

2015

Dynamic Directed Search

Gabriele Camera

Chapman University, camera@chapman.edu

Jaehong Kim

Chapman University

Follow this and additional works at: https://digitalcommons.chapman.edu/esi_working_papers



Part of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Camera, G., & Kim, J. (2015). Dynamic directed search. ESI Working Paper 15-06. Retrieved from http://digitalcommons.chapman.edu/esi_working_papers/155

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Dynamic Directed Search

Comments

Working Paper 15-06

Dynamic Directed Search[†]

Gabriele Camera
Chapman University
University of Basel

Jaehong Kim
Chapman University

March 20, 2015

Abstract

The directed search model (Peters, 1984) is static; its dynamic extensions typically restrict strategies, often assuming price or match commitments. We lift such restrictions to study equilibrium when search can be directed over time, without constraints and at no cost. In equilibrium trade frictions arise endogenously, and price commitments, if they do exist, are self-enforcing. In contrast to the typical model, there exists a continuum of equilibria that exhibit trade frictions. These equilibria support any price above the static price, including monopoly pricing in arbitrarily large markets. Dispersion in *posted* prices can naturally arise as temporary or permanent phenomenon despite the absence of pre-existing heterogeneity.

Keywords: frictions, matching, price dispersion, search

JEL: C70, D390, D490, E390

1 Introduction

The directed search model is a decentralized, general-equilibrium trading environment in which capacity constrained sellers post prices to influence buyers' search decisions (Peters, 1984). The game is played in two stages, over the course of one period. First, prices are posted for everyone to see, then buyers visit a seller of

[†] We thank three anonymous Referees and an anonymous Associate Editor for several helpful comments. G. Camera acknowledges partial research support through the NSF grant CCF-1101627. Correspondence address: Gabriele Camera, Economic Science Institute, Chapman University, One University Dr., Orange, CA 92866; tel.: 714-628-2086; FAX: 714-628-2881; e-mail: camera@chapman.edu.

their choice at no cost, knowing that random rationing is used to meet capacity constraints. Unlike other search models, search is costless and unrestricted and trade frictions arise *endogenously* only if buyers are indifferent where to shop, in equilibrium. For this reason, the model has been adopted to develop insights in the analysis of frictional labor and product markets (Burdett et al., 2001; Julien et al., 2000; Michelacci and Suarez, 2006; Montgomery, 1991; Camera and Selcuk, 2009; Virag, 2010).

The directed search literature has restricted attention to studying equilibrium when buyers follow symmetric strategies because such symmetry supports equilibrium trade frictions (Burdett et al., 2001). A central result is that directed search equilibrium is unique in small and large markets (Kim and Camera, 2014), and is inconsistent with posted price dispersion unless there is pre-existing heterogeneity or costs to make visits. A significant open problem is equilibrium analysis when market interactions are dynamic. For tractability reasons, the dynamic extensions in the literature restrict players' ability to fully exploit the temporal structure of the game. Typical assumptions include history-independent strategies, price commitments or match commitments, exogenous separation shocks, etc.¹ This study characterizes equilibria when these restrictions are lifted. The exercise is meaningful for two reasons. It fills an important gap regarding the type of equilibria that exist under dynamic directed search; for example, one would expect to see

¹See, for instance, Acemoglu and Shimer (1999); Albrecht et al. (2006); Julien et al. (2000); Galenianos and Kircher (2009). Essentially, these assumptions eliminate the need to consider how others would modify their strategies in the continuation game as a reaction to information about deviations. For example, with price commitments a seller who advertises a price below equilibrium does not have to worry that her competitors might react by sharply cutting their prices for the foreseeable future.

many other equilibria—exhibiting some degree of price collusion perhaps—when prices are repeatedly posted for everyone to see (e.g., Cason and Noussair, 2007; Anbarci and Feltovich, 2013). In addition, understanding how the directed search model performs when players can fully exploit the temporal structure of the game can open up its use to a richer set of applications compared to those that can be handled with the standard, static model.

We report that, unlike the typical model, a rich set of equilibria emerges that supports endogenous trade frictions. In these equilibria, match selection *and* duration are both endogenous and price commitments, if they do exist, are self-enforcing. There always exists a continuum of equilibria in which sellers post an identical price that lies in between the static Nash equilibrium price and the monopoly price. The size of the equilibrium set depends on market tightness and discounting. These “collusive” equilibria are supported by the sellers’ threat to play the static Nash equilibrium if any competitor cuts their price. The threat is credible because posted prices are public in the model, and static pricing is always an equilibrium of the dynamic game, but is the one that generates the lowest revenues. As a consequence, we obtain a result that mirrors the one in Diamond (1971); monopoly prices can be supported in markets where many sellers compete for a few buyers, even if buyers’ search is costless and unrestricted by external matching processes. Importantly, this result holds in arbitrarily large markets. On the other hand, buyers cannot induce sellers to cut prices below the static Nash value. Buyers cannot credibly threaten to lower a seller’s payoff by modifying their search behavior, punishing a seller by shopping elsewhere, for

example.

Another unique result is the existence of a continuum of equilibria that support posted price dispersion and trade frictions. These outcomes arise despite the absence of pre-existing heterogeneity or search costs, and can emerge as either a stable or a temporary phenomenon. In equilibrium, sellers who post different prices earn different payoffs because they all attract some buyers who, in equilibrium, are indifferent where they shop. These outcomes can be supported as equilibria because sellers can exploit the dynamic nature of the game reacting to undesirable changes in the distribution of posted prices by aggressively cutting prices in the continuation game. This finding contrasts with what is observed in the typical model, which does not admit trade frictions *and* dispersion in posted prices unless markets can be segmented by search costs or productivity differentials. For example, Montgomery (1991); Galenianos et al. (2011) study equilibria characterized by heterogeneous posted wages that are supported by exogenously different productivities. Burdett et al. (2001) discuss equilibria with dispersion in posted price that do not support trade frictions and, in fact, require buyers' perfect coordination in search strategies. Price-dispersion equilibria with endogenous trade frictions are discussed in Camera and Selcuk (2009), but dispersion in that study involves *realized* prices that can be renegotiated after sellers meet buyers. Finally, posted wage dispersion arises in Kircher (2009) as the market splits into separated sub-markets when buyers can pay a cost to simultaneously visit multiple markets.

