, Shi and Wright [4], Shi [31,32], Shimer [33]. Even in random search, often a rival meeting function is assumed where bad types negatively affect good types (see for example Albrecht and Vroman [2]), but alternatives with non-rival meeting technologies have recently been proposed (see for example Moscarini [21] and Albrecht and Vroman [2]).

ture [17,24] does ask what the equilibrium mechanism is, but it analyzes this in the presence of a particular purely non-rival meeting technology only. They show that in their setting second price auctions are always a weak best reply for an individual seller. They derive an equilibrium where buyers visit all sellers with equal probability, thus rendering visit strategies purely random. Once buyers turn up, sellers use the auctions for *ex-post screening*.

In contrast, when sellers are restricted to using simple price posting mechanisms, they will offer different prices for different buyer types to induce separation of buyers. Lower type buyers choose to visit sellers who offer low prices and corresponding low odds of trade, while high type buyers consume at high prices and enjoy a high probability of trade. Such trading mechanism leads to *ex-ante sorting*, with buyers endogenously revealing their type in equilibrium by choosing the price at which they want to trade. Due to the separation, each seller knows exactly the type of buyer he faces, and has ex-post no incentive to use type-revealing mechanisms.

A priori, it is not clear whether ex-ante sorting or ex-post screening will prevail in equilibrium. We are interested in how the nature of the meeting process affects the equilibrium trading mechanism. In particular, we investigate under which circumstances a simple price posting mechanism obtains in equilibrium. The key feature is to analyze posted prices as an equilibrium mechanism, even when other mechanisms are available. Our paper thus spans the literatures of directed/competitive search and competing mechanism design, and links the prevalence of posted price mechanisms tightly to the properties of the meeting technology.

We have four distinct results that highlight how the equilibrium trading mechanism depends on the meeting technology and the degree of heterogeneity in buyer preferences. *First*, in the absence of heterogeneity, we obtain an *equivalence result* independent of the exact nature of the meeting technology. For any mechanism, seller revenues are identical conditional on leaving the same expected surplus to buyers. Given revenue equivalence, sellers do not care whether they compete in posted prices, second price auctions or other mechanisms, and therefore competition in posted prices does constitute an equilibrium. There is a continuum of other equilibria in different mechanisms, but in all equilibria visit strategies are random and payoffs are invariant.

Our *second* result concerns purely rival meetings and establishes that fixed price mechanisms constitute an equilibrium and it is constrained efficient. Extending the usual notion of constrained efficiency for given fixed price mechanisms to encompass competition in larger classes of mechanisms, the key observation is that random visit strategies are not efficient because this leaves sellers unsure about the type of buyer they face. While there are still alternative ways of screening buyer types such as lotteries, incentive compatibility induces a cost in terms of wasteful destruction. More importantly, non-random visit strategies outperform random visit strategies because meetings are rival: When low types enter the same market as high types, then the probability of meeting a seller goes down for the high types. This makes it beneficial to keep buyers apart in separate markets. Ex-post screening is then no longer necessary because efficiency requires sellers to go to different markets. These efficiency concerns also drive equilibrium behavior, and sellers post different mechanisms that attract different buyer types. Since sorting leaves not residual uncertainty about the buyer types, again ex-post screening mechanisms are not necessary and prices suffice to shift resources between buyers and sellers.

Our *third* main result shows that the constrained efficiency and the equilibrium nature of fixed price mechanisms do *not* carry over to purely non-rival meeting technologies. This arises even though each seller knows exactly the type of buyers that he faces in equilibrium when only fixed prices are available, i.e., there is perfect sorting. Despite the fact that the well-known Hosios [11] condition for constrained efficiency is fulfilled for each market when only prices are available, other mechanisms generate higher surplus and are more profitable for an individual buyer. This has to do with the meeting technology: under ex-post screening, visit strategies are random in equilibrium which leads to strictly more meetings than non-random visit strategies. The gain from having more meetings is not eroded by the fact that low types search in the same market as high types, precisely because meetings are non-rival. Therefore, random visit strategies of all buyer types in one market and allocating the good via auctions are constrained efficient and arise in equilibrium. Even if all other sellers offer posted prices, a single deviant can exploit the constrained efficiency gain of random visit strategies by having all buyer types visit him. This is consistent with the result in McAfee [17] that second price auctions are weak best replies in his setting with non-rival meetings.

Our *fourth* result establishes that the prevalence of price posting as an equilibrium mechanism holds true generally for partially rival meeting technologies and does not exclusively hinge on the pure rival nature. As long as the degree of partial rivalry is high enough, ex-ante sorting via price posting will dominate ex-post screening. Even if there is some ex-post competition and multiple buyers meet the seller, it is not in the interest of the seller to announce a mechanism that screens ex-post. While ex-post screening through auctions would arise in a partial equilibrium setting, in the presence of competition from other sellers, an individual seller attracts more buyers by announcing a fixed price mechanism and thus generates a higher surplus.[3]

Our work relates to existing work on constrained efficiency in search markets. While the random search model is typically inefficient [11], Moen [18] shows that in the competitive search model with identical agents and rival meetings, the price posting equilibrium is constrained efficient. He considers a planner who faces the same meeting frictions and allocates the good with the same mechanisms (in his case, prices) as the agents in the decentralized economy. Then with quasi-linear preferences and in the presence of lump sum transfers, Pareto efficiency is equivalent to maximizing the surplus of trade.

Our contribution on the issue of efficiency is two-fold. From the equivalence result in the case of identical buyers, it follows that Moen's result extends beyond a non-rival meeting technology and for different trading mechanisms. For heterogeneous buyers, we show that for purely rival and purely non-rival meeting technologies, the equilibrium outcome – posted prices under purely rival meetings, second price auctions under purely non-rival meetings – is constrained efficient. We establish this result for the pure private value case and abstract from any common-value components. In particular, we do not consider adverse selection and associated lemons problems. Guerrieri, Shimer and Wright [10] combine differences in valuations of buyers, as in our setting, with adverse selection for the seller. Focusing on purely rival meetings, their equilibrium features separation of types, confirming more broadly the separation that we find in our environment. We conjecture that even in such a setting, the equilibrium will feature pooling once the rivalry becomes less severe, but since ours is the first step to see the implications of different meeting technologies, such results are beyond the scope of the present paper.

The next section outlines and analyzes the model with homogeneous buyers for a general meeting technology that allows for any degree of partial rivalry, spanning from purely rival to purely non-rival specifications. Section 3 extends the setup to encompass buyer heterogeneity. We focus on two buyer types for tractability, yet the analysis extends to any finite number of buyers. Section 4 analyzes the outcome when sellers are restricted to post prices. Section 5 analyzes the

[3] Observe that when price posting is an equilibrium, there may exist other equilibrium mechanisms that are payoff equivalent. For example, sellers may offer an auction with a reserve price that is high enough to attract only one type and that therefore effectively does not screen ex-post. This can generate the same expected surplus, which resembles our first result for homogeneous agents.

equilibria and efficiency when other mechanisms are also available. The analysis covers the cases of multilateral meetings (purely rival meetings), bilateral meetings (purely non-rival meetings), as well as those meeting technologies that are partially rival. Section 6 offers some concluding remarks.

2. Homogeneous buyers

2.1. The model

Consider an economy with a measure s of homogeneous sellers and a measure b of buyers. We start here with the case of homogeneous buyers, and later extend the setup to heterogeneous buyers. Sellers are endowed with one unit of a good for sale. Buyers have a valuation v for the good. If a buyer and a seller trade at a transfer t from the buyer to the seller, then the buyer's payoff is $v - t$ and the seller's payoff is t. Payoffs from no trade are normalized to zero. We are interested in trading procedures where search is directed in the sense that buyers can post a mechanism and sellers can choose which mechanism to participate in.

The market interaction. Each seller decides on a mechanism m from some Borel-measurable mechanism space $\mathcal{M}$, which we will define in detail below. The mechanisms posted by sellers can be summarized by the measure μ_s, where $\mu_s(M)$ denotes the measure of sellers that post mechanisms in set $M \subset \mathcal{M}$. Different mechanisms trade in different markets. Buyers can see which mechanisms are offered, and decide in which market to search.[4] This means that they can direct their search towards the mechanisms they find most attractive. Their trading decisions can be summarized by measure μ_b, where $\mu_b(M)$ denotes the measure of buyers that search in markets that offer a mechanism in set $M \subset \mathcal{M}$. The measures cannot exceed the overall measure of agents in the population.[5] The requirement that buyers can only search for mechanisms that are actually offered by some sellers is captured by the assumption that μ_b is absolutely continuous in μ_s. The Radon–Nikodym (RN) derivative $\lambda = d\mu_b/d\mu_s$ then delivers the buyer–seller ratio in each market place in the support of μ_s.[6] Since $\lambda(m)$ depends on the mechanism that is offered, sellers can affect the number of buyers that they attract by changing the mechanism that they offer. The buyer–seller ratio will be crucial in determining the meeting probabilities between buyers and sellers, and only the buyers that actually meet a seller can take part in the seller's posted mechanism.

Meetings. If the ratio of buyers to sellers in a particular market is λ, there is a probability $P_n(\lambda)$ that n buyers will show up at a given seller. P_n is assumed to be twice differentiable. From a buyer's perspective, there is a probability $Q_n(\lambda)$ that he arrives at a seller who has n buyers, with

[4] Following most of the literature, sellers are assumed to search only for one mechanisms. For simultaneous search for multiple mechanisms see for example Albrecht, Gautier and Vroman [3], Galenianos and Kircher [8] and Kircher [13].

[5] For sellers there always exists a weakly profitable mechanism (e.g., a positive price), and therefore we require $\mu_s(M) = s$. Sellers might not see any mechanism in the support of μ_s that gives them a higher payoff then abstaining, and therefore we only require $\mu_b(M) \leqslant b$.

[6] The Radon–Nikodym-derivative is almost everywhere unique on the support of μ in the sense any two RN-derivative coincide almost everywhere. To have well-defined payoffs, we also assume a selection device that selects a unique RN-derivative.

$Q_0(\lambda)$ being the probability of not finding any seller.[7] Since buyers and sellers meet jointly, there is a particular consistency condition that links P_n and Q_n. When there are σ sellers and β buyers in a market such that $\lambda = \beta/\sigma$, consistency requires for every $n > 0$ that

$$\sigma P_n(\lambda) = \frac{\beta Q_n(\lambda)}{n}.$$

The left-hand side gives the number of sellers in a meeting with n buyers. The numerator on the right-hand side gives the number of buyers in "n-buyer" meetings, and since there are n of them per seller the denominator scales it down to the number of sellers in such meetings. Rearranging this equation gives

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

For $n = 0$ the fact that probabilities add to unity implies that $Q_0(\lambda) = 1 - \sum_{n=1}^{\infty} Q_n(\lambda)$. We also require that for any $\lambda > 0$ there are some meetings, i.e. $0 < 1 - P_0(\lambda)$, but due to frictions not all sellers meet a buyer, i.e. $0 < P_0(\lambda)$. Moreover, we require that a higher buyer–seller ratio increases the meeting chances for a seller, i.e., we assume that the probability that no buyer shows up strictly decreases in λ and that the probability of meeting less than N buyers decreases (and the probability of meeting more than N buyers increases). Formally, we assume a shift in terms of first-order stochastic dominance:

$$\begin{aligned} &P_0'(\lambda) < 0 \quad \text{and} \\ &\sum_{n=0}^{N} P_n'(\lambda) \leqslant 0 \quad \text{for all } N. \end{aligned} \tag{2}$$

The first line is assumed throughout the literature to capture the idea that more buyers per seller makes it easier for a seller to find some buyer to trade with. Our main extension is in the second line which covers meetings with multiple buyers. The higher the buyer–seller ratio, the higher the number of buyers that will arrive (in stochastic terms). For the buyers, we assume a similar property. In particular, we assume that the distribution satisfies first-order stochastic dominance, i.e. $\sum_{n=N}^{\infty} Q_n'(\lambda) \leqslant 0$ for all N with strict inequality for some N. This immediately implies that $\sum_{n=1}^{\infty} \frac{Q_n'(\lambda)}{n} < 0$, which in turns implies the following relationship that will be useful later on: $1 + \lambda P_0'(\lambda^\star) - P_0(\lambda^\star) > 0$.[8]

Additionally we assume that

$$P_0 \text{ is strictly convex.} \tag{3}$$

This convexity property has to hold for some $\lambda \in [0, \infty)$ since $P_0(\lambda) \in [0, 1]$ is bounded, and we make the standard assumption that this property extends to the entire domain. Again, this condition is standard in the literature.

---

[7] Our meeting technology is inherently static, but it can capture the notion of a dynamic search process. For example, Pinheiro [29] proposes a dynamic model where the number of traders that meet the seller is a function of the time spent searching.

[8] Since $(1 - P_0(\lambda)) = \sum_{n=1}^{\infty} \lambda \frac{Q_n(\lambda)}{n}$ we have that $-\frac{1-P_0(\lambda)+\lambda P_0'(\lambda)}{\lambda^2} = \sum_{n=1}^{\infty} \frac{Q_n'(\lambda)}{n}$, and the result follows from $\sum_{n=1}^{\infty} \frac{Q_n'(\lambda)}{n} < 0$.

Examples. The standard example in the directed search literature is that buyers randomly visit one of the sellers that offers the mechanism that they like. Clearly the probability that the buyer meets a sellers is one, and the only question is how many other buyers are present. The randomness in a given market leads to a Poisson distribution $P_n(\lambda) = \frac{\lambda^n e^{-\lambda}}{n!}$ (and associated $Q_n(\lambda) = \frac{\lambda^{n-1} e^{-\lambda}}{(n-1)!}$ when $n > 0$). A specification from the monetary literature (see e.g. Kiyotaki and Wright [14]) is one where agents are randomly matched into pairs, and trade only if the pair includes a buyer and a seller. If a seller is paired with another seller, there is no meeting that can lead to trade. For a seller, the probability to be in a meeting with a buyer rather than with another seller when there are σ sellers and β buyers is $P_1(\lambda) = \frac{\sigma}{\sigma+\beta} = \frac{\lambda}{1+\lambda}$. The probability of a meeting without any buyers is $P_0(\lambda) = 1 - P_1(\lambda)$, and clearly it is impossible to meet more than one buyer so that $P_n(\lambda) = 0$ for $n > 1$.

