THE UNIVERSITY OF MELBOURNE

Department of Economics

Working Paper Series

Monotonicity, Non-Participation, and Directed Search Equilibria

James Bland, Simon Loertscher

June 2012

Research Paper Number 1147

ISBN:0819-2642

ISBN:97807 340 44976

Department of Economics
The University of Melbourne
Parkville VIC 3010
www.economics.unimelb.edu.au

Monotonicity, Non-Participation, and Directed Search Equilibria*

James Bland[†] Simon Loertscher[‡]

June 12, 2012

Abstract

We consider the canonical directed search framework in which sellers play pure strategies and assume that buyers play strategies that are monotone in prices, can remain inactive and choose to do so whenever their payoff from participating is zero regardless of what the other buyers do. We show that directed search equilibria, which have been the focus of the literature, are the only equilibria that satisfy these assumptions. Directed search equilibria are selected here not because buyers cannot coordinate – no such assumption is made – but because they fail to play strategies that require them to increase the demand for a seller's good as this good becomes more expensive.

Keywords: Directed search, monotone strategies, directed search equilibrium.

JEL-Classification: C72, D72.

*This paper has greatly benefitted from comments by Ian P. King and Randall Wright. A previous version of this paper circulated under the title "Matching with Frictions and the Law of Demand". Financial support by the Centre for Market Design (CMD) at the University of Melbourne is gratefully acknowledged.

[†]Department of Economics, Arts West, University of Melbourne, Victoria 3010, Australia. Email: james.bland@gmail.com

[‡]Department of Economics, Arts West, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au.

1 Introduction

Matching models with directed search have come to centre stage in economics. Perhaps most prominently, these models are applied to search behavior in labor markets and in markets for consumer goods. In the canonical model of Burdett, Shi, and Wright (2001, BSW herafter), agents on one side of the market – call them sellers – first simultaneously post prices. Observing these prices, the agents on the other side, called buyers, then simultaneously choose which seller to visit. If multiple buyers go to the same seller, each is served with equal probability and nets as her payoff the difference between her valuation and the price in case she is served. Otherwise, her payoff is zero.[1]

The second stage of this game being a coordination game, equilibria and equilibrium outcomes abound. Equilibrium selection is, therefore, inevitable. Arguing that coordination on equilibria in which buyers play pure strategies is difficult in large and anonymous markets, the literature has focused on what may be called *directed search equilibria (DSE)*, in which all buyers play mixed strategies on and locally off the equilibrium path.[2] In a DSE, sellers set mutual best response prices in the first stage, best responses being conditioned on the mixed strategy equilibrium in the second stage.

In this paper, we provide an alternative and complementary argument that leads to the selection of DSE in the canonical directed search model, which does not rest on the assumption that buyers experience difficulties in coordinating on a pure strategy equilibrium. All we require is that buyers play monotone strategies, with monotonicity meaning that every buyer's demand for the good offered by every seller is non-increasing in this seller's price and non-decreasing in the other sellers' prices, and that they abstain whenever the payoff from participating is zero regardless of what the other buyers do. Our analysis thus predicts that if buyers play monotone strategies, the DSE is obtained even if buyers could coordinate, implicitly or explicitly.

We now briefly illustrate how monotonicity of buyer strategies eliminates equilibria other than the DSE and explain why we think this monotonicity is a sensible and intuitive restriction. Consider the simple price posting game with two buyers, 1 and 2, whose valuations are normalized to one and two sellers, A and B, each with a capacity to produce one good at a

[1] See also Montgomery (1991), Moen (1997), and Peters (2000), who analyzed variants of this model.

[2] Various alternative names have been proposed for these equilibria. Peters (1984) calls the DSE in which all buyers play identical mixed strategies the "symmetric buyers equilibrium". Others refer to DSE as mixed strategy equilibria. Though within the standard directed search framework there are plethora of other equilibria, we think the term "directed search equilibria" is descriptive, and we minimize the potential for ambiguity by using alternative names for any other type of equilibrium.

cost of zero. As noted by BSW, this game has many equilibria besides the DSE. For the sake of the argument, consider the one where on the equilibrium path each seller sets a price of 2/3 and buyer 1 (2) visits seller A (B) with probability one. Upon a small upward deviation by one seller, say A, each buyer plays the non-degenerately mixed equilibrium strategy, that is, visits each seller with strictly positive probability (which is pinned down by the two prices). While buyer 1's strategy satisfies monotonicity because she visits the seller whose price is higher with lower probability than she does on the equilibrium path, buyer 2's strategy violates monotonicity because she now visits seller A with higher probability and seller B with lower probability. It is not hard to see that, for the same reason, all equilibria in which buyers play pure strategies on the equilibrium path and mixed strategies off the equilibrium path violate monotonicity, and we show that the same kind of logic extends to all other equilibria as well.

To see in what way the non-monotone equilibrium with prices of 2/3 is not plausible, consider buyer 2's problem after seller A's deviation. According to equilibrium, she should now visit seller A with positive probability whereas she visited him with probability zero before the price increase *despite* the fact that she has no strict incentive to do so: Given that buyer 1 plays the mixed strategy the equilibrium prescribes, buyer 2 is indifferent between visiting seller A and B. If buyer 1 is aware that buyer 2 will fail to play the non-monotone strategy and will not go the seller A with sufficiently high probability, her unique best response is to stay with seller A. But anticipating this, seller A has indeed an incentive to increase his price unilaterally.[3]

The paper is related to the already vast yet still growing literature on directed search models.[4] Though the approaches differ, our paper is complementary to Galenianos and Kircher (2012)'s contribution: Taking as given the existence of an equilibrium in pure strategies for sellers, whose existence Galenianos and Kircher establish for a special case of the mechanisms considered here, we provide a refinement, monotonicity, that selects mixed strategies in the buy-

[3] Buyers playing mixed strategies may seem to require little coordination on the equilibrium path, that is, when all sellers set the same prices, because in the mixed strategy equilibrium of the ensuing subgame each buyer randomizes uniformly across all sellers. However, off the equilibrium path when prices differ, substantial coordination of beliefs is required in the sense that each buyer must be very confident that every other buyer chooses exactly the right (that is, equilibrium) mixed strategy, despite the fact that no buyer has a strict inventive to get the mixture exactly right. In the example with price posting and two buyers and two sellers, the DSE price is 1/2. If seller A unilaterally deviates and sets $1/2 + \varepsilon$ with $\varepsilon > 0$ small, each buyer must visit seller A with probability $(1 - 4\varepsilon)/(2 - 2\varepsilon)$.