The model is applicable to a variety of retail markets for homogeneous goods

in which sellers compete in prices. If prices are transparent, then sellers can tacitly collude by regularly monitoring each other's price. A typical example is offered by the industry for retail gasoline, where gas station operators can coordinate on setting prices above the competitive level by threatening price wars. For example, consider the study in Slade (1992); it found evidence of tacit collusion among gas station operators, with stable and uniform prices during prolonged periods of normal demand (although admittedly, in that model capacity constraints do not play as crucial a role as in directed search). But one can think of other homogeneous product markets in which prices are transparent as in retail gasoline markets, such as retail consumer electronics or airline seats. Yet again, the model is applicable to labor markets in which firms compete in wages—the typical markets considered in the directed search literature (e.g., Albrecht et al., 2006; Julien et al., 2000).

The paper proceeds as follows. Section 2 presents the model and lays out some notation. Section 3 offers some preliminaries involving properties of prices in the static game, which are necessary to derive the results for the dynamic game, presented in Section 4. Section 5 concludes.

2 The model

The model follows the one in Peters (1984). In each period $t = 1, 2, \dots$ there is a constant population of $\mathcal{I} = \{1, \dots, I\}$ anonymous and identical buyers and $\mathcal{J} = \{1, \dots, J\}$ homogeneous sellers each of whom has an indivisible good to sell

in each period. All players are infinitely lived.

In each period t seller j can choose to *post a price* p_j^t . By posting price p_j^t seller j commits to sell at that price in period t to any buyer. However, sellers cannot commit to any future price. Trading at price p_j^t generates utility $v(p_j^t)$ to the buyer of the good, with $v' < 0$. Given the one-to-one relationship between prices and utilities, for convenience we will think of seller j as *promising utility* $v_j^t := v(p_j^t)$ to any buyer.² We will thus interchangeably use the phrases “post a higher (lower) price” or “promise a lower (higher) utility,” when no confusion arises. It is assumed that $v_j^t \in [\underline{v}, \bar{v}]$, where $0 \leq \underline{v} < \bar{v}$. Fixing t , denote the action profile of sellers for the period by $\mathbf{v}^t = (v_1^t, \dots, v_J^t) \in \mathbf{X}^J[\underline{v}, \bar{v}]$; also, let \mathbf{v}_{-j}^t denote \mathbf{v}^t where the j^{th} component is removed.

In each period t buyers choose to visit one seller, based on the utility promised by sellers, in the period. In symmetric equilibrium, the choice of a buyer for a period t is a probability distribution over sellers $(\pi_1^t, \dots, \pi_J^t)$. In any period t , given π_j^t and v_j^t , seller j 's payoff function for the period is $\mathcal{M}(\pi_j^t)\phi(v_j^t)$ where the function $\mathcal{M}(\pi_j^t)$ denotes the probability that seller j trades and the seller's utility function $\phi : [\underline{v}, \bar{v}] \rightarrow \mathbb{R}$ is concave, decreasing, with $\phi(\bar{v}) = 0$. Sellers discount future utility geometrically at rate $\beta \in (0, 1)$.

If a buyer visits seller j in period t , the buyer's payoff is $\mathcal{H}(\pi_j^t)v_j^t$, where the function $\mathcal{H}(\pi_j^t)$ denotes the probability that the buyer trades with seller j , in which case the buyer's utility is v_j^t . Buyers discount future utility geometrically, at rate $\beta_b \in (0, 1)$. Market participants observe all promised utilities and the realization

²Hence, “promised utility” here refers to the utility earned ex-post by a buyer who trades with seller j , not to be confused with the utility a buyer can *expect* ex-ante from visiting seller j .

of demand in their meeting. Sellers' identities are also observable.

3 Preliminaries: the static game

To start, we report some results for the static game, which will be useful to study the repeated game. Consider an outcome in which buyers adopt symmetric strategies. Let $q^i(I, \pi)$ denote the probability that the generic seller j meets $i = 1, \dots, I$ buyers when each of I buyers chooses the seller with probability π . Let $\rho(i) = \frac{1}{i}$ denote a random rationing rule at a seller who has been visited by i buyers. We can thus define

$$\mathcal{M}(\pi) := \sum_{i=1}^I q^i(I, \pi) = 1 - (1 - \pi)^I,$$

$$\mathcal{H}(\pi) := \sum_{i=0}^{I-1} q^i(I-1, \pi) \rho(i+1) = \frac{\mathcal{M}(\pi)}{I\pi} = \frac{1}{I} \sum_{i=0}^{I-1} (1 - \pi)^i.$$

The function $\mathcal{M}(\pi)$ denotes the *unconditional* probability that a seller trades, given that all buyers visit the seller with probability π . The function $\mathcal{H}(\pi)$ is the conditional probability that a buyer trades conditional on visiting a seller, when *every other buyer* visits that same seller with probability π . In symmetric equilibrium, $v_j = v$, $\pi_j = \frac{1}{J}$ for all j , and $\mathcal{M}(\frac{1}{J}) = 1 - (1 - \frac{1}{J})^I$.

We start by defining visiting probabilities in a symmetric Nash equilibrium of the static game.