Mechanisms. Sellers compete in mechanisms from some Borel-measurable set $\mathcal{M}$ with the following restrictions. We require that M includes the set of all fixed price mechanisms, i.e., sellers post a price and sell at this price to one of the buyers that show up. If several buyers show up the seller picks one at random to whom to sell. We require the set $\mathcal{M}$ to include only anonymous mechanisms that do not condition on the other mechanisms that are present.[9] In particular, a mechanism specifies for each number n of sellers some extensive form game Γ_m^n that induces some expected payoff π_n^m for the seller and some expected payoff u_n^m for each of the buyers.[10] Since only a surplus v is realized, we require that the payoffs at each end node of Γ_m^n sum be weakly less than v. This immediately implies $\pi_n^m + n u_n^m \leqslant v$. The mechanisms do not have to obey any participation constraints, as sellers can always choose to stay away from any mechanism if their expected payoffs (averaged over the expectation of n, the number of other buyers that are expected to turn up) are too low.[11] Obviously if there are no buyers, then $u_0^m = 0$ and $\pi_0^m = 0$.

Examples. Fixed price mechanisms where sellers post price p have payoffs $\pi_n = p$ and $u_n = [1 - p]/n$ conditional on being in a meeting with n buyers. Other feasible mechanisms include first price auctions, second price auctions, all-pay auctions with reserve price, as well as more esoteric mechanisms such as: If no more than 5 buyers show up, one of them gets the good for free. If more than 5 buyers show up, each has to pay a price p and one gets the good. This yields payoffs $\pi_n = 0$ and $u_n = 1/n$ for $n \in \{1, \dots, 5\}$ and $\pi_n = np$ and $u_n = 1/n - p$ for $n \in \{6, 7, \dots\}$.

Payoffs. Consider the expected payoffs for an individual agent when the trading strategies of the other agents are summarized by μ_s and μ_b. The agent understands how crowded all markets

[9] Since buyers observe all other mechanisms, the seller could elicit information about other sellers from the buyers. This is ruled out by assumption in most of the literature such as McAfee [17] or Peters [24,25]. In large economies or in a pure strategy equilibrium in a finite economy each seller knows the distribution of other mechanisms with certainty, and there is no benefit from conditioning on other mechanisms when no seller does so. Therefore, equilibria survive even if sellers can condition on the mechanisms offered by other sellers. Yet additional equilibria can arise. For a deeper discussion and modeling of mechanisms that condition on other mechanisms, see Epstein and Peters [7] and Peters [27].

[10] If the game has multiple equilibria and therefore multiple expected payoffs, we assume additionally that the seller can post an equilibrium selection device S_m. In general this is not an issue because there are alternative games (such as direct commitment to the payoffs) that give the desired outcome uniquely.

[11] If they arrive they are assumed to sign an agreement to participate before n is revealed.

are. Consider first the expected payoff to a buyer who chooses some mechanisms in the support of μ_s:

$$U(m|\mu_s,\mu_b)=\sum_{n=1}^{\infty}Q_n\big(\lambda(m)\big)u_n^m, \tag{4}$$

which simply captures the probability of being in a match with n buyers times the payoff from playing the mechanism. Similarly, a seller who offers a mechanisms in the support of μ_s obtains expected profits:

$$\Pi(m|\mu_s,\mu_b)=\sum_{n=1}^{\infty}P_n\big(\lambda(m)\big)\pi_n^m. \tag{5}$$

Finally, we have to specify the payoffs that a seller expects to obtain if he offered a mechanism that no other seller is offering, i.e., if he posts m outside of the support of μ_s. At this point the buyer–seller ratio is not tied down by the trading strategies μ_b and μ_s in the sense that the Radon–Nikodym-derivative is arbitrary. We complete the specification following the literature (e.g. [1,6]) by appealing to a notion of subgame perfection. It requires that the buyer–seller ratio that a deviant expects gives buyers as high payoff as they would obtain elsewhere in the market. Formally, let the queue length $\lambda(m)$ at a deviant who offers mechanisms m outside the support of μ_s satisfy

$$\sum_{n=1}^{\infty}P_n\big(\lambda(m)\big)\pi_n^m=\sup_{m\in\operatorname{supp}\mu_s}U(m|\mu_s,\mu_b), \tag{6}$$

if this equality can be achieved for some $\lambda(m)>0$, and otherwise let $\lambda(m)=0$.[12] This specification captures the following idea: A strictly positive buyer–seller ratio means that some buyers must be willing to search for the deviant. These sellers would only be willing to do so if the utility from searching in the market of the deviant (left-hand side) is at least as high as the utility they get on the equilibrium path (right-hand side). The utility at the deviant can also not be strictly higher than the utility on the equilibrium path, because in this case all buyers would want to search in his market which would drive up the buyer–seller ratio. While this is an informal argument, (6) can be derived as the subgame perfect equilibrium of particular finite market games when the market size grows large (see e.g. [22–24,26]). For many mechanisms such as price setting and standard auctions, (6) determines the queue length uniquely. In case of multiplicity, we follow McAfee [17] and others and assume that the seller believes that he can coordinate buyers in the way that is most desirable to him. Given $\lambda(m)$, payoffs are determined as in (5).

Equilibrium. We define an equilibrium as a large game [16], where each agent individually acts optimally, taking the overall trading strategies of the other agents as given. An equilibrium is a tuple (μ_s,μ_b) such that

1. Seller optimality: $\Pi(m|\mu_s,\mu_b)\geqslant\Pi(m'|\mu_s,\mu_b)$ for any m in the support of μ_s and any $m'\in\mathcal{M}$.

[12] In connection with footnote 6, existence of an equilibrium requires the selection device to fulfill (6) even on the support of μ_s if possible.

2. Buyer optimality: $U(m|\mu_s, \mu_b) \geqslant U(m'|\mu_s, \mu_b)$ for any m in the support of μ_b and any m' in the support of μ_s.

Constrained efficiency. Consider a social planner who faces the same restrictions as the decentralized economy. In particular, he can choose the trading strategies (μ_s, μ_b), but is subject to the same meeting and information frictions and the same set of feasible trading mechanisms as the decentralized economy. Due to quasi-linear preferences, Pareto optimality is equivalent to maximizing surplus (if lump-sum transfers are available). Therefore, a tuple (μ_s, μ_b) is constrained efficient if the associated RN-derivative λ generates surplus that is larger than the surplus generated by any other (μ_s', μ_b') and associated RN-derivative λ'. That means:

$$s \int_{\mathcal{M}} \left[\sum_n P_n\bigl(\lambda(m)\bigr)\bigl(\pi_n^m + n u_n^m\bigr) \right] d\mu_s \geqslant s \int_{\mathcal{M}} \left[\sum_n P_n\bigl(\lambda'(m)\bigr)\bigl(\pi_n^m + n u_n^m\bigr) \right] d\mu_s'.$$

Clearly, if the meeting frictions are reduced or the set of mechanisms is increased, the constraints on the planner become less severe and higher levels of surplus might become possible.

Random visit strategies. Some of our results hinge on the degree of directedness of the visit strategies of the buyers. In equilibrium, relatively more buyers might search for some mechanisms than for others, or they may search equally across all mechanisms. The latter means that the buyer–seller ratio $\mu_b(M)/\mu_s(M) = b/s$ is constant across all sets $M \subset \operatorname{supp} \mu_s$ offered mechanisms, which equivalently means that the RN-derivative $\lambda(m) = b/s$ is constant across all $m \in \operatorname{supp} \mu_s$. We call this "random visit strategies" since in equilibrium the strategies appear random, even though a deviating seller would attract a different amount of buyers.

2.2. Analysis

We will show that there is an equilibrium in fixed prices in this setting, and also every other class of mechanisms that allows the surplus to be shifted between buyers and sellers (e.g., by means of a reserve price in an auction) includes mechanisms that constitute an equilibrium.

Consider first the problem of an individual seller. In equilibrium, he maximizes (5) knowing that his queue length is determined according to (6). This holds off the equilibrium path by assumption and on the equilibrium path by the second equilibrium condition. Therefore, in equilibrium his choice of mechanism and the resulting queue length solve:

$$\max_{\lambda \in \mathbb{R}^+,\, m \in \mathcal{M}} \sum_{n=1}^{\infty} P_n(\lambda) \pi_n^m \tag{7}$$

$$\text{s.t.} \quad \sum_{n=1}^{\infty} Q_n(\lambda) u_n^m \leqslant U^\star, \quad \text{with equality if } \lambda > 0, \tag{8}$$

where $U^\star = \sup_{m \in \operatorname{supp} \mu_s} U(m|\mu_s, \mu_b)$ represents the utility that buyers can get at other sellers. The constraint represents (6), and so the program clearly reflects the equilibrium conditions when (6) determines the buyer–seller ratio uniquely. If this is not the case, the mechanism has to give the queue length that is best for the seller. The reason is that the seller always has the option to deviate and choose a mechanism that implements the desired queue length uniquely and gives at least the same revenues, for example by choosing a fixed price.[13]

Assume that the space of mechanisms includes mechanisms such that any u_n^m and π_n^m with $nu_n^m + \pi_n^m \leqslant 1$ is feasible.[14] For a given $U^\star$ we call a mechanism m a full-trade mechanism if $nu_n^m + \pi_n^m = 1$. That is, the mechanism induces trade with certainty.

Lemma 1. *Given $U^\star$, any mechanism $m \in \mathcal{M}$ that is not a full-trade mechanism is revenue-dominated by a full-trade mechanism $m' \in \mathcal{M}$ with queue $\lambda(m')$.*

Proof. Mechanism m has $nu_n^m + \pi_n^m < 1$. Let m' give identical utilities to buyers $u_n^{m'} = u_n^m$ but different profit $\pi_n' = 1 - nu_n^m + \pi_n^m$ to the seller. Clearly, buyers obtain identical payoffs and thus $\lambda(m') = \lambda(m)$ is feasible, but the seller achieves weakly higher profits when n buyers arrive. If $P_n(\lambda(m)) > 0$, the full-trade mechanism yields strictly higher profits. □

Therefore, for equilibrium play we only have to restrict attention to full-trade mechanisms. For the following proposition, let with slight abuse of notation $\Pi(m, \lambda) = \sum_{n=1}^{\infty} P_n(\lambda)\pi_n^m$ be the expected seller payoff for a mechanism m with feasible queue λ according to (8) and let $U(m, \lambda) = \sum_{n=1}^{\infty} Q_n(\lambda)u_n^m$ the expected buyer utilities.

Proposition 1 *(Equivalence result). For given $U^\star$, consider some full-trade mechanism $m \in \mathcal{M}$ with feasible $\lambda > 0$. Any other full-trade mechanism m' with $U(m', \lambda) = U(m, \lambda)$ achieves $\Pi(m', \lambda) = \Pi(m, \lambda)$.*

Proof. The queue length λ satisfying (8) for mechanism m means that $U(m, \lambda) = U^\star$. Since $U(m', \lambda) = U(m, \lambda)$, λ also solves (8) for mechanism m'. Since m is a full-trade mechanism, we have $1 - nu_n^m = \pi_n^m$ and therefore

$$\begin{aligned}\Pi(m, \lambda) = \sum_{n=1}^{\infty} P_n(\lambda)\pi_n^m &= \sum_{n=1}^{\infty} P_n(\lambda)\big[1 - nu_n^m\big] \\ &= 1 - P_0(\lambda) - \sum_{n=1}^{\infty} P_n(\lambda) n u_n^m \\ &= 1 - P_0(\lambda) - \lambda \sum_{n=1}^{\infty} Q_n(\lambda)u_n^m \\ &= 1 - P_0(\lambda) - \lambda U(m, \lambda),\end{aligned}$$

where the third line follows from (1). Similarly, since m' is also a full-trade mechanism we have $\Pi(m', \lambda) = 1 - P_0(\lambda) - \lambda U(m', \lambda)$. Since $U(m', \lambda) = U(m, \lambda)$ we have $\Pi(m', \lambda) = \Pi(m, \lambda)$. □

We call a class of mechanisms payoff-complete if for any $\lambda > 0$ and any U there exists a mechanism such that $U(m, \lambda) = U$. Clearly the class of price posting mechanisms is complete,

[13] It can be shown that under price posting (6) selects a queue length uniquely, and for a given queue length it yields higher profits than any other mechanism. See Proposition 1.

[14] Any sequence u_n^m and π_n^m can be implemented by a mechanism that charges price p_n to all buyers, and trade occurs with probability α_n (in which case one of the buyers is selected at random): $p_n = \pi_n^m / n$ and $\alpha_n = n[u_n^m + p - 1]$.

since $U(p,\lambda) = \sum_{n=1}^{\infty} Q_n(\lambda)[1-p]/n$ and $Q_1(\lambda) > 0$ [as $P_1(\lambda) > 0$]. Similarly, the class of second price auctions with reserve price is complete: since $Q_1(\lambda)$ is positive, the reserve price – which can be negative – can be adjusted to yield the right payoff to buyers. Also, the set of mechanisms that offer with probability γ a fixed price and with probability $(1-\gamma)$ give the good away for free is complete. Proposition 1 implies that for any $U^\star$, if there is an optimal mechanism, there is an optimal mechanism within any payoff-complete class of mechanisms.

Proposition 2. *There exists an equilibrium in any class of payoff-complete full-trade mechanisms, and the equilibrium is constrained efficient. It remains an equilibrium and remains constrained efficient even if additional mechanisms become available. The expected payoffs are identical in every equilibrium as long as sellers compete in a class of payoff-complete full-trade mechanisms.*

Proof. We will prove existence in the space of second-price auctions with reserve price. The second and third statements follow immediately from our equivalence Proposition 1: the same expected payoffs yield an equilibrium in any other payoff-complete class of full-trade mechanisms such as price posting. For a given reserve r the seller gets $\pi_1 = r$ and the buyer $u_1 = 1-r$ if one buyer arrives. If more buyers arrive the seller gets $\pi_n = 1$ and the buyers get $u_n = 0$. Payoffs with reserve r are therefore given by $\Pi(r,\lambda) = P_1(\lambda)r + (1 - P_1(\lambda) - P_2(\lambda))$ and $U(r\,\lambda) = Q_1(\lambda)(1-r)$. For a given $U^\star$, the first-order stochastic dominance condition (2) implies that under the optimal λ constraint (8) binds. Sellers therefore maximize

$$\max_r P_1(\lambda)r + \bigl(1 - P_0(\lambda) - P_1(\lambda)\bigr)$$
$$\text{s.t.} \quad \frac{P_1(\lambda)}{\lambda}(1-r) = U^\star.$$

Substituting out the constraint leaves

$$\max_\lambda 1 - \lambda U^\star - P_0(\lambda).$$

Since P_0 is strictly convex, each seller has a unique optimum, and therefore an equilibrium has all sellers posting the same reserve price r. It is characterized by the unique first-order condition

$$-P_0'(\lambda) = U^\star.$$

Since all sellers in equilibrium post the same r, they will face a queue length of $\lambda^\star = b/s$. Therefore equilibrium utility $U^\star = -P_0'(\lambda^\star)$, yielding equilibrium profits of $1 + \lambda P_0'(\lambda^\star) - P_0(\lambda^\star)$, which means that the equilibrium reserve price is

$$r^\star = 1 + \frac{\lambda^\star P_0'(\lambda^\star)}{P_1(\lambda^\star)}.$$

Efficiency follows because in equilibrium visit strategies are random. We elaborate on that point below. □

The reserve price $r^\star$ can in general be different from zero. Therefore, Peters' [24] conjecture that the reserve price in the homogeneous agent case may lie in the open set $(0, 1)$ obtains for many meeting specifications. For the specific urn-ball meeting technology of the directed search literature that Peters [24] considers this is not true, though.