[4] In labor market models such as Shimer (2005), workers are the agents who direct their search, but nothing of substance hinges on that distinction. Directed search models have been extensively used in the literatures on consumer, labor, and money markets with search (see e.g. Shimer, 2005; Guerrieri, 2008; Julien, Kennes, and King, 2008; Menzio and Shi, 2010; Hawkins, forthcoming).

ers' subgame. Galenianos and Kircher, on the other hand, assume that this is the equilibrium that will be selected in the buyers' subgame. In Section 2, we introduce a general mechanism, called γ-mechanism, which is a convex combination of the price posting mechanism analyzed, for example, by BSW and the auction mechanism introduced by Julien, Kennes, and King (2000, JKK hereafter). Consequently, price posting and auctions emerge as limit cases on the opposite ends of the spectrum of γ-mechanisms. The purpose of introducing γ-mechanisms is that it permits a unified analysis and exposition. In Section 3, we derive our results. Section 4 concludes. All the proofs are in the appendix.

2 The Model

There are N risk neutral buyers and M risk neutral sellers. All buyers have quasilinear preferences and a valuation of 1 for a homogenous good of known quality. All sellers have a capacity to produce 1 unit at marginal costs that are normalized to 0. The game has two stages. In the first stage, all sellers j simultaneously post prices p_j, $j = A, \ldots, M$. We assume that prices are restricted to be elements of $[0, 1]$, and denote by $\mathbf{p} \equiv (p_A, \ldots, p_M) \in [0, 1]^M$ the collection of all prices, and by $\mathbf{p}_{-j}$ all prices other than seller j's. Bold symbols denote vectors.

In the second stage, which will also be referred to as the buyer subgame, having observed $\mathbf{p}$ all buyers simultaneously choose which seller to visit. In a slight departure from the standard set of assumptions, we assume also that buyers have the option of remaining inactive. Remaining inactive generates a payoff of 0 and is formally modeled as an artificial and non-strategic seller Z (for zero). A – possibly degenerately – mixed strategy $\boldsymbol{\theta}^i$ for buyer i is therefore a mapping of prices into a probability distribution over sellers: $\boldsymbol{\theta}^i : [0, 1]^M \to \Delta$, where Δ is the M dimensional simplex (there being $M + 1$ sellers including the non-strategic seller Z). Denote the probability that buyer i visits seller j given prices $\mathbf{p}$ by $\theta^i_j(\mathbf{p})$. For every buyer i and every $\mathbf{p} < \mathbf{1}$, $\sum_{j=A}^{M} \theta^i_j(\mathbf{p}) = 1$ will hold in any subgame perfect equilibrium, where $\mathbf{x} < \mathbf{y}$ means $x_i \leq y_i$ for all i with at least one inequality strict.

Two kinds of mechanisms for determining transaction prices at a given seller, given a number of buyers that have approached this seller, have received particular attention in the literature. Under **price posting**, as analyzed by BSW, every seller j is committed to sell at the price p_j he posted regardless of the number of buyers that have approached him. This contrasts with the mechanism of an **auction**, introduced by JKK, according to which a buyer visiting seller j pays p_j if she is the only buyer at j and 1 otherwise, provided she is the one who is served by j. We introduce a new kind of mechanism, called **γ-mechanism**, which encompasses these two widely used mechanisms as special cases.[5]

In a $\boldsymbol{\gamma}$-mechanism, the transaction price at seller j and the allocation of seller j's good are determined as follows: If exactly one buyer approaches seller j, the good is sold to this buyer at price p_j just like with price posting and auctions. If at least two buyers approach seller j, the good is randomly allocated to one of these buyers with equal probability. The buyer who gets the good pays p_j with probability γ, and 1 with probability $1-\gamma$, where $\gamma \in [0,1]$ is common knowledge to all players at the beginning of the game, and common to all sellers. The directed search game with a γ-mechanism can therefore be completely described by the tuple $\langle M, N, \gamma \rangle$. The purpose of introducing γ-mechanisms is that this simplifies the exposition because it allows us to treat price posting and auctions as two opposite extremes of a single mechanism. Assuming that each of the $N-1$ other buyers approach a specific seller whose price is p with probability θ, the expected utility for the remaining buyer of approaching this seller, denoted $U(p,\theta,N-1,\gamma)$, is[6]

$$U(p,\theta,N-1,\gamma) = (1-p)\left[(1-\gamma)(1-\theta)^{N-1} + \gamma\frac{1-(1-\theta)^N}{N\theta}\right]. \tag{1}$$

Plainly, $U(p,\theta,N-1,\gamma)$ is continuous in and decreasing in p and θ. Observe also that

$$U(p,\theta,N-1,\gamma) = \gamma U(p,\theta,N-1,1) + (1-\gamma)U(p,\theta,N-1,0). \tag{2}$$

That is, the expected utility offered to an additional buyer by a seller who is visited by $N-1$ buyers with probability θ is a convex combination of the expected utility offered under price posting, $U(p,\theta,N-1,1)$, and auctions, $U(p,\theta,N-1,0)$, where the weight on price posting is γ. Similarly, under a γ-mechanism the expected profit of a seller who posts the price p and who is visited by each of the N buyers with probability θ, denoted $\Pi(p,\theta,N,\gamma)$, is

$$\Pi(p,\theta,N,\gamma) = Np\theta(1-\theta)^{N-1} + (1-\gamma(1-p))[1-(1-\theta)^N - N\theta(1-\theta)^{N-1}], \tag{3}$$

which is convex combination of the profit under price posting, $\Pi(p,\theta,N,1)$, and auctions, $\Pi(p,\theta,N,0)$, since

$$\Pi(p,\theta,N,\gamma) = \gamma\Pi(p,\theta,N,1) + (1-\gamma)\Pi(p,\theta,N,0). \tag{4}$$

[5]We would like to thank Ian King for suggesting this mechanism.

[6]To see this, notice $U(p,\theta,N-1,\gamma) = (1-p)\left[(1-\theta)^{N-1} + \gamma\sum_{k=1}^{N-1}\binom{N-1}{k}\frac{\theta^k(1-\theta)^{N-1-k}}{k+1}\right]$. Simplifying then yields (1).

Given (2) and (4), γ-mechanisms can therefore be said to be **convex combinations** of price posting and auctions.

While the motivation for introducing the $\boldsymbol{\gamma}$-mechanism is, as mentioned, of expositional nature, this mechanism also offers clear and plausible economic interpretations. Possibly the most straightforward interpretation is that all sellers employ, in principle, an auction, which leaves zero surplus to every buyer approaching a given seller, provided at least two buyers approach this seller. However, with probability γ the buyers who have approached seller j manage to coordinate, that is, to pretend that only one buyer has approached this seller.