Definition 1. Given $\mathbf{v} := (v_1, v_2, \dots, v_J)$, the distribution of probabilities $\pi(\mathbf{v})$ in

symmetric equilibrium must satisfy $\sum_{j \in \mathcal{J}} \pi_j(\mathbf{v}) = 1$; if $\pi_j(\mathbf{v}) > 0$ for $j \in \mathcal{J}$, then

$$\mathcal{H}(\pi_j(\mathbf{v}))v_j = \max_{k \in \mathcal{J}} \mathcal{H}(\pi_k(\mathbf{v}))v_k.$$

Since buyers are free to visit any seller, in symmetric equilibrium we need $\mathcal{H}(\pi_j)v_j = \mathcal{H}(\pi_l)v_l \geq \mathcal{H}(0)v_k$, for all $\pi_j, \pi_l > 0$, $\pi_k = 0$. Because we focus on a strongly symmetric equilibrium (as in the literature), let v_{-j} denote the identical strategy of the competitors of seller j . With a small abuse in notation, we may also use $\pi_j(v_j, v_{-j})$ instead of $\pi_j(\mathbf{v})$.

Proposition 1. *Consider a static game. Fix $v_{-j} = x > 0$. There is a unique $v_j(x) \in [\underline{v}, \bar{v}]$ that maximizes seller j 's profit*

$$\Phi(v_j, x) = \mathcal{M}(\pi_j(v_j, x))\phi(v_j),$$

where $\Phi(v_j, x)$ and $\Phi(v_j(x), x)$ are both decreasing in x . In addition, if $v_j(x) = x$ for a unique $x = v^* \in (0, \bar{v})$, then we have

- If $x < v^*$, then $v_j(x) > x$;
- If $x > v^*$, then $v_j(x) < x$.

The proof is in the Appendix. It is well-known that the symmetric equilibrium \mathbf{v} in the static game with homogeneous sellers is unique: $v_j = v^*$ for all j with $0 < v^* < \bar{v}$. Yet, Proposition 1 is helpful because it tells us how a seller would optimally react if his competitors would *all collude* on promising an identical utility x that is different from the equilibrium level v^* . The message is that a

seller should react to competitors who all promise $x > v^*$ by promising a utility below x ; instead, if all competitors promise utility $x < v^*$, then the seller should promise a utility above x .

This result is crucial when we consider the possibility of price collusion—or, equivalently, collusion in promised utilities—in the dynamic game. In the static game, instead, there cannot be collusion and the intuition is as follows: If every competitor attempted to increase their profits by promising a lower utility $v_{-j} = x < v^*$ (i.e., posting a higher price), then a rational seller should promise a utility $v_j(x) > x$ above her competitors' (the converse also holds). This is the central reason why price collusion cannot be sustained in the static game. Things are different once the directed search model is extended to a dynamic environment: here, sellers might wish to react to a deviation by changing their behavior in (part of) the continuation game. The section that follows shows how this can be done.

4 Price collusion in small and large markets

In this section we study sequential equilibrium in the dynamic model. In every period, buyers are free to visit any seller and sellers are free to promise any utility, i.e., there is neither commitment to prices nor to meetings. For the moment, consider outcomes that are stationary and symmetric in the sense discussed before; all sellers behave identically and all buyers behave identically. In such equilibria, the promised utilities v_j and the buyers' choices π_j must be time-invariant. In addition, $v_j = v$ and $\pi_j = 1/J$ for all $j \in \mathcal{J}$ in each period.

We start by defining two intuitive strategies for a generic seller $j \in \mathcal{J}$.

Definition 2 (Static Nash). *In each period the seller promises $v_j = v^*$, i.e., the strategy is (v^*, v^*, \dots) .*

The static Nash strategy is an open loop strategy. The seller ignores information on past pricing behaviors—as if it was not observed—and mechanically repeats static play, day in and day out. Such mechanical behavior has been considered, for instance, in the dynamic direct search extensions of Julien et al. (2000) and Albrecht et al. (2006). The static Nash strategy supports a sequential symmetric equilibrium because playing v^* is always the best response to play of v^* by every other competitor. If static Nash is the strategy adopted, then the seller’s payoff Π^* can be recursively defined by $\Pi^* = \mathcal{M}(\frac{1}{J})\phi(v^*) + \beta\Pi^*$, so that

$$\Pi^* = \frac{1}{1 - \beta} \mathcal{M}(\frac{1}{J})\phi(v^*).$$

Definition 3 (Collusive strategy). *Consider a seller j . In period $t = 1$, the seller promises $v_j^1 = v_c$. In all $t \geq 2$, the seller is in one of two states: colluding or punishing. A colluding seller promises v_c for the period; a punishing seller promises v^* . If the seller is colluding in t , then: (i) If $v_i^t = v_c$ for all $i \in \mathcal{J}$, then the seller keeps colluding in $t + 1$; (ii) otherwise, the seller permanently switches to punishing.*

Such a strategy is typical for models of cooperation in repeated games and is composed of two parts: a rule of desirable behavior (promises v_c) and a rule of punishment (promises v^* forever) that is selected only if a departure from desirable

behavior is observed. Collusion can be supported by the threat of an immediate, permanent and market-wide switch to static Nash play because in the directed search model prices are publicly posted. If $v_c < v^*$, then we interpret v_c as a collusive promised utility—equivalently, as price collusion.

If the collusive strategy is a social norm, i.e., if all sellers adopt it, then

$$\Pi_c := \frac{1}{1-\beta} \mathcal{M}\left(\frac{1}{J}\right) \phi(v_c)$$

denotes the seller's payoff from colluding. We now present a Folk Theorem-type result for the dynamic direct search model.

Theorem 2. *The collusive strategy supports a continuum of symmetric stationary sequential equilibria $v_c \in [\underline{v}, v^*]$. In particular,*

1. $v_c = v^*$ is always an equilibrium;
2. for $v_c \in [\underline{v}, v^*)$, there exists $\beta(v_c) < 1$ such that if $\beta \geq \beta(v_c)$, then the collusive strategy is an equilibrium;

Corollary 3. $v_c > v^*$ is never an equilibrium under the collusive strategy.

Proof of Theorem 2. We must consider the choices of a deviant seller in three cases, which depend on whether v_c is zero, or it is positive and below or above v^* .

Case 1: $0 < v_c \leq v^*$.