Corollary 1. *Under urn-ball meetings* ($P_n(\lambda) = \frac{\lambda^n e^{-\lambda}}{n!}$) *the equilibrium second price auction has a reserve price* $r^\star = 0$.

Together with the previous proposition this clarifies why under urn-ball meetings competition in fixed prices leads to the same surplus as in the case when sellers do not post anything but simply run second price auctions without reserve [15]: If sellers could compete in auctions, they would indeed choose to set a zero reserve under urn-ball meetings. Different meeting technologies would lead to a different reserve price, though. While we have compared different equilibria so far, the combination of Proposition 1 and Proposition 2 readily yields the following multiplicity of mechanisms within the same equilibrium. It arises when sellers have alternative mechanisms available that generate the same expected surplus for the buyers, such as fixed prices and second price auctions.

Corollary 2. *If the mechanism space contains two disjoint classes of payoff-complete full-trade mechanisms, then there exists a continuum of equilibria in which sellers announce different payoff-equivalent mechanisms.*

Finally, we relate our findings to the constrained efficiency of random visit strategies.

Corollary 3. *Since the equilibrium is in random visit strategies, it is constrained efficient in the sense of creating the highest number of trades at mechanisms that have full trade.*

To see the constrained efficiency result, observe that a necessary requirement for constrained efficiency is that sellers offer full trade mechanisms. Let $\mathcal{M}^F$ be the set of full-trade mechanisms, which includes all fixed-price mechanisms. Full-trade ensures that every time a seller is present, the good is exchanged. Apart from this, constrained efficiency then means to maximize the number of trades. Observe that the number of trades is maximized under

$$\max_{\mu_s,\mu_b} s \int \left[1 - P_0\bigl(\lambda(m)\bigr)\right] d\mu_s$$
$$\text{s.t.} \quad \lambda = d\mu_b/d\mu_s, \quad \mu_s\bigl(\mathcal{M}^F\bigr) = s, \quad \mu_b\bigl(\mathcal{M}^F\bigr) \leqslant b.$$

The strict concavity of $[1 - P_0(\lambda)]$ immediately implies that the solution to this program has $\lambda(m) = b/s$ constant for all mechanisms in the support of μ_s, i.e., the efficient allocation features random visit strategies.

Note again why in equilibrium all sellers always have the same queue length, no matter which class of payoff-complete full-trade mechanisms they compete in. Suppose this were not the case and there exist at least two sets of firms that face different queue lengths in equilibrium. From revenue equivalence, these firms could post second price auctions, possibly with different reserve prices. But because there is a unique solution to the first-order condition of a seller posting an auction (see proof of Proposition 2), and given concavity of the seller's profit function, there can only be one optimal auction and associated reserve price. By revenue equivalence, at least one of the mechanisms that the firms initially announced does not maximize profits, therefore contradicting that firms face different queue lengths.

Therefore, whenever the market utility assumption holds in the limit of finite economies as the population size grows large, the same expected equilibrium payoffs arise no matter the class of mechanisms in which the sellers compete (as long as the class is payoff-complete). This result highlights why competition in prices (e.g. the limit in Burdett, Shi and Wright [4]) yields the same expected profit as competition in auctions (see e.g. Julien, Kennes and King [12]). The market utility assumption holds for particular price posting games [22,23], and holds for competition in more general mechanisms as long as sellers compete in a bounded payoff space [34]. Of course, this equivalence in expected payoffs holds only in the limit. In a finite economy, changes in the announcement of an individual seller changes the market utility of the buyers due to market power, and multiple equilibria are possible.[15]

3. Heterogeneous buyers

Consider the extended model with two buyer types: measure $\underline{b}$ of low type buyers and measure $\bar{b}$ of high type buyers. As before, there is a measure s of homogeneous sellers. The low buyer type has a valuation $\underline{v} > 0$ for the good, the high buyer type has valuation $\bar{v} > \underline{v}$, and sellers have no value for the good or cost of production. Again, utilities are linear in the amount of money that is transferred, and payoffs from no trade are normalized to zero.

The market interaction. The market interaction is essentially unchanged from the previous section. $\mu_s(M)$ still denotes the measure of sellers that post mechanisms in set $M \subset \mathcal{M}$. The trading decisions by buyers are now summarized by two measures, $\underline{\mu}_b$ and $\bar{\mu}_b$, for the low and high buyer types, respectively. Both are required to be absolute continuous in μ_s. Now the buyer–seller ratios for each type are $\underline{\lambda} = d\underline{\mu}_b/d\mu_s$ and $\bar{\lambda} = d\bar{\mu}_b/d\mu_s$, and the total buyer–seller ratio is $\lambda = \underline{\lambda} + \bar{\lambda}$.

Meetings. Meetings happen exactly as in the previous section, i.e., the definitions of the probabilities $P_n(\lambda)$ and $Q_n(\lambda)$ are unchanged. We also assume that the type of the buyer does not affect his chances of meeting a seller. It might affect his probability of trade since the mechanism might distinguish between types, but the probability of getting to the mechanism is type-independent.[16] Therefore, from the seller's perspective, the probability of being in a match with $\underline{n}$ low and $\bar{n}$ high types is $\tilde{P}_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda}) = P_{\underline{n}+\bar{n}}(\underline{\lambda} + \bar{\lambda})B_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda})$ where $B_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda})$ is the binomial probability of drawing $\underline{n}$ low types out of all $\underline{n} + \bar{n}$ buyers, when the probability of drawing one low type is $\underline{\lambda}/(\underline{\lambda} + \bar{\lambda})$. For a buyer who meets a seller, the probability of being in a match with $\underline{n} + \bar{n}$ *other* buyers of which $\underline{n}$ are low types and $\bar{n}$ high types is $\tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda}) = Q_{\underline{n}+\bar{n}+1}(\underline{\lambda} + \bar{\lambda})B_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda})$.

Mechanisms. A mechanism now specifies for a given number $\underline{n}$ of low type buyers and $\bar{n}$ of high type buyers the expected payoff $\pi^m_{\underline{n},\bar{n}}$ for the seller and $\underline{u}^m_{\underline{n},\bar{n}}$ ($\bar{u}^m_{\underline{n},\bar{n}}$) for low (high) valuation buyers. The payoffs have to be implemented in a way such that buyers are willing to truthfully reveal their type. To specify this constraint, it is useful to decompose the payoffs into the probability of trade $\underline{x}^m_{\underline{n},\bar{n}}$ ($\bar{x}^m_{\underline{n},\bar{n}}$) for low (high) valuation buyers when $\underline{n}$ low valuation and $\bar{n}$ high valuation buyers are present. Clearly probabilities cannot add to more than unity, i.e. $\underline{n}\underline{x}^m_{\underline{n},\bar{n}} + \bar{n}\bar{x}^m_{\underline{n},\bar{n}} \leqslant 1$, and the seller cannot give out more surplus than is created

$$\pi^m_{\underline{n},\bar{n}} + \underline{n}\underline{u}^m_{\underline{n},\bar{n}} + \bar{n}\,\bar{u}^m_{\underline{n},\bar{n}} \leqslant \underline{n}\underline{x}^m_{\underline{n},\bar{n}}\underline{v} + \bar{n}\bar{x}^m_{\underline{n},\bar{n}}\bar{v}, \tag{9}$$

[15] For a finite agent model with a continuum of equilibria, including asymmetric ones where the queue length of one firm is larger than the other, see Coles and Eeckhout [5]. See also Galenianos, Kircher and Virag [9] for an elaboration on market power in finite economies.

[16] We do not model different search intensities for the buyers. If differences in search intensity were introduced, the number of buyers in the market would have to be weighted by their respective search intensity.

or in ex-ante terms

$$\sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} P_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\big[\pi^m_{\underline{n},\bar{n}} + \underline{n}\underline{u}^m_{\underline{n},\bar{n}} + \bar{n}\bar{u}^m_{\underline{n},\bar{n}}\big] \leqslant \sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} P_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\big[\underline{n}\underline{x}^m_{\underline{n},\bar{n}}\underline{v} + \bar{n}\bar{x}^m_{\underline{n},\bar{n}}\bar{v}\big]. \tag{10}$$

Incentive compatibility constraints require that a low type buyer is willing to reveal his type, i.e., the payoff he receives is higher than the payoff when pretending to be a high type. Ex-ante incentive compatibility is therefore:

$$\sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\underline{u}^m_{\underline{n}+1,\bar{n}} \geqslant \sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\big[\bar{u}^m_{\underline{n},\bar{n}+1} + \bar{x}^m_{\underline{n},\bar{n}+1}(\underline{v}-\bar{v})\big]. \tag{11}$$

The left-hand side is the payoff when announcing to be a low type. The right-hand side is the payoff when pretending to be the high type. The right-hand side comprises the payoff that the high type would obtain, adjusted for the fact that the low type only obtains $\underline{v}$ instead of $\bar{v}$ when he trades. Similarly, the high type is willing to reveal his type if

$$\sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\big[\underline{u}^m_{\underline{n}+1,\bar{n}} + \underline{x}^m_{\underline{n}+1,\bar{n}}(\bar{v}-\underline{v})\big] \leqslant \sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\bar{u}^m_{\underline{n},\bar{n}+1}. \tag{12}$$

Here the right-hand side is the payoff for the high type, and the left-hand side is the payoff of the low type adjusted for the fact that the high type obtains more when trading.

Payoffs. Given the trading strategies of all other agents, the payoffs from choosing $m \in \operatorname{supp}\mu_s$ for an individual seller and a low and high type buyer, respectively, are

$$\Pi(m|\mu_s,\underline{\mu}_b,\bar{\mu}_b) = \sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} P_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\pi^m_{\underline{n},\bar{n}}, \tag{13}$$

$$\underline{U}(m|\mu_s,\underline{\mu}_b,\bar{\mu}_b) = \sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\underline{u}^m_{\underline{n}+1,\bar{n}}, \tag{14}$$

$$\bar{U}(m|\mu_s,\underline{\mu}_b,\bar{\mu}_b) = \sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda},\bar{\lambda})\bar{u}^m_{\underline{n},\bar{n}+1}. \tag{15}$$

The buyer–seller ratio for a deviating seller is again given by the indifference of the buyers. The deviating seller expects $(\underline{\lambda}(m),\bar{\lambda}(m))$. His expectation can only include $\underline{\lambda}(m) > 0$ if buyers are indeed willing to come to his mechanism rather than to the mechanisms offered on the equilibrium path:

$$\sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}\big(\underline{\lambda}(m),\bar{\lambda}(m)\big)\underline{u}^m_{\underline{n}+1,\bar{n}} = \max_{m\in\mathcal{M}} \underline{U}(m|\mu_s,\underline{\mu}_b,\bar{\mu}_b). \tag{16}$$

Similarly, he can expect $\bar{\lambda}(m) > 0$ only if

$$\sum_{\underline{n}=0}^{\infty}\sum_{\bar{n}=0}^{\infty} \tilde{Q}_{\underline{n},\bar{n}}\big(\underline{\lambda}(m),\bar{\lambda}(m)\big)\bar{u}^m_{\underline{n},\bar{n}+1} = \max_{m\in\mathcal{M}} \bar{U}(m|\mu_s,\underline{\mu}_b,\bar{\mu}_b). \tag{17}$$

Similar to the previous section, for most mechanisms including price posting or standard auctions the pair $(\underline{\lambda}(m), \overline{\lambda}(m))$ that satisfies (16) and (17) is unique. For mechanisms where multiplicity is possible, we again follow McAfee [17] and others and assume that the seller believes that he can coordinate buyers and choose the combination of queue length that is most desirable to him.

Equilibrium. An equilibrium is a tuple $(\mu_s, \underline{\mu}_b, \overline{\mu}_b)$ such that

1. Seller optimality: Any $m \in \operatorname{supp} \mu_s$ solves $\max_{m' \in \mathcal{M}} \Pi(m' | \mu_s, \underline{\mu}_b, \overline{\mu}_b)$.
2. Low buyer optimality: Any $m \in \operatorname{supp} \underline{\mu}_b$ solves $\max_{m' \in \operatorname{supp} \mu_s} \underline{U}(m' | \mu_s, \underline{\mu}_b, \overline{\mu}_b)$.
3. High buyer optimality: Any $m \in \operatorname{supp} \overline{\mu}_b$ solves $\max_{m' \in \operatorname{supp} \mu_s} \overline{U}(m' | \mu_s, \underline{\mu}_b, \overline{\mu}_b)$.

Constrained efficiency. Consider again a planner who can determine the trading strategies, but is constrained by the same meeting frictions and available trading mechanisms as the decentralized equilibrium. A tuple $(\mu_s, \underline{\mu}_b, \overline{\mu}_b)$ is constrained efficient if the resulting RN-derivatives $\underline{\lambda}(m)$ and $\overline{\lambda}(m)$ generate a higher surplus $s \int_{\mathcal{M}} [\sum_{\underline{n}=0}^{\infty} \sum_{\overline{n}=0}^{\infty} P_{\underline{n},\overline{n}}(\underline{\lambda}, \overline{\lambda})(\pi_{\underline{n},\overline{n}}^m + \underline{n}\underline{u}_{\underline{n},\overline{n}}^m + \overline{n}\overline{u}_{\underline{n},\overline{n}}^m)]\, d\mu_s$ than the respective surplus under any other tuple $(\mu_s', \underline{\mu}_b', \overline{\mu}_b')$ and its resulting RN-derivatives.