In Section 3.2, we will impose the following two restrictions.

Non-Participation We say that a buyer's strategy satisfies **non-participation** if the buyer chooses to abstain – that is, to visit seller Z with probability 1 – if the payoff of visiting any other seller is 0 regardless of what the other buyers do. An equilibrium is said to satisfy non-participation if the strategies of all buyers satisfy non-participation. Clearly, focusing on subgame perfect equilibria, non-participation will only arise in the subgame that ensues upon $\mathbf{p} = \mathbf{1}$. By playing strategies that satisfy non-participation buyers can ensure that sellers set prices below 1 in equilibrium. Non-participation seems a natural restriction. Even absent spite or any other behavioral preferences, abstaining is an optimal action given $\mathbf{p} = \mathbf{1}$, and behavioral preferences will typically make this the uniquely optimal action.

Monotone Buyer Strategies We say that buyer i's strategy is **monotone** if, holding all other prices constant, an increase (decrease) in seller j's price means that she approaches seller j with weakly lower (higher) probability, and approaches any other seller $k \neq j$ with weakly higher (lower) probability. That is, buyer i's strategy is monotone if $\theta^i_j(\mathbf{p})$ is non-increasing in p_j and $\theta^i_k(\mathbf{p})$ is non-decreasing in p_j for all $j, k \in \{A, B, \ldots, M\}$ with $k \neq j$ and for all $\mathbf{p} \in [0, 1]^M$. Restricting buyers' strategies in this way seems fairly reasonable and intuitive: As the price of some seller decreases (increases), each individual buyer's demand for this seller's good is quite naturally expected to increase (decrease).[7] In contrast, if buyers play non-monotone strategies, they will sometimes approach a seller with higher probability after this seller's price has increased.

In the present two-stage game, the strategy set for every seller is simply the $[0, 1]$-interval.

[7] This monotonicity property is also known as the "law of demand".

A pure strategy for seller j is a price $p_j \in [0,1]$.[8] Consequently, any $\mathbf{p} \in [0,1]^M$ corresponds to one specific **subgame**. A strategy for every buyer i is $\boldsymbol{\theta}^i : [0,1]^M \to \Delta$. A **subgame perfect equilibrium**, or equilibrium for short, is a strategy profile such that all p_j and all $\boldsymbol{\theta}^i(\mathbf{p})$ for all $j = A,..,M$ and all $i = 1,..,N$ are mutually best responses, with the additional restriction that $\boldsymbol{\theta}^i(\mathbf{p})$ for all $i = 1,..,N$ constitute mutually best responses for every subgame $\mathbf{p} \in [0,1]^M$. Denoting by $\mathbf{p}^*$ the vector of prices the sellers set in some equilibrium, the subgame $\mathbf{p}^*$ is said to be **on the equilibrium path**. Accordingly, all the subgames $[0,1]^M \setminus \mathbf{p}^*$ are said to be **off the equilibrium path**. Buyer i's strategy is said to be **non-degenerately mixed** on the equilibrium path if $\boldsymbol{\theta}^i(\mathbf{p}^*)$ contains elements that are neither 0 nor 1, and is otherwise called pure (or degenerately mixed) on the equilibrium path. Similarly, we say that buyer i employs a non-degenerately mixed strategy **locally** off the equilibrium path if $\boldsymbol{\theta}^i(\mathbf{p})$ for some $\mathbf{p} \neq \mathbf{p}^*$ in some neighborhood of $\mathbf{p}^*$ contains elements that are neither 0 nor 1. The strategies buyers and seller play on the equilibrium path are sometimes also referred to as the **equilibrium outcome** of the game.

The buyers subgame is a coordination game that typically exhibits multiple equilibria. For example, in the $\langle 2,2,\gamma\rangle$ game, there are exactly three equilibria for any $\mathbf{p}$ such that $1-2(1-p_B)/\gamma \leq p_A \leq 1-\gamma(1-p_B)/2$. These subgames lie within the straight lines depicted in the left hand panel of Figure 1: The two pure strategy equilibria are $\theta_A^1(\mathbf{p}) = \theta_B^2(\mathbf{p}) = 1$ and $\theta_A^1(\mathbf{p}) = \theta_B^2(\mathbf{p}) = 0$ while the unique mixed strategy equilibrium is

$$\theta_j^i(\mathbf{p}) = \frac{2-\gamma-2p_j+\gamma p_{-j}}{(2-\gamma)(2-p_j-p_{-j})} \tag{5}$$

for $i = 1,2$ and $j = A,B$. Foreshadowing parts of our main result, notice that $\theta_j(\mathbf{p})$ satisfies the monotonicity property that it decreases in p_j and increases in p_{-j}. These properties hold for the general $\langle N,M,\gamma\rangle$-model as shown in the proof of the following lemma, which is in the appendix.

Lemma 1 *For any $\mathbf{p} < \mathbf{1}$ a symmetric mixed strategy equilibrium $\boldsymbol{\theta}^* \equiv (\theta_A^*(\mathbf{p}),..,\theta_M^*(\mathbf{p}))$ exists and satisfies the following properties:*

1. *$\theta_j^*(\mathbf{p})$ is continuous in $\mathbf{p}$, decreasing in p_j and increasing in p_k for all $k \neq j$ and all $j = A,..,M$.*

2. *$\theta_j^*(\mathbf{p}) = \frac{1}{M}$ for all $j = A,..,M$ is equivalent to $p_j = p$ for all $j = A,..,M$.*

[8] In line with the literature, we confine attention to equilibria in which sellers play pure strategies.

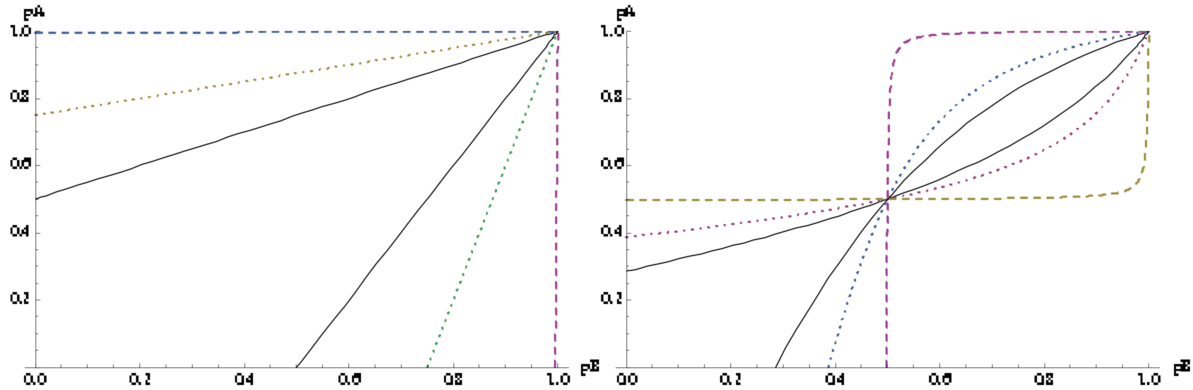

Figure 1: Regions of Buyer Subgames with Multiple Equilibria (lefthand panel) and Seller's Best Replies (righthand panel), assuming buyers play mixed strategies on and locally off the equilibrium path, for $\langle 2, 2, 1\rangle$ (solid lines), $\langle 2, 2, 1/2\rangle$ (dotted) and $\langle 2, 2, 1/100\rangle$ (dashed).