Note that $v_c = v^*$ is always the best response to v^* , in each period.

Now consider $0 < v_c < v^*$. We start by discussing choices in equilibrium. Let $v_d(v_c)$ denote the best possible deviation in an equilibrium where v_c is the collusive

promised utility, and let $\pi_d(v_c)$ denote the corresponding probability to visit the deviant seller. We omit the argument v_c when it is understood.

A seller does not defect in equilibrium if

$$\Pi_c \geq \Pi_d := \mathcal{M}(\pi_d)\phi(v_d) + \beta\Pi^* = \mathcal{M}(\pi_d)\phi(v_d) + \frac{\beta}{1-\beta}\mathcal{M}(\frac{1}{J})\phi(v^*).$$

Using the definition for Π_c , the above inequality holds if

$$\beta \geq \beta_c := \frac{\mathcal{M}(\pi_d)\phi(v_d) - \mathcal{M}(\frac{1}{J})\phi(v_c)}{\mathcal{M}(\pi_d)\phi(v_d) - \mathcal{M}(\frac{1}{J})\phi(v^*)}.$$

We have $\beta_c < 1$ for $v_c < v^*$, because $\phi(v_c) > \phi(v^*)$ by the properties of ϕ .

To find the best possible deviation $v_d(v_c)$, we must find the value v_d that maximizes Π_d , i.e., the value that maximizes $\mathcal{M}(\pi_d)\phi(v_d)$, because Π^* is given. Because we focus on strongly symmetric equilibrium, denote by v_{-j} the identical strategy of all sellers other than seller j . From Proposition 1, we know that the maximizer $v_d(v_c)$ is unique for all $v_{-j} = v_c < v^*$ and it is such that $v_d(v_c) > v_c$.

More specifically, the best deviation $v_d(v_c)$ is a solution to

$$\max_{v_j} \mathcal{M}(\pi_j(v_j, v_c))\phi(v_j) \quad \text{s.t.} \quad \mathcal{H}(\pi_j)v_j = \mathcal{H}(\frac{1-\pi_j}{J-1})v_c,$$

where the constraint ensures that buyers are indifferent (= indifference constraint).

The first order condition for an interior solution is

$$\mathcal{M}'(\pi_j)\frac{\partial \pi_j}{\partial v_j}\phi(v_j) + \mathcal{M}(\pi_j)\phi'(v_j) = 0.$$

Using the indifference constraint, we have

$$\frac{\partial \pi_j}{\partial v_j} = -\frac{(J-1)\mathcal{H}(\pi_j)}{\mathcal{H}'(\frac{1-\pi_j}{J-1})v_c + (J-1)\mathcal{H}'(\pi_j)v_j}.$$

A standard result is that $\mathcal{M}(\pi_j) = I\pi_j\mathcal{H}(\pi_j)$, which we can use to rearrange the first order condition together with $\frac{\partial \pi_j}{\partial v_j}$. We obtain that if there exists an interior solution $v_d(v_c)$, then $v_d(v_c) = \frac{\mathcal{H}(\frac{1-\pi_d}{J-1})v_c}{\mathcal{H}(\pi_d)}$ (from the indifference constraint) and $\pi_j = \pi_d \in (0, 1)$, where π_d must solve the rearranged first order condition

$$(J-1)\mathcal{M}'(\pi_d)\phi\left(\frac{\mathcal{H}(\frac{1-\pi_d}{J-1})v_c}{\mathcal{H}(\pi_d)}\right) - I\pi_d\phi'\left(\frac{\mathcal{H}(\frac{1-\pi_d}{J-1})v_c}{\mathcal{H}(\pi_d)}\right) \times \left[\mathcal{H}'(\frac{1-\pi_d}{J-1})v_c + (J-1)\mathcal{H}'(\pi_d)\frac{\mathcal{H}(\frac{1-\pi_d}{J-1})v_c}{\mathcal{H}(\pi_d)}\right] = 0.$$

Otherwise, we have a corner solution $\pi_j = 1$, with $v_d = \frac{\mathcal{H}(0)v_c}{\mathcal{H}(1)}$ which is when the constraint binds, so that the deviant seller gets all buyers.

Finally, consider the optimality of playing v^* out of equilibrium. Out of equilibrium, v^* maximizes the sellers' payoff when every other seller follows the punishment prescribed by the collusive strategy, i.e., $v_{-j} = v^*$. Hence it is never optimal to play $v_j \neq v^*$ out of equilibrium. The proof is by contradiction. Suppose $v_j = \tilde{v} \neq v^*$ is optimal out of equilibrium. Then we must have

$$\Pi^* = \mathcal{M}(\frac{1}{J})\phi(v^*) + \beta\Pi^* \leq \mathcal{M}(\pi(\tilde{v}, v^*))\phi(\tilde{v}) + \beta\Pi^*.$$

But v^* is the unique maximizer in static game, when $v_{-j} = v^*$ (Proposition 1), so $\mathcal{M}(\frac{1}{J})\phi(v^*) > \mathcal{M}(\pi(\tilde{v}, v^*))\phi(\tilde{v})$, which gives us the desired contradiction.

Case 2: $v_c = 0$.

This case differs from the previous one because there is no maximizer $v_d(v_c)$ due to a discontinuity of $\pi_j(v_j, 0)$. This has been discussed in Proposition 1. Hence, suppose that the deviant seller sets $v_j = \epsilon > 0$ when $v_{-j} = 0$. In this case $\pi_j = 1$. Thus the payoff to the deviant seller is

$$\Pi_d(\epsilon) = \mathcal{M}(1)\phi(\epsilon) + \beta\Pi^*.$$

It is suboptimal to deviate in equilibrium, if

$$\Pi_c > \lim_{\epsilon \rightarrow 0} \Pi_d(\epsilon) = \mathcal{M}(1)\phi(0) + \frac{\beta}{1-\beta}\mathcal{M}(\frac{1}{J})\phi(v^*),$$

which can be rearranged as $\beta > \beta_0$, where

$$\beta_0 := \frac{\mathcal{M}(1)\phi(0) - \mathcal{M}(\frac{1}{J})\phi(0)}{\mathcal{M}(1)\phi(0) - \mathcal{M}(\frac{1}{J})\phi(v^*)}.$$

We have $\beta_0 < 1$, because $\phi(0) > \phi(v^*)$. Finally, define

$$\beta(v_c) = \begin{cases} \beta_0, & \text{if } v_c = 0 \\ \beta_c, & \text{if } v_c > 0. \end{cases}$$

Case 3: $v_c > v^*$.