Preliminaries. Similar to the derivation of maximization problem (7) in the homogeneous buyer case, an individual seller now anticipates that his queue length arises according to (16) and (17). He maximizes his expected profits, understanding that the number of buyers is governed by (16) and (17). Given the market utility $\underline{U}^\star = \max_{m \in \mathcal{M}} \underline{U}(m | \mu_s, \underline{\mu}_b, \overline{\mu}_b)$ and $\overline{U}^\star = \max_{m \in \mathcal{M}} \overline{U}(m | \mu_s, \underline{\mu}_b, \overline{\mu}_b)$, this amounts to the problem

$$\max_{(\underline{\lambda}, \overline{\lambda}) \in \mathbb{R}_+^2,\, m \in \mathcal{M}} \sum_{\underline{n}=0}^{\infty} \sum_{\overline{n}=0}^{\infty} P_{\underline{n},\overline{n}}(\underline{\lambda}, \overline{\lambda}) \pi_{\underline{n},\overline{n}}^m \tag{18}$$

such that

$$\sum_{\underline{n}=0}^{\infty} \sum_{\overline{n}=0}^{\infty} \tilde{Q}_{\underline{n},\overline{n}}(\underline{\lambda}, \overline{\lambda}) \underline{u}_{\underline{n}+1,\overline{n}}^m \leqslant \underline{U}^\star, \quad \text{with equality if } \underline{\lambda} > 0, \tag{19}$$

$$\sum_{\underline{n}=0}^{\infty} \sum_{\overline{n}=0}^{\infty} \tilde{Q}_{\underline{n},\overline{n}}(\underline{\lambda}, \overline{\lambda}) \overline{u}_{\underline{n},\overline{n}+1}^m \leqslant \overline{U}^\star, \quad \text{with equality if } \overline{\lambda} > 0, \tag{20}$$

and incentive compatibility constraints for truthful type revelation (11) and (12) and the resource constraint (9) have to hold. The constraints ensure thus that (16) and (17) indeed apply.[17] This formulation resembles the standard mechanism design approach: Buyers offer transfers and trading probabilities such that types are revealed and participation constraints are met, only now the level of the participation constraint is endogenous.

[17] Again, the buyer–seller ratios are choice variables even if the constraints permit several solutions because we assumed that off the equilibrium path the seller can coordinate the buyers. On the equilibrium path, this still holds if the seller has deviations available to him that allow such a coordination off the equilibrium path, e.g. when other mechanisms achieve the same expected payoffs but implement the desired buyer–seller ratios uniquely. It can be shown that a combination of auctions and prices does indeed achieve this.

4. Price posting mechanisms

Consider first the set of equilibria in a price posting environment. Under price posting, the good is allocated indiscriminately and with equal probability to any one of the buyers that arrives. The surplus then is

$$\pi^m_{\underline{n},\overline{n}} + \underline{n}\underline{u}^m_{\underline{n},\overline{n}} + \overline{n}\overline{u}^m_{\underline{n},\overline{n}} = \frac{\underline{n}}{\underline{n}+\overline{n}}\underline{v} + \frac{\overline{n}}{\underline{n}+\overline{n}}\overline{v} \leqslant \underline{n}\underline{v} + \overline{n}\overline{v}.$$

Full surplus is only realized if $\underline{n}$ or $\overline{n}$ is equal to zero. A price p induces payoffs $\pi_{\underline{n},\overline{n}}(p) = p$ if either $\underline{n}$ or $\overline{n}$ or both are strictly positive, and one of the buyers obtains the good at random so that $\underline{u}_{\underline{n},\overline{n}}(p) = \frac{\underline{v}-p}{\underline{n}+\overline{n}}$ and $\overline{u}^m_{\underline{n},\overline{n}} = \frac{\overline{v}-p}{\underline{n}+\overline{n}}$. Taking $\underline{U}^\star$ and $\overline{U}^\star$ as given, we can consider (18) for an individual firm. For price posting mechanisms the maximization becomes particularly simple: the seller only cares about the probability of trade times the price $[1 - P_0(\underline{\lambda} + \overline{\lambda})]p$, while the buyers only care about their probability of trade times their gain $\sum_{n=1}^{\infty} Q_n(\underline{\lambda} + \overline{\lambda})\frac{v-p}{n}$. Using $nP_n(\lambda) = \lambda Q_n(\lambda)$ from (1) we get $\sum_{n=1}^{\infty} Q_n(\underline{\lambda} + \overline{\lambda})/n = [1 - P_0(\underline{\lambda} + \overline{\lambda})]/(\underline{\lambda} + \overline{\lambda})$ and therefore

$$\max_{(\underline{\lambda},\overline{\lambda})\in\mathbb{R}^2_+,\, p\in\mathbb{R}_+} \big[1 - P_0(\underline{\lambda} + \overline{\lambda})\big]p \tag{21}$$

such that

$$\frac{1 - P_0(\underline{\lambda} + \overline{\lambda})}{\underline{\lambda} + \overline{\lambda}}[\underline{v} - p] \leqslant \underline{U}^\star, \quad \text{with equality if } \underline{\lambda} > 0, \tag{22}$$

$$\frac{1 - P_0(\underline{\lambda} + \overline{\lambda})}{\underline{\lambda} + \overline{\lambda}}[\overline{v} - p] \leqslant \overline{U}^\star, \quad \text{with equality if } \overline{\lambda} > 0. \tag{23}$$

Note that in this case we can omit the truth-telling constraint as the seller will always choose to trade with the buyer type where he can achieve the highest queue length $\lambda = \underline{\lambda} + \overline{\lambda}$ and the other type then is either indifferent or strictly prefers not to appear. We first show

Lemma 2. *For given $\underline{U}^\star$ and $\overline{U}^\star$, consider the restricted problem* (21) *where the price is fixed and only* $(\underline{\lambda}, \overline{\lambda})$ *are chosen optimally*: *There exists $\hat{p}$ such that at prices below $\hat{p}$ the solution has $\overline{\lambda} = 0$ and only low types are attracted, while at prices above $\hat{p}$ the solution has $\underline{\lambda} = 0$ and only high types are attracted. Moreover, in the full problem* (21) *where the price is also chosen optimally, setting a price of $\hat{p}$ is never optimal, and a seller restricted to the set of prices* $[0, \hat{p}]$ (*or restricted to* $[\hat{p}, \infty)$) *has a unique optimal price within that set.*

Proof. Consider the restricted problem (21) where the price p is given. As long $p < \overline{v}$, the optimal queue lengths maximize $\lambda = \underline{\lambda} + \overline{\lambda}$ such that at least one constraint is still met at equality. If the constraint holds with equality we have $\frac{(1-P_0(\lambda))}{\lambda}[v - p] = C$ for some C which implicitly defines λ, and implicit differentiation gives

$$\frac{\partial \lambda}{\partial p} = -\frac{(1 - P_0(\lambda))\lambda}{(1 - P_0(\lambda) + \lambda P_0'(\lambda))(v - p)},$$

which is negative and strictly increasing in v. This single crossing property immediately implies that there exists some price $\hat{p}$ such that at price $p < \hat{p}$ the optimal λ has the low type's constraint (22) binding, while at $p > \hat{p}$ the optimal λ has the high type's constraint (23) binding. It also implies that no seller would like to cater to both low and high types by offering $\hat{p}$: if it is profitable to increase one's price up to $\hat{p}$ even though the queue length is falling a lot when trading with low types, it will clearly be profitable to increase it strictly above $\hat{p}$ because there the queue length is falling less drastically due to a price change. We will make this argument precise below.

We know that a seller who sells at a price below $\hat{p}$ will have a queue length determined by (22) holding with equality. Substituting the constraint, the optimal utility of such a seller is given by the optimization problem

$$\max_{\underline{\lambda}\in[\hat{\lambda},\infty)} \left[1 - P_0(\underline{\lambda})\right]\underline{v} - \underline{\lambda}\underline{U}^\star, \tag{24}$$

where $\hat{\lambda}$ is the queue length at $\hat{p}$ if (22) holds with equality. This problem is strictly concave. Therefore every buyer will choose the same queue length $\underline{\lambda}^\star(\underline{U}^\star, \overline{U}^\star)$ dependent on $\underline{U}^\star$ (and associated price) when trading with low types. The dependence on $\overline{U}^\star$ only arises because it might affect $\hat{\lambda}$. Similarly, the optimization problem for a seller that considers trading with high buyer types is

$$\max_{\overline{\lambda}\in[0,\hat{\lambda}]} \left[1 - P_0(\overline{\lambda})\right]\overline{v} - \overline{\lambda}\overline{U}^\star \tag{25}$$

which is also strictly concave. Therefore every buyer will choose the same queue length $\overline{\lambda}^\star(\underline{U}^\star, \overline{U}^\star)$ dependent on $\overline{U}^\star$ and $\underline{U}^\star$ (and associated price) when trading with low types.

It is easy to see that it is not optimal to trade at $\hat{\lambda}$. For an individual firm, trading at $\hat{\lambda}$ is only optimal if profits at higher queue length according to (24) are not profitable, i.e. the right derivative has to be $-P_0'(\hat{\lambda})\underline{v} - \underline{U}^\star \leqslant 0$. By a similar logic the left derivative of (25) has to be positive, i.e. $-P_0'(\hat{\lambda})\overline{v} - \overline{U}^\star \geqslant 0$, which together imply $-P_0'(\hat{\lambda})[\overline{v} - \underline{v}] \geqslant \overline{U}^\star - \underline{U}^\star$. Since $\hat{\lambda}$ is the point where both (22) and (23) hold, we have $\overline{U}^\star - \underline{U}^\star = [\overline{v} - \underline{v}](1 - P_0(\hat{\lambda}))/\hat{\lambda}$. Therefore the prior inequality becomes after some rearranging $1 - P_0(\hat{\lambda}) + \hat{\lambda}P_0'(\hat{\lambda}) \leqslant 0$, which delivers a contradiction. Therefore no seller caters to both buyer types. □

This leads to the following result:

Proposition 3. *Assume the mechanism space $\mathcal{M}$ includes fixed price mechanisms only. A unique equilibrium exists with one "market" for each buyer type that trades: If the low types trade in equilibrium, then exactly two prices are offered in equilibrium and all low type buyers trade at the low price and all high type buyers trade at the high price. If only high type sellers trade in equilibrium, only one price is offered at which they trade.*

Proof. Since $\hat{\lambda}$ is never an optimal choice for an individual seller, and neither zero nor infinity can be optimal choices, in equilibrium the optimal choice of a firm has to be characterized by the first-order condition. For (24) this is

$$-P_0'(\underline{\lambda})\underline{v} = \underline{U}^\star.$$

Since all sellers will choose the same queue length, the buyer–seller ratio is $\underline{\lambda} = \frac{b}{\gamma s}$ when γ is the fraction of sellers that trade with low types. With slight abuse of notation let $\underline{\lambda}^\star(\gamma) = \frac{b}{\gamma s}$. Profits are then

$$\underline{\pi}(\gamma) = \left[1 - P_0\big(\underline{\lambda}^\star(\gamma)\big) + \underline{\lambda}^\star(\gamma)P_0'\big(\underline{\lambda}^\star(\gamma)\big)\right]\underline{v}. \tag{26}$$

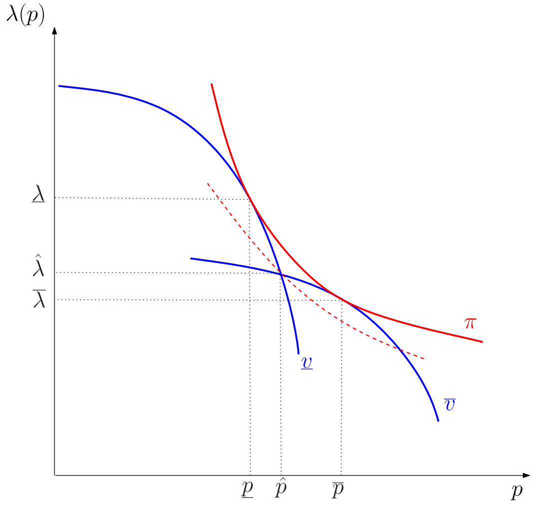

Fig. 1. The separation of buyer types.

Similar logic for high types implies $\overline{U}^\star = -P_0'(\bar{\lambda}^\star(1-\gamma))\bar{v}$ and $\bar{\lambda}^\star(1-\gamma) = \frac{\bar{b}}{(1-\gamma)s}$ and

$$\bar{\pi}(\gamma) = \left[1 - P_0\left(\bar{\lambda}^\star(1-\gamma)\right) + \bar{\lambda}^\star(1-\gamma)P_0'\left(\bar{\lambda}^\star(1-\gamma)\right)\right]\bar{v}. \tag{27}$$

In equilibrium it cannot be that only the low type buyers trade, as then $\overline{U}^\star = 0$ and it would be optimal to trade with the high types. If $\underline{\pi}(0) \leqslant \bar{\pi}(0)$ then even if it is most attractive to trade with low types it is not worthwhile and therefore in equilibrium $\gamma^\star = 0$. Otherwise the equilibrium has $\gamma^\star$ uniquely determined by $\underline{\pi}(\gamma^\star) = \bar{\pi}(\gamma^\star)$. This uniquely characterizes the equilibrium. □

The separation result is illustrated in Fig. 1. Indifference curves satisfy single crossing as derived in the proof of Lemma 2. For high valuation buyers $\bar{v}$, the IC is everywhere flatter than that of low valuation buyers $\underline{v}$. All sellers are identical and in equilibrium they will obtain equal profits catering to both types of buyers. Observe that if the price $\hat{p}$ is offered to both types of buyers, a seller has an incentive to deviate by offering either a strictly higher (or strictly lower) price, thereby catering exclusively to the high type (low type) buyers. Such a deviation is profitable since the isoprofit curve at $\hat{p}$ (dashed line) is not tangent, and maintaining the same utility level for at least one of the buyers, a seller can achieve a position on a strictly higher isoprofit. This continues to be the case as long as the isoprofit curve is not tangent to both indifference curves simultaneously. In equilibrium, a seller makes equal profits from both types of buyers, and the buyer types have strict preferences over which (p, λ) pair to choose. High type buyers prefer high prices and low queue lengths, and low type buyers prefer low prices and high queue lengths.

It is interesting to note that profit equations (26) and (27) fulfill the Hosios [11] condition, which is important for constrained efficiency in search models, because it gives agents the correct incentives to enter one market rather than another. Writing profits as $\pi = (1 - P_0)(1 - (-\lambda P_0'/(1 - P_0)))v$ shows that conditional on matching the share of the surplus that goes to the seller is equal to one minus the elasticity of the matching function (and full-trade ensures that trade takes place whenever at least one buyer is present). The matching function is the probability $1 - P_0(\lambda)$ of being able to trade. The Hosios condition induces constrained efficiency as long as we only consider non-screening mechanisms – but can sometimes be too weak a notion to guarantee constrained efficiency when larger classes of mechanisms are allowed. This depends on the meeting technology as we discuss in the following sections.