Notice that Lemma 1 is a straightforward generalization of a result first established by Peters (1984) for $\gamma = 1$.

3 Equilibrium Analysis

In this section, we first illustrate that the set of equilibria of the directed search game is large and include both continua of qualitatively similar and qualitatively very different equilibria. Second, we state our main result.

3.1 Families of Equilibria

As is well known, in finite dynamic games like the present one the multiplicity of equilibria in the last stage gives rise to a multiplicity of different equilibria and equilibrium outcomes for the full game. We now briefly illustrate three families of equilibrium outcomes in the present game.

Directed Search Equilibria (DSE) In the symmetric **directed search equilibrium (DSE)** of the game $\langle 2, 2, \gamma\rangle$, both buyers play the mixed strategy given in (5) for all $\mathbf{p}$ such that $1 - 2(1 - p_B)/\gamma \leq p_A \leq 1 - \gamma(1 - p_B)/2$. For $\mathbf{p}$ such that $p_A > 1 - \gamma(1 - p_B)/2$, $\theta_A^i(\mathbf{p}) = 0$ while $\theta_A^i(\mathbf{p}) = 1$ for $p_A < 1 - 2(1 - p_B)/\gamma$ for $i = 1, 2$. Anticipating that buyers play these strategies, seller j's expected profit $\pi_j^{mse}(\mathbf{p}) := \Pi(p_j, N, \theta_j(\mathbf{p}), \gamma)$ for given p_{-j} with $j = A, B$ and $-j \neq j$ is

$$\pi_j^{mse}(\mathbf{p}) = 2p_j\theta_j(\mathbf{p})(1 - \theta_j(\mathbf{p})) + \theta_j(\mathbf{p})^2(1 - \gamma + \gamma p_j), \tag{6}$$

whose unique maximizer $p_j^{mse.BR}(p_{-j})$ is j's best response to p_{-j}, anticipating the mixed strategy equilibrium in the buyers' subgame:

$$p_j^{mse.BR}(p_{-j}) := \frac{4 - 2\gamma - p_{-j}(4 - 6\gamma + \gamma^2) - \gamma(2-\gamma)p_{-j}^2}{8 - \gamma^2 - (8 - 2\gamma - \gamma^2)p_{-j}}. \tag{7}$$

The two best response functions are depicted in the right hand panel in Figure 1 for $\gamma = 1$, $\gamma = 1/2$ and $\gamma = 1/100$. Seller A's best response functions are the convex curves. Observe that the two best replies intersect at $1/2$ for any value of γ, so that $\mathbf{p}^* = (1/2, 1/2)$ is the DSE price vector for the game $\langle 2, 2, \gamma\rangle$. Equation (7) also shows that for any $\gamma > 0$, $p_A^{mse.BR}(1) = 1$, which is also illustrated by Figure 1, while $p_A^{mse.BR}(p_B) = 1/2$ for any $p_B \in [0,1]$ only if $\gamma = 0$. That is, only for auctions will the best responses be completely flat, but the best response functions become less responsive as γ decreases.

For the general $\langle N, M, \gamma\rangle$-model, the symmetric DSE can be derived in the usual manner (see Appendix B and BSW), taking as given that such an equilibrium exists, where symmetry means that all buyers play the same mixed strategy on and locally off the equilibrium path and that all the sellers set the same price.[9] Upon $\mathbf{p} = \mathbf{1}$ all buyers abstain, that is, $\theta_Z^i(\mathbf{1}) = 1$ for all $i = 1, .., N$. The price $p(N, M, \gamma)$ each seller sets in this equilibrium is given by

$$p(N,M,\gamma) = \frac{\mu^N(N-1)N + \gamma\left[\left(-1+\mu^N\right)M^3 + \mu^N(N-1)N - M\left(1+\mu^N(N-1)\right) + M^2\left(2 - 2\mu^N + \mu^N N\right)\right]}{-\mu^N N(M^2 + N - 2M) + \gamma\left[\left(-1+\mu^N\right)M^3 + \mu^N(N-1)N - M\left(1 - (M-1)\mu^N\right) + M^2\left(2 - 2\mu^N + \mu^N N\right)\right]}, \tag{8}$$

where $\mu = \frac{M-1}{M}$. It is readily checked that

$$p(N,M,1) = \frac{1 - \left(1 + \frac{N}{M-1}\right)\left(\frac{M-1}{M}\right)^N}{1 - \left(1 + \frac{N}{M(M-1)}\right)\left(\frac{M-1}{M}\right)^N},$$

which is the DSE price derived by BSW for price posting (see their equation (5)), while

$$p(M,N,0) = \frac{N-1}{(M-1)^2 + N - 1}$$

is the price JKK derive for the sellers' reserve price in the auction mechanism (see their equation (8)).[10] A model that has been of particular interest in the literature is the limit economy as

[9] JKK and Galenianos and Kircher (2012) prove the existence of competitive search equilibria for the special cases of $\gamma = 0$ and $\gamma = 1$. In the $\langle 2, 2, \gamma\rangle$ game, the sellers' best response functions, given the mixed strategy equilibrium in the buyer subgame, can be explicitly solved for and can be shown to have a unique point of intersection.

[10] Two comments are in order. First, in the JKK model the agents whose number is denoted N make the price offer, and the agents whose number is M observe the price offers (and randomize), so when adapting their formula one needs to swap N's and M's. Second, the parameters denoted α and γ in JKK are 0 here and their parameter θ is 1 in our setup.

$N \to \infty$, keeping the ratio with $M/N = k$ fixed at some level $k > 0$. Letting $p_\infty(k,\gamma) := \lim_{N\to\infty} p(N,M,\gamma)$ we get:

$$p_\infty(k,\gamma) = 1 - \frac{1}{1-\gamma+\gamma k\left(e^{\frac{1}{k}}-1\right)},$$

which specializes to $p_\infty(k,0) = 0$ as in JKK and to $p_\infty(k,1) = 1-(1/k)/(e^{1/k}-1)$ as in BSW.