We will show that this cannot be an equilibrium, by means of a contradiction. If in equilibrium a seller deviates to $v_d(v_c)$, then the deviant seller is visited with

probability π_d by any buyer in the period when the deviation takes place. The deviant seller plays v^* in all subsequent periods; hence from then on the seller is visited with probability $\frac{1}{J}$. Since $v_c > v^*$, and $v_d(v_c)$ is the maximizer, we have $\mathcal{M}(\frac{1}{J})\phi(v_c) < \min\{\mathcal{M}(\pi_d)\phi(v_d(v_c)), \mathcal{M}(\frac{1}{J})\phi(v^*)\}$. This implies

$$\Pi_c = \mathcal{M}(\frac{1}{J})\phi(v_c) + \beta\Pi_c < \mathcal{M}(\pi_d)\phi(v_d(v_c)) + \beta\Pi^*,$$

because $\Pi_c = \frac{\mathcal{M}(\frac{1}{J})\phi(v_c)}{1 - \beta} < \Pi^* = \frac{\mathcal{M}(\frac{1}{J})\phi(v^*)}{1 - \beta}$. Hence, if $v_c > v^*$, then there is a profitable deviation $v_d(v_c)$, which gives us the desired contradiction. \square

Corollary 4. *Fix the number of buyers $I = Jr$ with $r > 0$. A version of Theorem 2 holds in the limit as $J \rightarrow \infty$.³*

The lesson is that in the dynamic directed search model, surplus can be easily redistributed from buyers to sellers because directed search is based on public monitoring of prices. Consequently, sellers who are sufficiently patient can collude on *any* price higher than the static Nash price, independent of market tightness. This is a unique result because market tightness is the central determinant of prices in the typical directed search model.

An illustration of Theorem 2 is provided in Figures 1-2. Figure 1 illustrates the mapping between the lower-bound discount factor $\beta(v_c)$, for markets where the number of sellers and the number of buyers vary between 2 and 100, when sellers collude on monopoly prices, $v_c = 0$ and $\phi(v) = 1 - v$ for $v \in [0, 1]$. In symmetric equilibrium $\pi_j = 1/J$ for all sellers $j = 1, \dots, J$ and we obtain by

³The proof of Corollary 4 mirrors the one for small markets. It is contained in an additional appendix available upon request.

direct calculation the utility level that is promised in the static Nash equilibrium

$$v^*(I, J) = \frac{I(J-1)^I}{J^{I+1} - J(J-1)^I - I(J-1)^{I-1}}.$$

The minimal discount factor needed to support monopoly prices, $v_c = 0$, is

$$\beta(0) = \frac{1}{1 + \left[\left(\frac{J}{J-1} \right)^I - 1 \right] v^*(I, J)}.$$

The minimal discount factor grows as the number of sellers increases and as the number of buyers falls. Figure 2 illustrates the value of the static Nash promised utility $v^*(I, J)$ as I and J vary; $v^*(I, J)$ grows as the number of sellers increases and as the number of buyers falls.

FIGURES 1 and 2 APPROX. HERE

Equilibria with promised utilities $v_c < v^*$ can be supported by the (implicit) threat to switch to play the static Nash value v^* as a punishment for any deviation. The static Nash equilibrium gives every seller a lower payoff compared to colluding on v_c ; such a threat is sufficient to deter defections if sellers are adequately patient. As a consequence, we obtain a result akin to the “Diamond paradox” (Diamond, 1971): a large number of sellers can collude on the promised utility $v_c = 0$ (corresponding to the monopoly price) even if they compete for *very few* buyers, and despite the fact that—unlike in Diamond (1971)—there are no search costs; an illustration is the case $I = 2$ and $J = 100$ in Figure 1. This finding sharply

contrasts with the typical directed search model where, on the contrary, monopoly pricing is not an equilibrium, in a finite market (Burdett et al., 2001), and can be an equilibrium in the case of infinite markets only if the queue is arbitrarily large (Albrecht et al., 2006; Camera and Selcuk, 2009; Julien et al., 2000). In fact, it can be shown that in an infinite market the equilibrium set in Theorem 2 can be reduced to just two elements, monopoly and static Nash pricing, which is done in the next subsection.

4.1 Narrowing the equilibrium set in large markets

Consider the case of countable many players with $I = Jr$, when $J \rightarrow \infty$ for $r > 0$, as is typical in the literature. When $\pi = \frac{1}{J}$, we have $I\pi \rightarrow r$, while the matching probabilities are $\mathcal{M}(\pi) \rightarrow 1 - e^{-r}$ and $\mathcal{H}(\pi) \rightarrow \frac{1 - e^{-r}}{r}$. If sellers collude on promising utility v_c , then the seller's payoff is $\frac{1 - e^{-r}}{1 - \beta} \phi(v_c)$. Consider the following strategy that helps players to select the collusive outcome.

Definition 4 (Queue-based collusive strategy). *Amend the strategy in Definition 3 as follows. Let $\mu_{j,t}$ denote the number of buyers expected by seller j in period t . If seller j colludes in t , then: (i) If $\mu_{j,t} \geq r := I/J$, then seller j colludes in $t + 1$; (ii) otherwise, seller j switches to punishing.*

The behaviors described in Definitions 3 and 4 coincide for the case of finite economies, because any deviation in a finite economy affects the queue μ_{-j} at non-deviant sellers. However, this is not true in the limiting case of an infinite market when v_c is positive, so that a defection will not generate an externality to

non-deviant sellers. This characteristic can be exploited by sellers to collude as if they were monopolists, as we demonstrate next.