To establish such constrained efficiency in the class of non-screening mechanisms formally, we consider a planner who takes the frictions as given and can only use mechanisms that give the good to one of the buyers at random. We will show that even if the planner assigns to each seller i its personalized mechanism with queue lengths $\underline{\lambda}^p(i)$ and $\overline{\lambda}^p(i)$ of low and high buyer types and even if the planner does not have to consider any incentive constraints, he will not achieve a higher surplus than the decentralized economy. Here i can be viewed as the name of the mechanism of this seller. These functions $\underline{\lambda}^p(i)$ and $\overline{\lambda}^p(i)$ are feasible if they are measurable and indeed reflect the ratio of buyers to sellers in the sense that

$$\int_0^s \underline{\lambda}^p(i)\,di \leqslant \underline{b} \quad \text{and} \quad \int_0^s \overline{\lambda}^p(i)\,di \leqslant \bar{b}. \tag{28}$$

The realized surplus is

$$S(\underline{\lambda}^p, \overline{\lambda}^p) = \int_0^s \left[1 - P_0\left(\underline{\lambda}^p(i) + \overline{\lambda}^p(i)\right)\right]\left[\frac{\underline{\lambda}^p(i)\underline{v} + \overline{\lambda}^p(i)\overline{v}}{\underline{\lambda}^p(i) + \overline{\lambda}^p(i)}\right] di, \tag{29}$$

where the first brackets reflect the probability that the seller has a buyer, in which case his probability of selecting a high type is $\overline{\lambda}^p(i)/(\underline{\lambda}^p(i) + \overline{\lambda}^p(i))$ and $\overline{v}$ is realized, and with the complementary probability $\underline{v}$ is realized.

The price posting equilibrium in which a fraction γ caters to the low valuation sellers can be represented as a planner's strategy $\underline{\lambda}^p(i) = \underline{\lambda}^\star(\gamma)$ for all $i \in [0, \gamma s]$ and zero otherwise, and $\overline{\lambda}^p(i) = \overline{\lambda}^\star(\gamma)$ for $i \in (\gamma s, s]$ and zero otherwise. We will show that this "price posting" assignment is constrained efficient in the space of mechanisms that give away the good at random.

Proposition 4. *Consider a mechanism space $\mathcal{M}$ that includes fixed prices and possibly other mechanisms, but only includes non-screening mechanisms. Then the equilibrium when sellers can only compete in fixed prices is constrained efficient even under $\mathcal{M}$. In particular, the realized surplus in the price posting equilibrium is higher than under any other feasible queue length functions $\underline{\lambda}(.)$ and $\overline{\lambda}(.)$.*

Proof. Consider some queue length functions $\underline{\lambda}^p(.)$ and $\overline{\lambda}^p(.)$. We will first show that it is sufficient to concentrate on queue length functions that assign strictly positive $\underline{\lambda}^p(\sigma)$ only if $\overline{\lambda}^p(.) = 0$. To see this, write the surplus as

$$\begin{aligned} S(\underline{\lambda}^p, \overline{\lambda}^p) &= \underline{v}\int_0^s \left[1 - P_0\left(\check{\lambda}(i)\right)\right]\underline{f}(i)di + \overline{v}\int_0^s \left[1 - P_0\left(\check{\lambda}(i)\right)\right]\overline{f}(i)\,di \\ &= \underline{v}\int_0^s \left[1 - P_0\left(\check{\lambda}(i)\right)\right]d\underline{F}(i) + \overline{v}\int_0^s \left[1 - P_0\left(\check{\lambda}(i)\right)\right]d\overline{F}(i), \end{aligned}$$

where $\check{\lambda}(i) = \underline{\lambda}^p(i) + \overline{\lambda}^p(i)$, $\underline{f}(i) = \frac{\underline{\lambda}^p(i)}{\underline{\lambda}^p(i)+\overline{\lambda}^p(i)}$ and $\underline{F}(i) = \int_0^s \underline{f}(i)\,di$, likewise for $\overline{f}, \overline{F}$. The first term can be interpreted as assigning all low valuation buyers to a measure $\underline{F}(s)$ of firms and all high valuation buyers to a measure $\overline{F}(s)$ if firms match according to the assignment function $\check{\lambda}(i)$. Another way of seeing this is to rewrite the expression as

$$S(\underline{\lambda}^p, \overline{\lambda}^p) = \underline{v} \int_0^{\underline{F}(s)} \left[1 - P_0\left(\check{\lambda}\left(\underline{F}(x)\right)\right)\right] dx + \overline{v} \int_{\underline{F}(s)}^{s} \left[1 - P_0\left(\check{\lambda}\left(\overline{F}\left(x - \underline{F}(s)\right)\right)\right)\right] dx$$

which is identical to a feasible assignment function that assigns queue length $\underline{\lambda}(i) = \check{\lambda}(\underline{F}(i))$ to all firms in $[0, \underline{F}(s)]$ and $\underline{\lambda}(i) = 0$ otherwise, and $\overline{\lambda}(i) = \check{\lambda}(\overline{F}(i - \underline{F}(s)))$ to all firms in $[\underline{F}(s), s]$ and $\underline{\lambda}(i) = 0$ otherwise.

Next, observe that $1 - P_0(\lambda)$ is strictly concave, and therefore it is optimal to assign an identical queue length $\underline{\lambda}$ to all firms that sell to low valuation buyers and an identical queue length $\overline{\lambda}$ to all firms that sell to high valuation buyers. If a fraction α of the firms sells to low valuation buyers, those achieve a buyer seller ratio of $\underline{b}/(\alpha s)$. This ratio depends on α, and we will denote it with slight abuse of notation by $\underline{\lambda}(\alpha) = \underline{b}/(\alpha s)$. For the fraction $(1-\alpha)$ of firms that sell to the high types, the buyer–seller ratio is $\overline{b}/(1-\alpha)$. Again, with slight abuse of notation we denote this as $\overline{\lambda}(1-\alpha) = \overline{b}/(1-\alpha)$. The maximization problem of the planner reduces to

$$\max_{\alpha \in [0,1]} \alpha\left(1 - P_0\left(\underline{\lambda}(\alpha)\right)\right)\underline{v} + (1-\alpha)\left(1 - P_0\left(\overline{\lambda}(1-\alpha)\right)\right)\overline{v}.$$

The planner clearly chooses $\alpha < 1$. If he chooses to trade with low valuation buyers, the first-order condition for optimality is

$$\begin{aligned}&\left(1 - P_0\left(\underline{\lambda}(\alpha)\right) + \underline{\lambda}(\alpha) P_0'\left(\underline{\lambda}(\alpha)\right)\right)\underline{v} \\ &\quad = \left(1 - P_0\left(\overline{\lambda}(1-\alpha)\right) + \overline{\lambda}(1-\alpha) P_0'\left(\overline{\lambda}(1-\alpha)\right)\right)\overline{v},\end{aligned} \tag{30}$$

where we used the fact that $(\partial\underline{\lambda}(\alpha)/\partial\alpha) = -\underline{\lambda}(\alpha)/\alpha$ and $(\partial\overline{\lambda}(1-\alpha)/\partial\alpha) = \overline{\lambda}(1-\alpha)/(1-\alpha)$. Note that this is identical to the condition that profits according to (26) and (27) have to be equal. Moreover, it is easy to show that $\alpha = 0$ if the left-hand side of (30) evaluated as $\alpha = 0$ is smaller then the right-hand side of (30) evaluated at $\alpha = 0$, which again coincides with the equilibrium. □

In the following, we will extend the mechanisms space $\mathcal{M}$ to allow for ex-post screening mechanisms such as auctions. This changes the available mechanisms for the sellers and, thus, possibly the nature of the equilibrium. It also relaxes the constraints on the planner, and higher surplus might be possible. In particular we investigate under which meeting technologies the extended mechanism space actually changes equilibria and constrained efficiency, and under which it does not.

5. Competition in mechanisms

5.1. *Purely non-rival meetings*

Here we consider meeting technologies in which a seller can meet multiple buyers simultaneously. More specifically, meetings are purely non-rival in the sense that the meeting probability for a buyer is not affected by the presence of other buyers in the market: the probability $1 - Q_0(\lambda)$ of a meeting is independent of λ. This means in particular that an increase in the number of low types in a market does not diminish the chances for high types of meeting a seller. It might diminish their trading probability e.g., in the case when prices are used and a seller tries to attract both types of buyers. Yet in terms of simply meeting one of the sellers there are no spillovers. This is for example true under urn-ball meetings where buyers choose one seller in their desired market at random. In that case the probability of not meeting any seller is zero. The probability of trading will change when there are more buyers because it is harder to obtain the good, but the probability of meeting is constant. The urn-ball meeting technology is the standard assumption of the competing mechanism design literature [24,25]. The set of non-rival meeting functions is substantially larger than the urn-ball meeting function, though.[18]

Under the non-rival meetings assumption we will prove that an equilibrium in which all sellers use second price auctions with a reserve below $\underline{v}$ and buyers use random visit strategies yields strictly higher surplus than the equilibrium when sellers were restricted to use prices only. With slight modification it can be shown that any class of mechanisms that does not allow for ex-post screening will not be constrained efficient. The result obtains even though we have seen that sellers can perfectly screen between buyers by using prices, and therefore no seller in a price posting environment has any uncertainty about the type of buyer he is facing. Note that under purely non-rival meetings, the probability of having a match with n other high type depends only on $\overline{\lambda}$ and is $Q_n(\overline{\lambda})$.

Proposition 5. *Consider non-rival meetings. The equilibrium in a price posting environment, i.e., an environment where $\mathcal{M}$ only includes fixed price mechanisms, is not constrained efficient under mechanism set $\mathcal{M}'$ that also includes second price auctions*: *Under $\mathcal{M}'$ all sellers posting second price auctions with reserve below $\underline{v}$ and buyers using random visit strategies are constrained efficient.*

Proof. We compare the outcome of the equilibrium where sellers could only use fixed prices to the outcome when all sellers post second price auctions with reserve price below $\underline{v}$ and buyers use random visit strategies. Consider the case when both buyer types can trade in a price posting environment. The proof has two steps. First, we prove that there are strictly more trades in the auction environment. Second, we prove that high valuation buyers have strictly more trades. Together this establishes that the auction environment is more efficient.

For the first part of the argument, observe that in a price posting environment some fraction α of sellers has a low queue length $\underline{\lambda} = \underline{b}/\alpha s$ while a fraction $1-\alpha$ has a high queue length $\overline{\lambda} = \overline{b}/(1-\alpha)s$. The total number of trades in a price posting economy is

$$T(\alpha) = \alpha s\left[1 - P_0\left(\frac{\underline{b}}{\alpha s}\right)\right] + (1-\alpha)s\left[1 - P_0\left(\frac{\overline{b}}{(1-\alpha)s}\right)\right].$$

The first-order condition for the optimal number of trades is

$$\left[1 - P_0\left(\frac{\underline{b}}{\alpha s}\right)\right] + \frac{\underline{b}}{\alpha s}P_0'\left(\frac{\underline{b}}{\alpha s}\right) - \left[1 - P_0\left(\frac{\overline{b}}{(1-\alpha)s}\right)\right] - \frac{\overline{b}}{(1-\alpha)s}P_0'\left(\frac{\overline{b}}{(1-\alpha)s}\right) = 0,$$

[18] We do not have a full characterization for this rather large set of meeting technologies. Even starting from the urn-ball meeting function, lots of other meeting functions can be constructed that are non-rival and differ substantially from urn-ball. For example, all buyers in one-buyer matches can be coordinated into pairs, which increases the number of two-buyer matches and decreases the number of one-buyer matches: Let Q^U be the urn-ball meeting technology. Then $Q_0 = Q_0^U$, $Q_1 = 0$, $Q_2 = Q_2^U + Q_1^U/2$ and $Q_n = Q_n^U$ for $n > 2$ is also a non-rival meeting technology. Along this line, many other transformations are possible.

which is, under P_0 convex, only satisfied if $\underline{b}/\alpha s = \bar{b}/(1-\alpha)s$, i.e., when agents' visiting probabilities are the same at all firms. This is not achieved in a price posting environment which induces non-random visit strategies. Since $T''(\alpha) < 0$ at all α, the first-order condition indeed characterizes the optimal number of trades. The random visit strategies of the auction environment indeed yield the highest number of trades.

For the second part of the argument, we show that the high valuation buyers have strictly more trades under random visit strategies with auctions than in a price posting environment. Under auctions, the queue length for high buyer types is $\bar{\lambda}^a = \frac{\bar{b}}{s}$. The trading probability is

$$\sum_{\bar{n}=1}^{\infty} Q_{\bar{n}}\big(\bar{\lambda}^a\big)\frac{1}{\bar{n}} = \sum_{\bar{n}=1}^{\infty} \frac{P_{\bar{n}}(\bar{\lambda}^a)}{\bar{\lambda}^a} = \frac{1-P_0(\bar{\lambda}^a)}{\bar{\lambda}^a}.$$

Similarly, under price posting only a fraction $(1-\alpha)$ of firms attract high types which induces a queue of $\bar{\lambda}^p = \bar{b}/[(1-\alpha)s]$. A high buyer's trading probability is then by a similar logic $(1-P_0(\bar{\lambda}^p))/\bar{\lambda}^p$. Since P_0 is convex we have the trading probability increasing in λ, and since $\bar{\lambda}^a > \bar{\lambda}^p$ the high types can trade more often in the auction environment and therefore more often the high valuation surplus is realized.

Finally, consider the case when only the high valuation buyers can trade in a price posting equilibrium. They do not trade more often than in the auction environment, as we have shown that the auction environment maximizes the number of trades for high types. In fact, it achieves an identical amounts of trade for high types compared to price posting where only high types trade. Moreover, the auction environment allows some low types to trade when no high type is present, which strictly raises the surplus. Constrained efficiency of the auction environment follows trivially from the fact that auctions maximize the overall number of trades as well as the number of trades for high valuation buyers. □

Observe that any mechanism that does not screen buyers ex-post will generate less surplus than an equilibrium in second price auctions. We know from McAfee [17] and Peters [24] that an equilibrium in second price auctions exists under urn-ball meetings, and their approach can be extended to our setting by generalizing the meeting technology. They do not prove uniqueness. Here we show that there does not exist another equilibrium in which all sellers post fixed prices (or a mechanism without ex-post screening) when a larger set of ex-post screening mechanisms is allowed. The reason is that a seller can offer buyers the same payoff as under price posting, while enjoying exactly the efficiency gains that an auction environment implements.