Moreover, for $N > 2$ and $M > 2$ there can also exist asymmetric directed search equilibria. In these DSE, all buyers play mixed strategies on and locally off the equilibrium, but their mixed strategies are not symmetric. For an illustration, consider the $\langle 4,4,\gamma\rangle$-game with sellers $j = A,..,D$. There is an asymmetric DSE in which all sellers set the price $p(2,2,\gamma)$ and, on and locally off the equilibrium path buyers 1 and 2 play the mixed strategy given in (5) for $j = A, B$ $\theta^i_j(\mathbf{p}) = 0$ for $j = C, D$ and $i = 1, 2$ while buyers 3 and 4 play the mixed strategy given in (5) for $j = C, D$ $\theta^i_j(\mathbf{p}) = 0$ for $j = A, B$ and $i = 3, 4$. In this equilibrium, there are thus two submarkets, each consisting of two buyers and two sellers. Upon a small deviation by any seller, all buyers keep playing the mixed strategy defined with respect to their submarkets with "small" meaning small enough for visiting a given seller not to become a strictly dominated strategy.

Trigger Equilibria (TE) In a **trigger equilibrium (TE)** of the $\langle N,N,\gamma\rangle$-game, all buyers play pure strategies on the equilibrium path. With $M = N$, every seller is visited by exactly one buyer with probability 1. Locally off the equilibrium path, buyers play the same mixed strategy $\boldsymbol{\theta}$. In a TE each seller j is successfully deterred from deviating by buyers' credible threat of coordinating on the mixed strategy equilibrium in the second stage if one seller sets a higher price than "prescribed" by the equilibrium. Of course, there are a few restrictions these prices have to satisfy.

Depending on γ and the price p^*_j seller j is supposed to set, a downwards deviation may a priori be profitable if it induces the mixed strategy equilibrium in the buyers subgame. For example, with γ close to 0 and $N = M = 2$ the best response price anticipating the mixed strategy equilibrium, $p^{mse.BR}_j(p_{-j})$, is close to $1/2$ for a large set of p_{-j}'s. However, such deviations can be deterred by buyers by refusing to coordinate on the mixed strategy equilibrium following a downward deviation. Therefore, the deviations that need to be deterred by triggering coordination on the mixed strategy equilibrium are upwards deviations. Consider the $\langle 2,2,\gamma\rangle$-game and let $\pi^{mse+}_j(p^*_j, p_{-j}) := \max_{p_j|p_j\geq p^*_j} \pi^{mse}_j(p_j, p_{-j})$ be j's maximal profit,

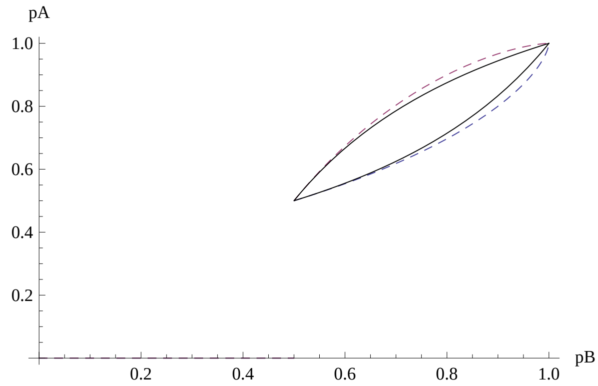

Figure 2: The Sets of TE prices is given by the prices inside the two dashed lines for auctions ($\gamma = 0$) and inside the two solid lines for price posting ($\gamma = 1$) for $N = M = 2$.

anticipating the mixed strategy equilibrium in the buyers' subgame, when best responding to his competitor's price p_{-j} subject to the constraint that his best response price be no less than p_j^*.[11] Solving $p_j^* = \pi_j^{mse+}(p_j^*, \mathbf{p}_{-j})$ for p_j^* yields a function $p_j^*(p_{-j})$ that determines the lower bound of prices j can be induced to set as part of a TE. The sets of price vectors $\mathbf{p}^*$ that can be supported as (part of) the outcome of a TE can be very large. For example, for the $\langle 2, 2, \gamma \rangle$-game the sets of prices that are part of a TE outcome is given by lens-shaped regions as shown in Figure 2 for auctions (dashed lines) and price posting (solid lines). More generally, letting k be an integer no less than 1, TE always exist for $\langle kM, M, \gamma \rangle$-games. TE may also exist when $N \neq kM$, in which case some sellers will be visited by more buyers than others with probability 1.

Mixed Trigger Equilibria (MTE) A priori and depending on N, M and γ, there may also exist what we call **mixed trigger equilibria (MTE)**, which are such that buyers play mixed strategies on the equilibrium path and pure strategies off the equilibrium path. In these equilibria, the price every seller sets is not locally a best response to the mixed strategy buyers play on the equilibrium path (meaning that if buyers also played this mixed strategy locally off the equilibrium path, sellers would be better off by either raising or lowering their price from $\mathbf{p}^*$), and the prices set by other sellers. To deter upward and downward deviations, some buyers play pure off the equilibrium path. For γ close to zero, buyers playing pure can be harmful to sellers because they miss out on more than one buyer approaching, in which case the sellers would extract the full surplus almost with probability one.

[11]The restriction $p_j \geq p_j^*$ is less stringent for γ close to 1 than for γ close to 0.

Hybrid Equilibria (HE) Depending on N, M and γ, one can also construct **hybrid equilibria (HE)**, in which some but not all buyers play mixed strategies on and locally off the equilibrium path. On the equilibrium path, a subset of buyers approach exactly one seller with probability one each, and the others play non-degenerately mixed strategies. As in a DSE, sellers who are not approached with probability one in equilibrium are deterred from raising or lowering prices by these buyers adjusting their mixed strategy. As in a TE, sellers who are approached with probability one in equilibrium are deterred from raising prices by all buyers playing their fully mixed strategies in the ensuing subgame. These sellers are deterred from lowering prices because buyers do not adjust their probabilities of visiting a specific seller locally off the equilibrium path.