Proposition 5. *Let $0 \leq \underline{v} < \bar{v}$. The strategy in Definition 4 supports two symmetric stationary sequential equilibria: $v_c = v^*$, always, and $v_c = 0$ if $\beta \in [\beta(0), 1)$.*

The proof is in the Appendix. As the market grows infinitely large, a single deviation from $v_c > 0$ does not affect the payoffs of non-deviant sellers. This is so, because setting $v_j \neq v_c > 0$ cannot change the distribution of buyers at non-deviant sellers, i.e., there is no change in the queue, $\mu_{-j} = r$. However, if $v_c = 0$, then a single deviation $v_j > 0$ can soak up infinitely many buyers away from non-deviant sellers, simply because buyers' payoff is 0. This is true in a finite market and so it holds true in the limit as a market grows infinitely large. In this case, a single deviation to $v_j \neq 0$ changes the queue at all non-deviant sellers from $\mu_{-j} = r$ to $\mu_{-j} < r$ (possibly 0). Consequently, a single deviation from $v_c = 0$ triggers the punishment phase; if players are sufficiently patient, this will deter any deviation.

4.2 Can prices fall below the static equilibrium price?

There are many punishment strategies that let sellers obtain payoffs above the static Nash, apart from the strategy in Definition 3. For instance, sellers can resort to a penal code-type of punishment, whereby sellers revert to collusion after a sufficiently long punishment spell. But, the promised utility cannot end up *above* the static value. There are two reasons for this finding. First, Proposition 1 has demonstrated that there is *always* an incentive to deviate from $v_c > v^*$ because

there exists a best deviation $v_d(v_c) < v_c$ and a corresponding visiting strategy π_d such that $\frac{\mathcal{M}(\pi_d)\phi(v_d)}{1-\beta} > \frac{\mathcal{M}(\frac{1}{J})\phi(v_c)}{1-\beta}$. Second, competitors never have an incentive to punish the best deviation $v_d(v_c) < v_c$ as $\frac{\mathcal{M}(\frac{1-\pi_d}{J-1})\phi(v_c)}{1-\beta} > \frac{\mathcal{M}(\frac{1}{J})\phi(v_c)}{1-\beta}$, i.e., the deviation is profitable to the deviator and his competitors: it generates a positive externality for all non-deviant sellers because their expected demand increases. As such, $v_c > v^*$ cannot be a symmetric equilibrium because sellers have no reason to sustain it. The open question is: could buyers exploit the dynamic structure of the game to motivate sellers to promise utilities *above* the static Nash value v^* ? For example, could buyers successfully threaten to ostracize sellers who deviate from promising a utility $v_c > v^*$?

This threat is not credible since buyers do not have a way to commit to carrying it out (and sellers have no incentive to punish a deviation, as seen above). To see why this is so, suppose that buyers punish a seller who deviates to $0 < v_d < v_c$ by shopping more frequently elsewhere, forever. That is, buyers collectively punish the deviant seller with $0 \leq \tilde{\pi}_d < \pi_d$ forever; this may not be optimal for buyers. Given permanent punishment, the deviant seller promises v_d utility in all subsequent periods, hence his payoff is $\frac{\mathcal{M}(\tilde{\pi}_d)\phi(v_d)}{1-\beta}$. But $\tilde{\pi}_d < \pi_d$ is never optimal for buyers, i.e., they will not ostracize a deviant seller who promises positive utility. Buyers can earn a greater payoff by visiting the deviant sellers since $\mathcal{H}(\tilde{\pi}_d)v_d > \mathcal{H}\left(\frac{1-\tilde{\pi}_d}{J-1}\right)v_c$; this follows from the symmetric equilibrium indifference condition $\mathcal{H}(\pi_d)v_d = \mathcal{H}\left(\frac{1-\pi_d}{J-1}\right)v_c$ and the monotonicity of the matching function \mathcal{H} . Hence, an argument similar to the above demonstrates that $v_c > v^*$ cannot be an equilibrium; a seller can improve his payoff by deviating by optimally choosing

some $0 < v_d < v_c$.

5 Frictional equilibria with price dispersion

The directed search model with homogeneous players does not support equilibrium frictions *and* dispersion in posted prices or, equivalently, in promised utilities. It supports either one, or the other separately, unless there are some exogenous heterogeneity elements, as in Montgomery (1991) and Kim and Camera (2014). For example, in Burdett et al. (2001) dispersion in posted prices may occur if buyers coordinate their actions and *choose not to* direct their search at random. This section demonstrates that equilibrium (posted) price dispersion and trade frictions naturally arise when market participants can interact repeatedly over time. We will show that the degree of price heterogeneity may be time-invariant or not.

To prove this unique finding, we take two steps. To develop intuition, we first show that small degrees of price dispersion are compatible with symmetric equilibrium if sellers can publicly observe the outcome of a random pricing mechanism. The mechanism randomizes prices to be posted over a small, pre-specified interval. We show that this is never an equilibrium in the static model, but it can be an equilibrium when sellers interact repeatedly. In this equilibrium sellers *choose* to ignore price variations in the market, as long as they are sufficiently small and the mean price is sufficiently high. We call these “small dispersion” equilibria.

Then, we remove this public randomization device entirely. In an additional

section, we show that significant degrees of price dispersion can be sustained in equilibrium when sellers behave asymmetrically. However, buyers still play symmetric strategies, optimally choosing to randomize their visits across all sellers in the market. In this sense, equilibrium still exhibits endogenous trade frictions.