Proposition 6. *Consider non-rival meetings. If all sellers post prices, it is strictly profitable for a deviating seller to post a second price auction coupled with a show-up fee.*

Proof. The equilibrium when sellers can only compete in prices delivers some utility $\underline{U}^\star$ and $\overline{U}^\star$ for the low and high type buyers, as well as some profit $\pi^\star$ for the sellers. Now consider a seller who posts a second price auction with show-up fee $-f$ and reserve r. The reserve r is due only when a single buyer is present, and we assume that buyers commit to paying this reserve when they are the only buyers, even if $r > \underline{v}$. With such a mechanism, the seller obtains a payoff $\pi^{f,r}_{\underline{n},\bar{n}} = r - f$ when there is one low or high buyer, obtains $\pi^{f,r}_{\underline{n},\bar{n}} = \underline{v} - (\underline{n}+\bar{n})f$ when $\bar{n} = 1, \underline{n} \geqslant 1$, or when $\underline{n} \geqslant 2$ and $\bar{n} = 0$; and obtains $\pi^{f,r}_{\underline{n},\bar{n}} = \bar{v} + (\underline{n}+\bar{n})f$ when $\bar{n} \geqslant 2$. In all other cases his payoff is zero. The seller solves the following program

$$\max_{(\underline{\lambda},\overline{\lambda})\in\mathbb{R}^2_+,(f,r)\in\mathbb{R}^2}\sum_{\underline{n}=0}^{\infty}\sum_{\overline{n}=0}^{\infty}P_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})\pi^{f,r}_{\underline{n},\overline{n}} \tag{31}$$

such that

$$\big(1-Q_0(\underline{\lambda}+\overline{\lambda})\big)f+\tilde{Q}_{0,0}(\underline{\lambda},\overline{\lambda})[\underline{v}-r]\leqslant\underline{U}^\star,\quad \text{with equality if } \underline{\lambda}>0, \tag{32}$$

$$\big(1-Q_0(\underline{\lambda}+\overline{\lambda})\big)f+\tilde{Q}_{0,0}(\underline{\lambda},\overline{\lambda})[\overline{v}-r]+\sum_{\underline{n}=1}^{\infty}\tilde{Q}_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})[\overline{v}-\underline{v}]\leqslant\overline{U}^\star,$$
$$\text{with equality if } \overline{\lambda}>0, \tag{33}$$

where the first constrained arises because the low type buyers obtain the good only when they are alone. The second constraint arises because high types only make positive profits when there are no other high types in the auctions, in which case they either pay the bid of the low type buyer if one is present, or they pay the reserve price. Incentive compatibility is trivially fulfilled for second price auctions.

Clearly the seller can implement the average queue length $\underline{\lambda}=\underline{b}/s$ and $\overline{\lambda}=\overline{b}/s$ with some combination of his instruments f and r. Now assume all sellers would use such an auction. Then the average queue length would be implemented at all sellers, which is feasible. We know from Proposition 5 that the surplus in this environment is higher than under price posting. But we have constructed the environment in such a way that the buyers obtain exactly the same utility as in the price posting environment, which means that the sellers must get higher profits. This in turn means that our deviating seller enjoys strictly higher profits than those sellers who do not screen buyers ex-post. □

5.2. *Purely rival meetings*

Here we assume that meetings are purely rival. Pure rivalry can be captured by the condition $\lambda(1-Q_0(\lambda))=1-P_0(\lambda)$. For a unit measure of sellers, the right-hand side gives the number of them that meet at least one buyer. The left-hand side gives the number of buyers in the market that meet at least one seller. Taking the total derivative with respect to λ and rearranging reveals the following: When more buyers enter the market, the meetings for existing buyers change by $\lambda Q_0'(\lambda)$, which equals the change in overall meetings $P_0'(\lambda)$ minus the meetings $(1-Q_0(\lambda))$ for new buyers. Every new buyer therefore takes one-for-one away from the meeting prospects for existing buyers.

It is easy to see from (1) that this happens if and only if meetings are bilateral. That is, $P_n(\lambda)=0$ for all $n>1$. This immediately implies that only $Q_0(\lambda)$ and $Q_1(\lambda)$ can be strictly positive. In this case we have

$$1-Q_0(\lambda)=Q_1(\lambda)=\frac{P_1(\lambda)}{\lambda}=\frac{1-P_0(\lambda)}{\lambda}.$$

Since $P_0'(\lambda)<0$ we cannot have $Q_0'(\lambda)=0$, and therefore the meeting probability of say high types changes when more low types enter the market. This is unavoidable in models with bilateral meetings. It constitutes the main externality in this setting.

We want to show that price posting is optimal for the sellers, i.e., no other mechanism performs better, and it is an equilibrium. Other mechanisms could still elicit buyers' types via wasteful destruction, yet the key insight with purely rival meetings is that ex-ante separation of types achieves type revelation without wasteful destruction.

Proposition 7. *Under purely rival meetings, price posting is always a best response by a seller.*

Proof. A seller now solves the program

$$\max_{(\underline{\lambda},\overline{\lambda})\in\mathbb{R}_+^2,\, m\in\mathcal{M}} P_1(\underline{\lambda}+\overline{\lambda})\left[\frac{\underline{\lambda}\pi_{1,0}^m+\overline{\lambda}\pi_{0,1}^m}{\underline{\lambda}+\overline{\lambda}}\right]$$
$$=\max_{(\underline{\lambda},\overline{\lambda})\in\mathbb{R}_+^2,\, m\in\mathcal{M}} Q_1(\underline{\lambda}+\overline{\lambda})\left[\underline{\lambda}\pi_{1,0}^m+\overline{\lambda}\pi_{0,1}^m\right] \tag{34}$$

such that

$$Q_1(\underline{\lambda}+\overline{\lambda})\underline{u}_{1,0}^m \leqslant \underline{U}^\star, \quad \text{with equality if } \underline{\lambda}>0, \tag{35}$$
$$Q_1(\underline{\lambda}+\overline{\lambda})\overline{u}_{0,1}^m \leqslant \overline{U}^\star, \quad \text{with equality if } \overline{\lambda}>0, \tag{36}$$

and the incentive compatibility constraints for truthful type revelation have to hold:

$$Q_1(\underline{\lambda}+\overline{\lambda})\underline{u}_{1,0}^m \geqslant Q_1(\underline{\lambda}+\overline{\lambda})\left[\overline{u}_{0,1}^m-\overline{x}_{0,1}^m(\overline{v}-\underline{v})\right],$$
$$Q_1(\underline{\lambda}+\overline{\lambda})\left[\underline{u}_{1,0}^m+\underline{x}_{1,0}^m(\overline{v}-\underline{v})\right] \leqslant Q_1(\underline{\lambda}+\overline{\lambda})\overline{u}_{0,1}^m$$

as well as the resource constraints $\pi_{1,0}^m+\underline{u}_{1,0}^m \leqslant \underline{x}_{0,1}^m\underline{v}$ and $\pi_{0,1}^m+\overline{u}_{0,1}^m \leqslant \overline{x}_{0,1}^m\overline{v}$. The incentive compatibility conditions can be rewritten as

$$\overline{x}_{0,1}^m \geqslant \frac{\overline{u}_{0,1}^m-\underline{u}_{1,0}^m}{\overline{v}-\underline{v}} \geqslant \underline{x}_{0,1}^m.$$

Optimality clearly has $\overline{x}_{0,1}^m=1$. Thus, any difference $\Delta u=\overline{u}_{0,1}^m-\underline{u}_{1,0}^m<\overline{v}-\underline{v}$ involves no trade with low types with probability $\frac{\Delta u}{\overline{v}-\underline{v}}$. Therefore, $\frac{\underline{\lambda}[\pi_{1,0}^m+\underline{u}_{1,0}^m]+\overline{\lambda}[\pi_{0,1}^m+\overline{u}_{0,1}^m]}{\underline{\lambda}+\overline{\lambda}} \leqslant \frac{\underline{\lambda}\Delta u\underline{v}+\overline{\lambda}\overline{v}}{\underline{\lambda}+\overline{\lambda}}$, it is easy to see that it is optimal to have this hold with equality. This reduces the problem to

$$\max_{(\underline{\lambda},\overline{\lambda})\in\mathbb{R}_+^2,\, m\in\mathcal{M}} Q_1(\underline{\lambda}+\overline{\lambda})\left[\underline{\lambda}\frac{\overline{u}_{0,1}^m-\underline{u}_{1,0}^m}{\overline{v}-\underline{v}}\left(\underline{v}-\underline{u}_{1,0}^m\right)+\overline{\lambda}\left(\overline{v}-\overline{u}_{0,1}^m\right)\right]$$

such that

$$Q_1(\underline{\lambda}+\overline{\lambda})\underline{u}_{1,0}^m \leqslant \underline{U}^\star, \quad \text{with equality if } \underline{\lambda}>0,$$
$$Q_1(\underline{\lambda}+\overline{\lambda})\overline{u}_{0,1}^m \leqslant \overline{U}^\star, \quad \text{with equality if } \overline{\lambda}>0,$$
$$0 \leqslant \overline{u}_{0,1}^m-\underline{u}_{1,0}^m \leqslant \overline{v}-\underline{v}.$$

Assume the optimal program has some contract m and $\underline{\lambda}>0$ and $\overline{\lambda}>0$. If $\underline{\lambda}\frac{\overline{u}_{0,1}^m-\underline{u}_{1,0}^m}{\overline{v}-\underline{v}}(\underline{v}-\underline{u}_{1,0}^m) \leqslant \overline{\lambda}(\overline{v}-\overline{u}_{0,1}^m)$ then there exists an optimal contract m' and $\overline{\lambda}'>0$ and $\underline{\lambda}'=0$. The reason is that at least the same payoff can be obtained by choosing $\overline{\lambda}'=\underline{\lambda}+\overline{\lambda}$ and $\overline{u}_{0,1}^{m'}=\overline{u}_{0,1}^m$ and $\overline{u}_{0,1}^m-\underline{u}_{1,0}^{m'}=\overline{v}-\underline{v}$. Yet this can be achieved with a price posting contract where the posted price p is such that $\overline{u}_{0,1}^m=\overline{v}-p$.

If $\underline{\lambda}\frac{\overline{u}_{0,1}^m-\underline{u}_{1,0}^m}{\overline{v}-\underline{v}}(\underline{v}-\underline{u}_{1,0}^m) > \overline{\lambda}(\overline{v}-\overline{u}_{0,1}^m)$ then there exists an optimal contract m' and $\overline{\lambda}'=0$ and $\underline{\lambda}'>0$. The reason is again that at least the same payoff can be obtained by choosing $\underline{\lambda}'=\underline{\lambda}+\overline{\lambda}$ and $\underline{u}_{0,1}^{m'}=\underline{u}_{0,1}^m$ and $\overline{u}_{0,1}^{m'}-\underline{u}_{1,0}^{m'}=\overline{v}-\underline{v}$. This again can be achieved with a price posting contract with a price such that $\underline{u}_{0,1}^m=\underline{v}-p$. □

Given that an equilibrium exists when sellers can only use price posting strategies, and given that no other mechanism achieves higher profits, we have

Corollary 4. *Under purely rival meetings, an equilibrium exists if the set of mechanisms $\mathcal{M}$ includes all price posting mechanisms. One equilibrium is identical to the price posting equilibrium of Section* 4.

Our final result concerns the constrained efficiency of the equilibrium. Again, constrained efficiency involves finding functions $\underline{\lambda}^p(\sigma)$ and $\overline{\lambda}^p(\sigma)$ such that (28) is fulfilled. We could additionally specify functions that destroy some of the surplus in order to induce truthful type revelation, but it is clear that a social planner would always want the seller to trade once a buyer shows up. The realized surplus that is to be maximized is then

$$S\big(\underline{\lambda}^p, \overline{\lambda}^p\big) = \int_0^s \big[P_{1,0}\big(\underline{\lambda}^p(\sigma) + \overline{\lambda}^p(\sigma)\big)\underline{v} + P_{0,1}\big(\underline{\lambda}^p(\sigma) + \overline{\lambda}^p(\sigma)\big)\overline{v}\big]\, d\sigma. \tag{37}$$

Proposition 8. *The price posting equilibrium is constrained efficient.*

Proof. Since $P_{1,0}(\underline{\lambda}, \overline{\lambda}) = P_1(\underline{\lambda} + \overline{\lambda})\frac{\underline{\lambda}}{\underline{\lambda}+\overline{\lambda}}$ and $P_{0,1}(\underline{\lambda}, \overline{\lambda}) = P_1(\underline{\lambda} + \overline{\lambda})\frac{\overline{\lambda}}{\underline{\lambda}+\overline{\lambda}}$ and under purely rival meetings $P_1(\underline{\lambda} + \overline{\lambda}) = 1 - P_0(\underline{\lambda} + \overline{\lambda})$ the surplus (37) is identical to the surplus specified in (29) in Section 4, for which we have shown that the queue length that arises under price posting is optimal. □

5.3. *Partially rival meetings*

For a synthesis that clearly highlights the constrained efficiency considerations that drive the choice of mechanisms, consider the intermediate case where meetings are partially rival. That means that neither $1 - Q_0(\lambda)$ is constant as in the purely non-rival case, nor $\lambda(1 - Q_0(\lambda)) - 1 + P_0(\lambda)$ is constant as in the purely rival case. Even though the literature including much of our present paper has focused on the extreme cases, such intermediate cases constitute an important avenue for future research.

Partially rival meeting functions are particularly important in our context because they show that the dominance of price posting is *not* an artifact of bilateral meetings, where price posting is optimal even in a non-competitive setting [30]. Rather, it is directly linked to the negative externality of low types on high types *before* they meet a seller. Partially rival meetings imply that at least some of the meetings are multilateral (since purely rival is equivalent to bilateral meetings). Therefore, in any partially rival meeting process buyers can use auctions at least for those meetings in which multiple buyers are present. And a monopolist seller who is allocated exogenously some random number of buyers would optimally use auctions to screen between buyers whenever multiple buyers are present. Here we show that with competing sellers, it is efficient to separate buyer types into separate markets when rivalry is strong, and in this case the equilibrium also features separation of types and prices constitute an equilibrium mechanism.