As an illustration, consider supporting prices $\mathbf{p}^* = \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)$ in the $\langle 3, 3, 1 \rangle$-game.[12] On the equilibrium path, suppose buyer 1 approaches seller A with probability one, and buyers 2 and 3 mix uniformly between sellers B and C. On the equilibrium path, buyer 1 strictly prefers approaching seller A to B or C, and buyers 2 and 3 strictly prefer mixing between B and C to approaching seller A. Therefore any profitable seller deviation must be a large deviation from equilibrium. It can be shown that in order to attract buyer 2 or 3, seller A must lower his price to at least $p_A = 0$, in which case he would receive no profit. Similarly, if seller B were to lower his price to zero, buyer 1 would be indifferent between approaching seller A or C, both of which would yield a higher utility than approaching B. Therefore no seller could profitably deviate from posting his equilibrium price. Note also that upon an upward deviation by seller A, buyers 2 and 3 must approach seller A with positive probability in order to make buyer 1 indifferent between approaching all three sellers.

Exhaustive Categorization of Equilibria Any candidate equilibrium can be placed into exactly one of four categories, defined by (i) whether or not at least one buyer plays a pure strategy on the equilibrium path, and (ii) whether or not one buyer plays a pure strategy locally off the equilibrium path. These categories are shown in Table 1, with (i) labeled "ON" and (ii) labeled "Locally OFF". The equilibrium in which all sellers set prices equal to one and all buyers play pure strategies on and off the equilibrium path, denoted $\mathbf{p} = \mathbf{1}$, does not satisfy non-participation. The label "MCE" stands for "mis-coordination equilibria", which are described after Proposition 1 below, where it is also shown that these can only exist for

[12]Note that the symmetric DSE price is $p(3, 3, 1) = \frac{7}{15} < \frac{1}{2}$.

	Locally OFF	
ON	All mix	At least one plays pure
All mix	DSE	MTE
At least one plays pure	TE, HE	$\mathbf{p} = \mathbf{1}$, MCE

Table 1: Four categories of candidate equilibria, defined by buyers' strategies on and locally off the equilibrium path.

$\gamma = 0$.

3.2 Monotonicity

We now impose the two restrictions introduced in Section 2 – that buyers abstain whenever the payoff from participating is zero regardless of what the other buyers do and that buyers play monotone strategies – and we show that they yield a powerful uniqueness result.

Proposition 1 *For any $\gamma \in (0, 1]$, any equilibrium that satisfies non-participation and monotonicity is a DSE.*

The reason why the case $\gamma = 0$ is excluded from Proposition 1 is that it can give rise to a plethora of what may be called **mis-coordination equilibria (MCE)**. Consider the game $\langle kM, M, 0\rangle$, where $k \geq 2$ is an integer and denote the set of buyers, in slight departure from the notation used hitherto, by $\{a_1, a_2, ..a_k, b_1, b_2, .., b_k, , .., m_1, m_2, .., m_k\}$. For any $\mathbf{p} \in [0, 1)^M$ and any $i = a, .., m$, let all buyers $i = i_1, .., i_k$ play the strategy $\theta_i^i(\mathbf{p}) = 1$. As these strategies are not responsive to changes in p_i, they clearly satisfy monotonicity. Observe also that every buyer's payoff is 0 when all other buyers play these strategies regardless of $\mathbf{p}$ yet playing $\theta_i^i(\mathbf{p}) = 1$ is not weakly dominated because if all other $k-1$ buyers of label i_j visited seller i with smaller probability, the expected payoff of playing $\theta_i^i(\mathbf{p}) = 1$ would be strictly positive. Notice then that every seller i's payoff is 1 for any $p_i \in [0, 1)$ if the buyers play these strategies. Therefore, for $\gamma = 0$ any $\mathbf{p} \in [0, 1)^M$ is part of the outcome of an equilibrium in which buyers play monotone (and undominated) strategies satisfying non-participation. To the best of our knowledge, the observation that such mis-coordination equilibria can exist for auctions has not been made before. These equilibria are completely knife-edge as they do not exist for any $\gamma > 0$ because then sellers are no longer indifferent between any $p \in [0, 1)$, holding fixed the buyers' probabilities of visiting them, but strictly prefer higher prices.

4 Conclusions

We introduce two new features into the canonical directed search model, γ-mechanisms and buyers' option to remain inactive. The primary purpose of introducing γ-mechanisms is expositional because their introduction permits us to perform a unified analysis of price posting and auctions and, in a well-defined sense, every mechanism in between. However, γ-mechanisms also improve our understanding of price posting and auctions as limit cases. Endowing buyers with the option of not participating seems fairly natural in a market economy and in many other instances too. It allows us to get rid of all the equilibria in which at least one seller posts a price equal to one and is visited with probability one by at least one buyer.

The main observation of our paper is that all equilibria other than directed search equilibria violate a monotonicity property of buyers' strategies that seems fairly natural and intuitive because all these other equilibria require at least one buyer to visit a seller with higher probability after this seller increases his price. Therefore, monotone buyer strategies provide an alternative rationale for focusing on competitive search equilibria. Our monotonicity argument is complementary to the standard rationale for focusing on the competitive search equilibrium, which is based on buyers' difficulties to coordinate. Indeed, if buyers are restricted to play monotone strategies, monotonicity may be the source of what appear to be frictions to coordinate: Provided buyers can only play monotone strategies, the directed search equilibrium would still emerge as the unique equilibrium even if buyers could coordinate, implicitly or explicitly.

Ultimately, the question whether or not buyers play monotone strategies is, of course, an empirical one. Further research that aims at answering this question would be very valuable. An avenue for such research that seems particularly fruitful to us would be to design and run laboratory experiments with a specific focus on buyers' probabilities of visiting sellers as a function of the posted prices and differentiating treatments according to whether or not buyers can communicate.[13]

[13] See Anbarci and Feltovich (2012) for one of the first experimental studies of directed search models à la BSW and Coles and Eeckhout (2003).

A Proofs

Proof of Lemma 1: The existence proof is a straightforward application of Brouwer's fixed point theorem. Let $\boldsymbol{\theta} = (\theta_A, .., \theta_M)$ be a symmetric mixed strategy (excluding seller Z). Notice that $\sum_{j=A}^{M} \theta_j = 1$, so $\boldsymbol{\theta} \in \Delta_{M-1}$, where Δ_{M-1} is the $(M-1)$-dimensional simplex. So as to invoke Brouwer's theorem, first define a continuous function $f : \Delta_{M-1} \to \Delta_{M-1}$. Let $f(\boldsymbol{\theta}) = (f_A(\boldsymbol{\theta}), .., f_M(\boldsymbol{\theta}))$ be the vector valued function whose i-th element is

$$f_i(\boldsymbol{\theta}) = \frac{\theta_i + \max\{0, U(p_i, \theta_i) - \max_{j \neq i} U(p_j, \theta_j)\}}{1 + \sum_{h=A}^{M} \max\{0, U(p_h, \theta_h) - \max_{j \neq h} U(p_j, \theta_j)\}}. \tag{9}$$