5.1 Equilibria with small price dispersion

Suppose for a moment that sellers will not react to price differentials observed on the market as long as they are “small,” i.e., within a certain range of a “target” price. For example, if \$4.97 is the target price, then sellers will not react to prices posted elsewhere as long as they do not exceed \$4.99 or fall below \$4.95. We can think of this as noise in the implementation of the price-posting process; but it may also be in the observation of the price—at this point it really does not matter what the source of noise is.⁴ So, we simply suppose that sellers can calibrate a public randomization device with a mean promised utility value x , when the noise factor $\epsilon_x > 0$ that determines the interval range is exogenously given. Let us denote (v_1, \dots, v_J) the promised utilities that result from this process.

Observe that symmetric equilibrium in the static game is **not** robust to some small price dispersion in the following sense. If sellers promise utility $x \in (0, v^*)$, plus or minus some exogenous noise factor $\epsilon_x > 0$ with zero mean, then the representative seller j should respond by promising something *greater* than everyone else, i.e., $v_j > x + \epsilon_x$ (see Proposition 9 in the Appendix). This result is helpful to prove that in the repeated game equilibria with small price dispersion can, in fact,

⁴For a search model of noise in beliefs about the distribution of prices see Rauh (2001).

be supported. In order to do so, consider a strategy according to which sellers can choose to promise the utility proposed by a public random device that, in each period, generates a mean value $x < v^*$ plus or minus a noise factor up to ϵ_x .

Definition 5 (ϵ -pricing). *Fix a pricing mechanism that in each period t generates a random value $v_j^t \in (x - \epsilon, x + \epsilon)$ for each seller j , with $x < v^*$ and $\epsilon > 0$. In $t = 1$, seller j promises v_j^1 . In all $t \geq 2$, seller j either colludes or punishes. A seller who colludes promises v_j^t ; a seller who punishes promises v^* . If seller j colludes in $t \geq 1$, then: (i) If $v_i = v_i^t \in (x - \epsilon, x + \epsilon)$ for all $i \in \mathcal{J}$, then seller j colludes in $t + 1$; (ii) otherwise, seller j switches to punishing, which is an absorbing state.*

Now consider an outcome in which all sellers adopt the strategy above.

Theorem 6 (ϵ -equilibrium). *Fix $x < v^*$. There exists $\epsilon > 0$ and β_ϵ such that if $\beta > \beta_\epsilon$, then the strategy in Definition 5 supports a price-dispersion equilibrium.*

The Theorem demonstrates that dispersion in posted prices can be supported in symmetric equilibrium as long as dispersion is small. In equilibrium buyers are indifferent where they shop and sellers keep varying prices over time. Hence, promised utilities vary within an interval around x and the degree of dispersion may fluctuate from period to period. But this equilibrium can take other forms. In particular, sellers could choose their initial promised utility and then stick to that value forever, i.e., punishment would occur only if $v_i^t \neq v_i^1 \in (x - \epsilon, x + \epsilon)$. This is a special case in which sellers' strategies are stationary and the market is partitioned into cheap and expensive sellers of a homogeneous product. The following section explores this idea further by showing that repeated directed search can lead to

significant price dispersion, even without public randomizing devices.

5.2 Equilibria with significant price dispersion

Consider a market in which sellers are divided into two groups, according to the price they choose to post. One group is composed of s_L sellers and the other of s_H sellers, where $s_L + s_H = J$. We remain agnostic about what generates this difference and simply presume that it exists (but see later). For convenience, let sellers $(1, 2, \dots, s_L)$ be in the first group and sellers $(s_L + 1, s_L + 2, \dots, J)$ be in the other. Suppose that in each period sellers are free to post any price they wish.

Conjecture that pricing behavior is stationary, so that the promised utility vector in each period is $\mathbf{v} = (v_1, v_2, \dots, v_J)$ where $v_i < v^*$ for all i and

$$v_i = \begin{cases} v_L, & \text{if } i = 1, 2, \dots, s_L \\ v_H \neq v_L, & \text{otherwise.} \end{cases}$$

Conjecture also that it is optimal for a seller to promise utility v^* forever after having observed a deviation from \mathbf{v} . Note that pricing strategies in this conjectured outcome are not symmetric. However, we will retain the focus on symmetric strategies by buyers. Hence, asymmetric behavior occurs in equilibrium only in the seller's game, not in the buyer's game. This implies that, if the equilibrium exists, then it is characterized by trade frictions.

Theorem 7. *There is a continuum of price dispersion equilibria where some sellers promise utility $v_L < v^*$ to any buyer, others promise $v_L \neq v_H < v^*$, and trade*

frictions arise endogenously.

To prove this Theorem we start by considering a pair of sellers, choosing one from each group. Without loss of generality we pick seller 1 and J . Indifference for buyers means that the distribution of demand at each seller must satisfy

$$\mathcal{H}(\pi_1)v_L = \mathcal{H}(\pi_J)v_H, \quad s_L\pi_1 + s_H\pi_J = 1, \text{ and } 0 \leq \pi_1, \pi_J.$$

From the implicit function theorem, for $0 < \pi_1, \pi_J$ we have the following properties:

$$\frac{d\pi_J}{dv_H} = -\frac{\mathcal{H}(\pi_J)}{\mathcal{H}'(\pi_J)v_H + \frac{s_H}{s_L}\mathcal{H}'(\pi_1)v_L} > 0,$$

$$\begin{aligned} \frac{d^2\pi_J}{dv_H^2} &= -\frac{d\pi_J}{dv_H} \times \\ &\frac{\{2(\mathcal{H}(\pi_J))^2 - \mathcal{H}(\pi_J)\mathcal{H}''(\pi_J)\}v_H + 2\mathcal{H}'(\pi_J)\mathcal{H}'(\pi_1)(\frac{s_H}{s_L})v_L + \mathcal{H}(\pi_J)\mathcal{H}''(\pi_1)(\frac{s_H}{s_L})^2v_L}{(\mathcal{H}'(\pi_J)v_H + \frac{s_H}{s_L}\mathcal{H}'(\pi_1)v_L)^2} < 0. \end{aligned}$$