To state our results, it will be useful to introduce some notation first. Let the full information surplus of a market with unit measure of buyers, $\underline{\lambda}$ low buyers and $\overline{\lambda}$ high buyers be

$$S^F(\underline{\lambda},\overline{\lambda}) = \sum_{\underline{n}=0}^{\infty}\sum_{\overline{n}=1}^{\infty} P_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})\overline{v} + \sum_{\underline{n}=1}^{\infty} P_{\underline{n},0}(\underline{\lambda},\overline{\lambda})\underline{v}. \tag{38}$$

This is the surplus when every seller who has at least one high type sells the good to a high type, and every seller who has only low types sells the good to one of the low types. It is the highest surplus that can be achieved in this market. Now consider the surplus if a fraction α of the sellers attracts all the low type buyers to some market and sells whenever they have at least one buyer, and the remaining fraction $1-\alpha$ of sellers attracts all the high type buyers to some other market and sells whenever they have at least one buyer. In this case the total surplus is

$$(1-\alpha)S^F\big(0,\overline{\lambda}/(1-\alpha)\big) + \alpha S^F(\underline{\lambda}/\alpha, 0). \tag{39}$$

We will consider the case where for any $\underline{\lambda} > 0$ and $\overline{\lambda} > 0$ there exists α such that (39) is strictly larger than (38), so that it is efficient to separate markets. In this case we show that price posting always constitutes an equilibrium. This condition is demanding because it requires an ordering of (38) and (39) for all $\underline{\lambda} > 0$ and $\overline{\lambda} > 0$. A less demanding condition is needed to rule out a price posting equilibrium. Consider the overall buyer–seller ratios in the market $\underline{\lambda} = \underline{b}$ and $\overline{\lambda} = \overline{b}$, and the optimal fraction α^p of firms that cater to the low types is characterized in (30). If for these values (38) is larger than (39), than it is efficient to join markets and screen ex-post, and price posting cannot be part of any equilibrium. Before stating and proving this result formally, it might be instructive to discuss which type of meeting functions obtain such an ordering of (38) and (39).

Consider a meeting technology P for sellers that is a convex combination of a purely rival meeting technology P^R and a purely non-rival meeting technology P^N, i.e.,[19]

$$P = (1-\gamma)P^R + \gamma P^N. \tag{40}$$

This is a permissible meeting technology since consistency (1), stochastic dominance (2) and convexity (3) are preserved. Now fix some $\underline{\lambda}$ and $\overline{\lambda}$. In the previous analysis, Section 5.1 considered the extreme case where $\gamma = 1$ and we showed that (38) is larger than (39), while Section 5.2 analyzed the extreme case of $\gamma = 0$ and we showed that (38) is smaller than (39). By continuity, (38) is larger than (39) if the most of the weight is on P^N and (39) is larger than (38) if most weight is on P^R. To rule out price posting, we have to show that a deviant can generate higher profits. For some specifications of the rival and non-rival meeting technology, α^P as characterized in (30) does not depend on γ, and since $\underline{\lambda} = \underline{b}$ and $\overline{\lambda} = \overline{b}$ do not depend on γ either it is obvious that (38) is larger than (39) for γ sufficiently large.[20] To sustain price posting, we have to ensure that a deviant does not obtain higher profits. The deviant might attract very different buyer–seller ratios, and therefore we require efficiency of separation for *all* possible combinations of high and low types in a market. In Appendix A we discuss that there indeed exist P^R

[19] A more structural reason for hybrid meeting technologies is the following. Consider buyers who can send one letter to one of the buyers indicating that they want to trade (i.e., think of workers that have time to fill out one job application). But assume that sellers only have time to open and read up to N envelopes. The case $N = 1$ amounts to bilateral meetings: The seller can only see one buyer and cannot screen between any applicants any longer. With $N \to \infty$ we are in the standard urn-ball meeting environment with purely non-rival (multilateral) meetings, and sellers have full control whom to give the object to by choosing the appropriate auction format. To the extent that one can make N divisible, e.g., by interpreting $N = 1.5$ as a 50%–50% chance that the seller has time to open one or two envelopes, then $N \to 1$ has the same effect as $\gamma \to 0$ in the following.

[20] Note also that (39) dominates (38) for given partially rival meeting technology and given v_1 if v_2 is sufficiently large, since the rival nature in the meeting function will make it too costly to put both types in the same market.

and P^N and $\gamma > 0$ such that (38) uniformly dominates (39). Therefore, if meetings are predominantly governed by a rival meeting function, price posting constitutes an equilibrium even though sellers can commit to auctions when meetings are multilateral.

The following shows formally that the nature of the efficiency ordering between (38) and (39) crucially affects the types of mechanisms used in equilibrium. For the second part of the proposition, recall that α^p is the fraction of firms that cater to low types when only fixed price mechanisms are available.

Proposition 9. *If the meeting technology is such that for all* $\bar{\lambda} > 0$ *and* $\underline{\lambda} > 0$ *there exists an* $\alpha \in [0, 1]$ *such that*

$$S^F(\underline{\lambda}, \bar{\lambda}) < \alpha S^F(\underline{\lambda}/\alpha, 0) + (1-\alpha) S^F\big(0, \bar{\lambda}/(1-\alpha)\big), \tag{41}$$

then there exists an equilibrium in which all sellers post prices.

If the meeting technology is such that

$$S^F(\underline{b}, \bar{b}) > \alpha^p S^F\big(\underline{b}/\alpha^p, 0\big) + \big(1-\alpha^p\big) S^F\big(0, \bar{b}/\big(1-\alpha^p\big)\big), \tag{42}$$

then there exists no equilibrium in which all sellers post prices if the mechanisms space is rich enough (*e.g., includes second price auctions with reserve and participation fee*).

Proof. For the first part, assume (41) holds. Consider a candidate equilibrium. We will show that no matter what mechanism m an individual seller considers to post, he can make weakly higher profits by posting a price. This establishes the result.

Consider any mechanism m. It has to satisfy resource constraint (10). This constraint is most relaxed if the good is always given to the high type when possible, i.e. $\bar{x}^m_{\underline{n},\bar{n}} = 1$ if $\bar{n} > 0$ and zero otherwise, and $\underline{x}^m_{\underline{n},\bar{n}} = 1$ if $\bar{n} = 0$ but $\underline{n} > 0$ and zero otherwise. In this case constraint (10) reduces to

$$\sum_{\underline{n}} \sum_{\bar{n}} P_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda})\big[\pi^m_{\underline{n},\bar{n}} + \underline{n}\underline{u}^m_{\underline{n},\bar{n}} + \bar{n}\bar{u}^m_{\underline{n},\bar{n}}\big] \leqslant S^F(\underline{\lambda}, \bar{\lambda}). \tag{43}$$

Define the expected profit using this mechanism as $\Pi = \sum_{\underline{n}} \sum_{\bar{n}} P_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda})\pi^m_{\underline{n},\bar{n}}$. Further, note that in any optimal mechanism (if it attracts the low type) the participating constraint (19) binds $\sum_{\underline{n}} \sum_{\bar{n}} \tilde{Q}_{\underline{n},\bar{n}}(\underline{\lambda}, \bar{\lambda})\underline{u}^m_{\underline{n}+1,\bar{n}} = \underline{U}$. Similar for high types. Using the relationship between P and Q in (1) and the other properties on the meeting technology, we show in Appendix A that (43) can equivalently be written as

$$\Pi + \underline{\lambda}\underline{U} + \bar{\lambda}\bar{U} \leqslant S^F(\underline{\lambda}, \bar{\lambda}). \tag{44}$$

This has the clear-cut interpretation that the expected profit for the seller plus the expected number of low types times their individual expected utility plus the expected number of high types times their individual expected utility has to be less than the expected surplus generated by an individual seller. Since we started with the premise that (41) holds with strict inequality if both $\underline{\lambda}$ and $\bar{\lambda}$ are strictly positive, we have

$$\Pi + \underline{\lambda}\underline{U} + \bar{\lambda}\bar{U} < \alpha S^F\left(\frac{\underline{\lambda}}{\alpha}, 0\right) + (1-\alpha) S^F\left(0, \frac{\bar{\lambda}}{1-\alpha}\right), \tag{45}$$

if both $\underline{\lambda}$ and $\bar{\lambda}$ are strictly positive, and by continuity it holds with a weak inequality otherwise.

Now consider an individual seller who contemplates to post a price. He considers price $\underline{p}$ such that participation constraint (19) for the low types is exactly met at buyer–seller ratio $\underline{\lambda}/\alpha$. Therefore, his profit is at least $\Pi(\underline{p}) = (1 - P_0(\underline{\lambda}/\alpha))\underline{p}$.[21] Since there is always trade if at least one buyer shows up, all the rest of the surplus goes to the buyers and it is easy to show that $S^F(\frac{\underline{\lambda}}{\alpha}, 0) = \Pi(\underline{p}) + \frac{\underline{\lambda}}{\alpha}\underline{U}$.

Alternatively, he can contemplate posting a high price $\overline{p}$ such that participation constraint (20) for the high types is exactly met at buyer–seller ratio $\overline{\lambda}/(1-\alpha)$. He makes at least profit $\Pi(\overline{p}) = (1 - P_0(\frac{\overline{\lambda}}{1-\alpha}))\overline{p}$. Again the rest of the surplus goes to the buyers and $S^F(0, \frac{\overline{\lambda}}{1-\alpha}) = \Pi(\overline{p}) + \frac{\overline{\lambda}}{1-\alpha}\overline{U}$. Therefore, we can write (45) as

$$
\begin{aligned}
&\Pi + \underline{\lambda}\underline{U} + \overline{\lambda}\overline{U} < \alpha\big[\Pi(\underline{p}) + (\underline{\lambda}/\alpha)\underline{U}\big] + (1-\alpha)\big[\Pi(\overline{p}) + \big(\overline{\lambda}/(1-\alpha)\big)\overline{U}\big] \\
&\quad\Leftrightarrow\quad \Pi < \alpha\Pi(\underline{p}) + (1-\alpha)\Pi(\overline{p}) \\
&\quad\Rightarrow\quad \Pi < \max\big\{\Pi(\underline{p}), \Pi(\overline{p})\big\},
\end{aligned}
$$

where the inequality is weak if either $\underline{\lambda}$ or $\overline{\lambda}$ is zero. Therefore, price posting is always at least as profitable as posting any other mechanism. Thus, equilibrium in which sellers compete only in prices remains an equilibrium even if other mechanisms are available. Moreover, posting the optimal price is strictly more profitable than any mechanism that attracts both buyer types with positive probability. Therefore, any equilibrium has strict separation of types in different markets.

The second part of the proof is essentially a reversal of the above arguments. The logic of Proposition 6 applies, and the price posting equilibrium cannot survive. □

This highlights that price posting is not an artifact of purely rival meetings. Rather, it arises from the interaction of buyers in the search process. If the externalities in the search process are strong in the sense that bad types induce an externality such that good types find it difficult to reach the mechanism, sellers do not find it optimal to attract both types of buyers even if they could screen them apart in the event that multiple buyers reach the mechanism. Such externalities arise, e.g., when the seller is time constrained and cannot interact with all potential buyers (see also footnote 19). The analysis of competition in mechanisms therefore crucially relies on the properties of the underlying meeting process that hitherto has not been considered in the literature.

6. Conclusion

Posted prices are prevalent in many economic environments, yet theory from the literature on search and competing mechanism design tells us that auctions generate higher surplus. In this paper we have shown that the characteristics of the meeting technology are crucially important for which equilibrium sales mechanism is used. When meetings are rival, low buyer types significantly affect the prospects of the high types obtaining the good and ex-post screening as in auctions is very costly. Instead, under price posting different buyer types adequately sort ex-ante, ensuring that high types trade with sufficiently high probability.

[21] Participation constraint (19) is binding if $(1 - Q_0(\underline{\lambda}/\alpha))(\underline{v} - \underline{p}) = \underline{U}$. By (16) this will indeed be the queue length that this seller attracts if $(1 - Q_0(\underline{\lambda}/\alpha))(\overline{v} - \underline{p}) \leqslant \overline{U}$, since in this case he does not attract any high types. He could attract even more buyers if $(1 - Q_0(\underline{\lambda}/\alpha))(\overline{v} - \underline{p}) > \overline{U}$, because then the buyer–seller ratio is determined by indifference of the high types according to (17), and a higher buyer–seller ratio at the same price means even higher profits.

In order to rationalize the prevailing equilibrium mechanism observed in markets, we offer this as a novel explanation. Much of the mechanism design literature focuses on the role of variations in the mechanism space. For example, can competing sellers condition their mechanism on the mechanism of other sellers (see Epstein and Peters [7])? As an alternative, our results show that the characteristics of the search frictions can explain the nature of the observed equilibrium trade mechanism. Because of frictions, valuable information is obtained from ex-ante selection even before the mechanism is called to act. Such selection is an integral feature of competition.

We finish by pointing out a parallel in our findings to standard price theory. In general equilibrium, when goods are rival the price mechanism works well in the allocation process, while other mechanisms are required when goods are non-rival. In our setting, goods are purely rival, but the meeting technology may not be. When meetings are rival, prices allocate resources well. In contrast, when meetings are non-rival, other mechanisms outperform the price mechanism.

Appendix A

A.1. Elaboration on rival and non-rival meeting technologies P^R and P^N in (40)

As a special case of a purely non-rival meeting technology, consider the urn-ball technology with $P_n^N(\lambda) = \frac{e^{-\lambda}\lambda^n}{n!}$. We can transform this into a purely rival meeting technology by assuming that whenever several buyers are present, only one is selected at random to enter the mechanism and the others are excluded. This yields the purely rival meeting technology P^R such that $P_0^R(\lambda) = P_0^N(\lambda) = e^{-\lambda}$, $P_1^R(\lambda) = \sum_{n=1}^{\infty} P_n^N(\lambda) = 1 - e^{-\lambda}$, and $P_n^R(\lambda) = 0$ for $n > 0$.

Now consider the convex combination $P = \gamma P^R + (1-\gamma)P^N$. For this meeting technology, it is clear that the surplus from having buyers in separated markets: $\alpha S^F(\underline{\lambda}/\alpha, 0) + (1-\alpha)S^F(0, \overline{\lambda}/(1-\alpha))$, is independent of γ for all $\underline{\lambda}$, $\overline{\lambda}$ and α. This is apparent because in separate markets it does not matter which buyer is selected by the mechanisms because in each market buyers are homogeneous. Therefore, whether the mechanism selects the buyer as in P^N or one buyer is chosen at random as under P^R does not affect efficiency. The full information efficiency when having both types search in the same market does change with γ, though. When γ is low, then the meeting function randomly selects a buyer and a low type might be chosen instead of a high type, and separation would be preferable (Proposition 8). When γ is high than most buyers enter some mechanism and the higher type is chosen when both types are present, and pooling is preferable (Proposition 5). Applying the logic of Proposition 6 rules out the existence of a price posting equilibrium in the latter case.