where $U(p_j, \theta_j)$ is short-hand for $U(p_j, \theta_j, N-1, \gamma)$. Notice that $f_i(\boldsymbol{\theta}) \geq 0$ and $\sum_{i=A}^{M} f_i(\boldsymbol{\theta}) = 1$. Thus, f maps elements from Δ_{M-1} into Δ_{M-1}. Moreover, f is continuous in θ because the $U(p_j, \theta_j)$'s are continuous in the θ_j's. Thus, a fixed point $\theta^*(\mathbf{p}) = f(\theta^*(\mathbf{p}))$ exists. Such a fixed point is an equilibrium because for every i and j such that $\theta_i^*(\mathbf{p}) > 0$ and $\theta_j^*(\mathbf{p}) > 0$, $U(p_i, \theta_i^*(\mathbf{p})) = U(p_j, \theta_j^*(\mathbf{p}))$ holds while for every h with $\theta_h^*(\mathbf{p}) = 0$, $U(p_h, \theta_h^*(\mathbf{p})) \leq U(p_j, \theta_j^*(\mathbf{p}))$ holds. Thus, a symmetric mixed strategy equilibrium exists.

Continuity of $\theta^*(\mathbf{p})$ is inherited from the continuity of the $U(p_i, \theta_i)$'s in p_i. For any i such that $\theta_i^*(\mathbf{p}) > 0$ before and after the price change, an increase in p_i must result in a decrease of $\theta_i^*(\mathbf{p})$ so as to maintain the indifference condition $U(p_i, \theta_i^*(\mathbf{p})) = U(p_k, \theta_k^*(\mathbf{p}))$ for k such that $\theta_k^*(\mathbf{p}) > 0$. Since a decrease in θ_i^* implies that all the other sellers k must now be visited with (weakly) higher probability, $\theta_k^*(\mathbf{p})$ increases in p_j for $j \neq k$. This completes the proof of 1.

That $p_j = p$ for all j implies $\theta_j^*(\mathbf{p}) = 1/M$ follows from the monotonicity property of $U(p_i, \theta_i)$ in θ_i: If, say, $\theta_i^* < \theta_k^*$, $U(p_i, \theta_i^*) > U(p_k, \theta_k^*)$ would hold, implying $f_i(\theta^*) > \theta_i^*$, a contradiction. Conversely, that $\theta_j^*(\mathbf{p}) = \frac{1}{M}$ for all $j = A, .., M$ implies $p_j = p$ for all j follows from the monotonicity properties of $\theta_j^*(\mathbf{p})$ in p_j and p_k stated in 1. This completes the proof. ■

Proof of Proposition 1: That DSE satisfy monotonicity follows from property 1 of Lemma 1. They satisfy non-participation because $\theta_Z^i(\mathbf{1}) = 1$ for all $i = 1, .., N$. The remainder of this proof shows that no other equilibrium satisfies non-participation and monotonicity.

As shown in the main text, equilibria with $p_j = 1$ for at least one seller j violate non-participation. In the following, we can therefore confine attention to candidate equilibria with $\mathbf{p}^* \ll 1$, where "$\ll$" means that every element of the price vector is strictly less than one.

a) All TE and HE (with $\mathbf{p}^* \ll 1$) are not monotone equilibria: Let buyer 1 be a buyer who derives the (weakly) greatest equilibrium utility of all buyers who play pure strategies on the equilibrium path ($\mathbf{p}^* \ll \mathbf{1}$ and $\gamma > 0$ guarantees that this utility is strictly positive). Let A be the seller she visits with probability 1. Define n_j as the (ex-post or realized) number of buyers who approach seller j on the equilibrium path. Therefore equilibrium utility for buyer i who approaches seller j with probability 1 is $1 - p_j$ if $n_j = 1$, and $\gamma \frac{1-p_j}{n_j}$ if $n_j \geq 2$.

So as to deter small upward deviations by A it must be the case that 1 visits A with probability strictly less than 1 when $p_A > p_A^*$. However since buyers are restricted to monotone strategies, no more buyers can approach seller A after this deviation. As buyer 1 receives the highest equilibrium utility of all buyers who play pure on the equilibrium path, she cannot approach another seller (say, B) who is always approached in equilibrium, because this yields strictly lower utility than approaching A:

$$U(p_B, 1, n_B, \gamma) = \gamma \frac{1 - p_B}{n_B + 1} < U(p_B, 1, n_B - 1, \gamma) \leq U(p_A, 1, 0, \gamma) = 1 - p_A, \tag{10}$$

where $U(p_B, 1, n_B, \gamma)$ is buyer 1's utility from deviating to seller B, $U(p_B, 1, n_B - 1, \gamma)$ is the equilibrium utility of buyers who approach seller B in equilibrium, and $U(p_A, 1, 0, \gamma)$ is buyer 1's equilibrium utility.

We first show that TE are not monotone equilibria. In a TE, there are no sellers who are visited with probability less than 1 on the equilibrium path. If buyers' strategies are monotone, seller A can thus raise his price by a small amount without losing any buyers because of (10). Therefore, TE violate monotonicity.

Consider next HE. In a HE, $\mathbf{p}^*$ must be such that in equilibrium buyer 1 is indifferent between approaching seller A and any seller in a directed search submarket. But then, any individual seller in this submarket could benefit by unilaterally lowering his price by an infinitesimally small amount, attracting an additional buyer into the submarket (for small downward deviations, the increased probability of a sale dominates the effect of a lower selling price). Therefore, monotone buyer strategies eliminate all HE.

b) All mixed trigger equilibria (MTE) are not monotone equilibria:

Take any seller A whose equilibrium price p_A^* is supported by such a buyer strategy, and let $p_A^{\text{mse}}(\mathbf{p}_{-A}^*)$ be his best response to all other sellers announcing prices $\mathbf{p}_{-A}^*$ when all buyers play mixed equilibrium strategies. If $p_A^* < p_A^{\text{mse}}(\mathbf{p}_{-A}^*)$, small upward deviations are deterred by one buyer approaching seller A with probability 1 (call this buyer 1), and all others reducing the probability in which they approach seller A. However buyer 1's strategy is not monotone. If all buyers were to approach seller A with lower probability in this case, then seller A could offer a price $\hat{p}_A \in (p_A^*, p_A^{\text{mse}}(\mathbf{p}_{-A}^*)]$ that yields strictly higher profit than equilibrium (this follows because $p_A^* \neq p_A^{\text{mse}}(\mathbf{p}_{-A}^*)$, so there exists a small profitable deviation if buyers play mixed strategies).

On the other hand, if $p_A^* > p_A^{\text{mse}}(\mathbf{p}_{-A}^*)$, downward deviations must be deterred by a pure strategy.[14] Suppose again that it is buyer 1 who plays pure. All other buyers must be much less likely to approach seller A when he deviates downward (a large increase in buyer 1's probability, and only a small decrease in other buyers' probabilities would generate a profitable downward deviation for the seller, again because $p_A^* \neq p_A^{\text{mse}}(\mathbf{p}_{-A}^*)$). Therefore there must be at least one buyer $i \neq 1$ whose strategy is not monotone.

Finally, note that since the above argument need only consider one seller in isolation, if an equilibrium exists where some (but not all) prices are supported with mixed trigger buyer strategies, then at least one buyer's strategy must not be monotone. All other equilibria are either DSE, TE or HE. This completes the proof. ■

B Derivation of the Symmetric DSE Price

In this appendix, we briefly describe the standard procedure (developed by BSW for $\gamma = 1$) to derive the symmetric DSE price in equation (8), taking as given that such an equilibrium exists. We begin by stipulating that all sellers other then j set the same price p^* and that, on and off the equilibrium path, all buyers play the symmetric mixed strategy, that is, $\boldsymbol{\theta}^i = \boldsymbol{\theta}$ for all $i = 1, .., N$. Letting p_j be the price the price seller j sets and $\theta_j(p_j, p^*)$ the probability with which each buyer approaches seller j, this implies that each of the $M-1$ sellers other than j are approached with probability $(1-\theta_j(p_j, p^*))/(M-1)$. Provided p_j and p^* are such that $\theta_j(p_j, p^*) \in (0,1)$, $\theta_j(p_j, p^*)$ must be such that

$$U_j^i = U^i \quad \text{for all } j \in \{A, B, \ldots, M\} \text{ and all } i \in \{1, 2, \ldots, N\}, \tag{11}$$

where U_j^i is buyer i's expected payoff from visiting seller j and U^i is her expected payoff of visiting any other seller. Unless $\gamma = 0$ as in JKK, one can in general not solve (11) for an explicit expression for $\theta_j(p_j, p^*)$. However, invoking the Implicit Function Theorem, one can determine $d\theta_j(p^*, p^*)/dp_j$.

[14] It cannot be that all buyers who approach A on the equilibrium path reduce their θ_A, otherwise any of these buyers could benefit by unilaterally setting $\theta_A = 1$.

Seller j's profit when setting price p_j all others set price p^*, anticipating the symmetric mixed strategy equilibrium in the buyers subgame, is

$$\pi_j^{mse}(p_j, p^*) = Np_j\theta_j(p_j, p^*)(1-\theta_j(p_j, p^*))^{N-1} + (1-\gamma(1-p_j))\left[1-(1-\theta_j(p_j, p^*))^N - N\theta_j(p_j, p^*)(1-\theta_j(p_j, p^*))^{N-1}\right]. \tag{12}$$

The first-order condition for a profit maximum for j is $\frac{\partial \pi_j^{mse}(p_j, p^*)}{\partial p_j} = 0$. Imposing symmetry, that is $p_j = p^*$, implies $\theta_j(p^*, p^*) = 1/M$ and using in particular the expression for $d\theta_j(p^*, p^*)/dp_j$, one can solve $\frac{\partial \pi_j^{mse}(p^*, p^*)}{\partial p_j} = 0$ for the unique p^* that satisfies this condition. This solution is given by the expression $p(N, M, \gamma)$ displayed in equation (8) in the main text.

An open issue is whether in general the profit function $\pi_j^{mse}(p_j, p^*)$ is quasi-concave, so that the solution to the first-order condition indeed characterizes the best response. That it is quasiconcave can be shown for any $\gamma \in [0, 1]$ for $\langle 2, 2, \gamma \rangle$. For $\gamma = 1$, Galenianos and Kircher (2012, Theorems 1 and Lemma 4) show that this is indeed the case in general and, moreover, that the DSE is unique (see their Lemma 4).[15] For $\gamma = 0$, JKK establish that the profit function is globally concave regardless of N and M.

References

Anbarci, N., and N. Feltovich (2012): "Directed search, coordination failure and seller profits: an experimental comparison of posted pricing with single and multiple prices," *Working Paper, Deakin University.*

Burdett, K., S. Shi, and R. Wright (2001): "Pricing and Matching with Frictions," *Journal of Political Economy*, 108(5), 1060–1085.

Coles, M., and J. Eeckhout (2003): "Indeterminacy and directed search," *Journal of Economic Theory*, 111, 265–276.

Galenianos, M., and P. Kircher (2012): "On Game-theoretic Foundations of Competitive Search Equilibrium," *International Economic Review*, 53(1), 1–21.

[15]The arguments in Galenianos and Kircher (2012) rely on a multiplicative separability of the expected payoff v_j a seller j promises to every buyer if this buyer is served by seller j, and every buyer's probability of being served by seller j conditional on having chosen seller j. This multiplicative separability is not given under a γ-mechanism for $\gamma < 1$ because $v_j = 1 - p_j$ if a single buyer is matched to j and $v_j = \gamma(1 - p_j)$ if multiple buyers visit j.

Guerrieri, V. (2008): "Inefficient Unemployment Dynamics under Asymmetric Information," *Journal of Political Economy*, 116(4), 667–708.

Hawkins, W. B. (forthcoming): "Competitive Search, Efficiency, and Multi-worker Firms," *International Economic Review*.

Julien, B., J. Kennes, and I. King (2000): "Bidding for Labor," *Review of Economic Dynamics*, 3, 619 – 649.

——— (2008): "Bidding for money," *Journal of Economic Theory*, 142, 196 – 217.

Menzio, G., and S. Shi (2010): "Block recursive equilibria for stochastic models of search on the job," *Journal of Economic Theory*, 145, 1453 – 1494.

Moen, E. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105, 385 – 411.

Montgomery, J. (1991): "Equilibrium Wage Dispersion and Interindustry Wage Differentials," *Quarterly Journal of Economics*, 106, 163 – 179.

Peters, M. (1984): "Bertrand Equilibrium with Capacity Constraints and Restricted Mobility," *Econometrica*, 52(5), 1117–1127.

——— (2000): "Limits of Exact Equilibria for Capacity Constrained Sellers with Costly Search," *Journal of Economic Theory*, 95, 139 – 168.

Shimer, R. (2005): "The Assignment of Workers to Jobs in an Economy with Coordination Frictions," *Journal of Political Economy*, 113(5), 996–1025.