Showing that a price posting equilibrium does exist requires by the first part of Proposition 9 that the surplus from separating surpasses the surplus from pooling for all combinations of buyer–seller ratios of the two types. We will briefly discuss that there indeed exists a $\gamma > 0$ such that $P = \gamma P^N + (1-\gamma)P^R$ has this property. Consider first the surplus from separation. We have

$$\alpha S^F\left(\frac{\underline{\lambda}}{\alpha}, 0\right) + (1-\alpha)S^F\left(0, \frac{\overline{\lambda}}{1-\alpha}\right) = \alpha\left(1 - e^{-\frac{\underline{\lambda}}{\alpha}}\right)\underline{v} + (1-\alpha)\left(1 - e^{-\frac{\overline{\lambda}}{1-\alpha}}\right)\overline{v}.$$

The optimal α^P is $\alpha^P = 0$ if $v_1 - (1 - e^{-\lambda_2} - \lambda_2 e^{-\lambda_2})v_2 < 0$. Otherwise it is uniquely characterized by the first-order condition

$$\left(1 - e^{-\frac{\lambda_1}{\alpha^P}} - \frac{\lambda_1}{\alpha^P}e^{-\frac{\lambda_1}{\alpha}}\right)v_1 - \left(1 - e^{-\frac{\overline{\lambda}}{1-\alpha^P}} - \frac{\overline{\lambda}}{1-\alpha^P}e^{-\frac{\overline{\lambda}}{1-\alpha^P}}\right)\overline{v} = 0,$$

which corresponds to (30) in the main text. The surplus at the optimal α^P is a function of $\underline{\lambda}$ and $\overline{\lambda}$, but is not affected by γ. Now the difference between the pooling and the separating surplus can be written as

$$\Delta(\underline{\lambda},\overline{\lambda},\gamma) = S^F(\underline{\lambda},\overline{\lambda}) - \alpha^P S^F\left(\frac{\underline{\lambda}}{\alpha^P},0\right) - (1-\alpha)S^F\left(0,\frac{\overline{\lambda}}{1-\alpha^P}\right) \tag{46}$$

$$= \gamma\left(1-e^{-\overline{\lambda}}\right)\overline{v} + \gamma e^{-\overline{\lambda}}\left(1-e^{-\underline{\lambda}}\right)\underline{v} \tag{47}$$

$$+ (1-\gamma)\left(1-e^{-\overline{\lambda}-\underline{\lambda}}\right)\left(\frac{\underline{\lambda}}{\underline{\lambda}+\overline{\lambda}}\underline{v} + \frac{\overline{\lambda}}{\underline{\lambda}+\overline{\lambda}}\overline{v}\right) \tag{48}$$

$$- \alpha^P\left(1-e^{-\frac{\underline{\lambda}}{\alpha^P}}\right)\underline{v} - \left(1-\alpha^P\right)\left(1-e^{-\frac{\overline{\lambda}}{1-\alpha^P}}\right)\overline{v}, \tag{49}$$

where lines (47) and (48) give the surplus under pooling: (47) is the surplus under non-rival meetings when the good is given to the high type whenever at least one high type is present and to the low type only if no high type is present, while (48) represents the surplus under the rival part of the meeting function whenever one buyer is selected at random when at least one buyer is present. For given $\underline{\lambda} > 0$ and $\overline{\lambda} > 0$, $\Delta(\underline{\lambda},\overline{\lambda},0) < 0$ by Proposition 8, while $\Delta(\underline{\lambda},\overline{\lambda},1) > 0$ by Proposition 5. By continuity there exists $\gamma(\underline{\lambda},\overline{\lambda}) > 0$ such that $\Delta(\underline{\lambda},\overline{\lambda},\gamma(\underline{\lambda},\overline{\lambda})) = 0$.

Condition (41) has to hold for all $\underline{\lambda} > 0$ and $\overline{\lambda} > 0$, which is equivalent to requiring that $\gamma(\underline{\lambda},\overline{\lambda})$ remains bounded away from zero for all $\underline{\lambda} > 0$ and $\overline{\lambda} > 0$. This is not obvious since $\Delta(\underline{\lambda},\overline{\lambda},1)$ converges to zero when either $\underline{\lambda}$ or $\overline{\lambda}$ converges to zero. The reason is that in the presence of (essentially) only a single type the gain from pooling over separation is rather low. But to the same extent $\Delta(\underline{\lambda},\overline{\lambda},0)$ converges to zero because the gain from sorting in different markets rather than taking a random buyer is also rather low. Both together mean that increasing γ yields only moderate benefits in (47) but also induces only moderate costs in (48) relative to the constant surplus from separation in (49). While (46) remains analytically intractable, to assess whether $\gamma(\underline{\lambda},\overline{\lambda})$ remains bounded away from zero we have evaluated $\gamma(\underline{\lambda},\overline{\lambda})$ on a (logarithmic) grid for $(\underline{\lambda},\overline{\lambda})$. The grid spans values for $(\underline{\lambda},\overline{\lambda})$ from $(0.0001, 0.0001)$ up to $(1000, 1000)$. The value of γ is smallest when both queue lengths are small. Varying the values of $(\underline{\lambda},\overline{\lambda})$ by a factor of 100 [going from $(0.01, 0.01)$ to $(0.001, 0.001)$ to $(0.0001, 0.0001)$] changes the value of γ by only less than half a percent [from 0.17215 to 0.17163 to 0.17158], suggesting that γ is bounded away from zero.

A.2. Derivation of inequality (44) from (43)

In the following derivation we repeatedly use (1), $\tilde{Q}_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) = Q_{\underline{n}+\overline{n}+1}(\underline{\lambda}+\overline{\lambda})B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})$, and the fact that $B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})\frac{\underline{n}}{\underline{n}+\overline{n}}\frac{\underline{\lambda}+\overline{\lambda}}{\underline{\lambda}} = B_{\underline{n}-1,\overline{n}}(\underline{\lambda},\overline{\lambda})$ and $B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})\frac{\overline{n}}{\underline{n}+\overline{n}}\frac{\underline{\lambda}+\overline{\lambda}}{\overline{\lambda}} = B_{\underline{n},\overline{n}-1}(\underline{\lambda},\overline{\lambda})$:

$$\underbrace{\sum_{\underline{n}}\sum_{\overline{n}}}_{\underline{n}+\overline{n}>0} P_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})\left[\underline{n}\underline{u}^m_{\underline{n},\overline{n}} + \overline{n}\overline{u}^m_{\underline{n},\overline{n}}\right]$$

$$= \underbrace{\sum_{\underline{n}}\sum_{\overline{n}}}_{\underline{n}+\overline{n}>0} P_{\underline{n}+\overline{n}}(\underline{\lambda}+\overline{\lambda})B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda})\left[\underline{n}\underline{u}^m_{\underline{n},\overline{n}} + \overline{n}\overline{u}^m_{\underline{n},\overline{n}}\right]$$

$$
\begin{aligned}
&= \underbrace{\sum_{\underline{n}} \sum_{\overline{n}}}_{\underline{n}+\overline{n}>0} Q_{\underline{n}+\overline{n}}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \frac{\underline{\lambda}+\overline{\lambda}}{\underline{n}+\overline{n}} \big[\underline{n}\underline{u}^m_{\underline{n},\overline{n}} + \overline{n}\overline{u}^m_{\underline{n},\overline{n}}\big] \\
&= \underline{\lambda} \sum_{\underline{n}>0} \sum_{\overline{n}} Q_{\underline{n}+\overline{n}}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \frac{\underline{n}}{\underline{n}+\overline{n}} \frac{\underline{\lambda}+\overline{\lambda}}{\underline{\lambda}} \underline{u}^m_{\underline{n},\overline{n}} \\
&\quad + \overline{\lambda} \sum_{\underline{n}} \sum_{\overline{n}>0} Q_{\underline{n}+\overline{n}}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \frac{\overline{n}}{\underline{n}+\overline{n}} \frac{\underline{\lambda}+\overline{\lambda}}{\overline{\lambda}} \overline{u}^m_{\underline{n},\overline{n}} \\
&= \underline{\lambda} \sum_{\underline{n}>0} \sum_{\overline{n}} Q_{\underline{n}+\overline{n}}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n}-1,\overline{n}}(\underline{\lambda},\overline{\lambda}) \underline{u}^m_{\underline{n},\overline{n}} + \overline{\lambda} \sum_{\underline{n}} \sum_{\overline{n}>0} Q_{\underline{n}+\overline{n}}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n},\overline{n}-1}(\underline{\lambda},\overline{\lambda}) \overline{u}^m_{\underline{n},\overline{n}} \\
&= \underline{\lambda} \sum_{\underline{n}} \sum_{\overline{n}} Q_{\underline{n}+\overline{n}+1}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \underline{u}^m_{\underline{n}+1,\overline{n}} \\
&\quad + \overline{\lambda} \sum_{\underline{n}} \sum_{\overline{n}} Q_{\underline{n}+\overline{n}+1}(\underline{\lambda}+\overline{\lambda}) B_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \overline{u}^m_{\underline{n},\overline{n}+1} \\
&= \underline{\lambda} \sum_{\underline{n}} \sum_{\overline{n}} \tilde{Q}_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \underline{u}^m_{\underline{n}+1,\overline{n}} + \overline{\lambda} \sum_{\underline{n}} \sum_{\overline{n}} \tilde{Q}_{\underline{n},\overline{n}}(\underline{\lambda},\overline{\lambda}) \overline{u}^m_{\underline{n},\overline{n}+1} \\
&= \underline{\lambda}\underline{U} + \overline{\lambda}\overline{U}.
\end{aligned}
$$

References

[1] Daron Acemoğlu, Robert Shimer, Holdup and efficiency with search frictions, Int. Econ. Rev. 47 (1999) 651–699.

[2] Jim Albrecht, Susan Vroman, A matching model with endogenous skill requirements, Int. Econ. Rev. 43 (1) (2002) 283–305.

[3] James Albrecht, Pieter A. Gautier, Susan Vroman, Equilibrium directed search with multiple applications, Rev. Econ. Stud. 73 (2006) 869–891.

[4] Kenneth Burdett, Shouyong Shi, Randall Wright, Pricing and matching with frictions, J. Polit. Economy 109 (2001) 1060–1085.

[5] Melvyn Coles, Jan Eeckhout, Indeterminacy and directed search, J. Econ. Theory 111 (2003) 265–276.

[6] Jan Eeckhout, Philipp Kircher, Sorting and decentralized price competition, Econometrica (2010), forthcoming.

[7] Larry Epstein, Michael Peters, A revelation principle for competing mechanisms, J. Econ. Theory 88 (1) (1999) 119–161.

[8] Manolis Galenianos, Philipp Kircher, Directed search with multiple job applications, J. Econ. Theory 114 (2) (2009) 445–471.

[9] Manolis Galenianos, Philipp Kircher, Gabor Virag, Market power and efficiency in a search model, Int. Econ. Rev. (2011), forthcoming.

[10] Veronica Guerrieri, Robert Shimer, Randall Wright, Adverse selection in competitive search equilibrium, mimeo, 2009.

[11] Arthur Hosios, On the efficiency of matching and related models of search and unemployment, Rev. Econ. Stud. 57 (1990) 279–298.

[12] Benoit Julien, John Kennes, Ian King, Bidding for labor, Rev. Econ. Dynam. 3 (4) (2000) 619–649.

[13] Philipp Kircher, Efficiency of simultaneous search, J. Polit. Economy 117 (5) (2009) 861–913.

[14] Nobuhiro Kiyotaki, Randall Wright, A search-theoretic approach to monetary economics, Amer. Econ. Rev. 83 (1) (1993) 63–77.

[15] Klaus Kultti, Equivalence of auctions and posted prices, Games Econ. Behav. 27 (1999) 106–113.

[16] Andreu Mas-Colell, On a theorem by Schmeidler, J. Math. Econ. 13 (1984) 201–206.

[17] Preston McAfee, Mechanisms design by competing sellers, Econometrica 61 (1993) 1281–1312.

[18] Espen R. Moen, Competitive search equilibrium, J. Polit. Economy 105 (1997) 385–411.
[19] Espen R. Moen, Asa Rosen, Incentives in competitive search equilibrium and wage rigidity, mimeo, Norwegian School of Management, 2006.
[20] Dale T. Mortensen, Randall Wright, Competitive pricing and efficiency in search equilibrium, Int. Econ. Rev. 43 (2002) 1–20.
[21] Giuseppe Moscarini, Excess worker reallocation, Rev. Econ. Stud. 69 (3) (2001) 593–612.
[22] Michael Peters, Ex ante pricing in matching games: non steady states, Econometrica 59 (1991) 1425–1454.
[23] Michael Peters, On the equivalence of Walrasian and non-Walrasian equilibria in contract markets: the case of complete contracts, Rev. Econ. Stud. 64 (1997) 241–264.
[24] Michael Peters, A competitive distribution of auctions, Rev. Econ. Stud. 64 (1997) 97–123.
[25] Michael Peters, Competition among mechanism designers in a common value environment, Rev. Econ. Design 4 (1999) 273–292.
[26] Michael Peters, Limits of exact equilibria for capacity constrained sellers with costly search, J. Econ. Theory 95 (2000) 139–168.
[27] Michael Peters, Common agency and the revelation principle, Econometrica 69 (5) (2001) 1349–1372.
[28] Michael Peters, Sergei Severinov, Competition among sellers who offer auctions instead of prices, J. Econ. Theory 75 (1997) 141–179.
[29] Roberto Pinheiro, Venture capital and underpricing: capacity constraints and early sales, University of Pennsylvania mimeo, 2008.
[30] John G. Riley, Richard J. Zeckhauser, Optimal selling strategies: when to haggle, when to hold firm, Quart. J. Econ. 98 (1983) 267–289.
[31] Shouyong Shi, Frictional assignment I: efficiency, J. Econ. Theory 98 (2001) 232–260.
[32] Shouyong Shi, A directed search model of inequality with heterogeneous skills and skill-biased technology, Rev. Econ. Stud. 69 (2002) 467–491.
[33] Robert Shimer, The assignment of workers to jobs in an economy with coordination frictions, J. Polit. Economy 113 (5) (2005) 996–1025.
[34] Gabor Virag, Collusive equilibria in directed search models, Rochester mimeo, 2007.
[35] Ruqu Wang, Auctions versus posted-price selling, Amer. Econ. Rev. 83 (4) (1993) 838–851.