Since $s_L\pi_1 + s_H\pi_J = 1$, we also have $\frac{d\pi_1}{dv_H} < 0$ and $\frac{d^2\pi_1}{dv_H^2} > 0$. This means that if

$\pi_i = \pi_L$ for $i \leq s_L$ and π_H otherwise, then

$$\frac{d\pi_H}{dv_H} > 0 > \frac{d\pi_L}{dv_H}, \quad \text{and} \quad \frac{d^2\pi_L}{dv_H^2} > 0 > \frac{d^2\pi_H}{dv_H^2}.$$

The important step is to ensure that sellers wish to participate in this pricing

scheme. This amounts to verifying that

$$\frac{\mathcal{M}(\pi_i)\phi(v_i)}{1-\beta} > \frac{\mathcal{M}(\frac{1}{J})\phi(v^*)}{1-\beta}, \quad i = 1, \dots, J. \quad (1)$$

If one of these participation constraints does not hold, then v^* cannot be used as a threat to deter deviations. One needs to check that the set of all possible utility pairs (v_L, v_H) that satisfy such participation inequalities is not empty. This is the case; to prove it, note that the set

$$\left\{ (v_1, \dots, v_J) \mid \frac{\mathcal{M}(\pi_i)\phi(v_i)}{1-\beta} > \frac{\mathcal{M}(\frac{1}{J})\phi(v^*)}{1-\beta}, \quad i = 1, \dots, J \right\}$$

is open, and (v_c, \dots, v_c) is in the set if $v_c < v^*$. Moreover $\pi_i > 0$ for $i = 1, J$.

Hence, consider pairs (v_L, v_H) that satisfy the requisite in (1) and study the incentive constraints for sellers, i.e., find parameters such that sellers do not want to deviate. Suppose the best deviation that seller i can do is \tilde{v}_i , $i = 1, J$. We have

$$\tilde{v}_i := \arg \max_x \mathcal{M}(\tilde{\pi}_i(x, \mathbf{v}_{-i}))\phi(x), \quad i = 1, J$$

with $\tilde{\pi}_i + s_L \tilde{\pi}_L + s_H \tilde{\pi}_H - (J - i) \left(\frac{\tilde{\pi}_L - \tilde{\pi}_H}{J - 1} \right) - \tilde{\pi}_H = 1$, $i = 1, J$ where

$$\left\{ \begin{array}{ll} \tilde{\pi}_i = 1, \tilde{\pi}_L = \tilde{\pi}_H = 0 & , \text{ if } \mathcal{H}(1)\tilde{v}_i > \mathcal{H}(0)v_L, \mathcal{H}(0)v_H \\ \tilde{\pi}_i, \tilde{\pi}_k > 0, \tilde{\pi}_l = 0 & , \text{ if } \mathcal{H}(\tilde{\pi}_i)\tilde{v}_i = \mathcal{H}(\tilde{\pi}_k)v_k > \mathcal{H}(0)v_l, k \neq l \in \{L, H\} \\ \tilde{\pi}_i, \tilde{\pi}_L, \tilde{\pi}_H > 0 & , \text{ if } \mathcal{H}(\tilde{\pi}_i)\tilde{v}_i = \mathcal{H}(\tilde{\pi}_L)v_L = \mathcal{H}(\tilde{\pi}_H)v_H. \end{array} \right.$$

No seller deviates if the following inequalities hold

$$\mathcal{M}(\tilde{\pi}_i(\tilde{v}_i, \mathbf{v}_{-i}))\phi(\tilde{v}_i) + \frac{\beta}{1-\beta}\mathcal{M}(\frac{1}{J})\phi(v^*) \leq \frac{1}{1-\beta}H(\pi_i(\mathbf{v}))\phi(v_i), \quad i = 1, J. \quad (2)$$

For $i = 1, J$, if $\beta = 0$, then the left-hand side of (2) is no less than the right-hand side because \tilde{v}_i is the best deviation. For large enough $\beta < 1$, the right-hand side of (2) is greater than the left-hand side; this is immediate from (1). Therefore there are $\beta_1 < 1$, and $\beta_J < 1$ that satisfy (2) with equality for each seller i .

To conclude the proof of Theorem 7, let $\bar{\beta} := \max\{\beta_1, \beta_J\}$. If $\beta > \bar{\beta}$, then \mathbf{v} is an equilibrium. Moreover, there is an open set $O(\mathbf{v})$ around \mathbf{v} such that $\mathbf{v}' \neq \mathbf{v}$ can also be sustained as an equilibrium for $\mathbf{v}' \in O$. That is, there is a continuum of equilibria that support trade frictions and dispersion in posted prices.

It is possible to derive explicit bounds for the v_H/v_L ratio in the limit as the market grows large.⁵ To do so, consider a large market with market tightness λ and an equilibrium in which some sellers act as monopolists—promising the lowest feasible utility level $\underline{v} > 0$ —while others promise a greater utility level. We want to determine the maximal difference in promised utilities that can be supported as an equilibrium. Let η_L be the proportion of sellers who promise utility $v_L = \underline{v}$, and $1 - \eta_L$ be the proportion of sellers who promise $v_H \in [v_L, v^*]$. Denote the respective queues as λ_H and λ_L , with $\lambda = (1 - \eta_L)\lambda_H + \eta_L\lambda_L$. We wish to find the upper bound of $\frac{v_H}{v_L}$. In a large market $\mathcal{M}(\lambda) = 1 - e^{-\lambda}$ and $\mathcal{H}(\lambda) = \frac{1 - e^{-\lambda}}{\lambda}$.

⁵We thank a referee for making this suggestion.

By the participation constraints in (1), we must have

$$(1 - e^{-\lambda_j})\phi(v_j) > (1 - e^{-\lambda})\phi(v^*), \quad j = H, L.$$

Since the expected profit $(1 - e^{-\lambda})\phi(v)$ is quasi-concave in v , the solution of the inequality for $j = H$ is convex. Let $[\underline{v}_H, \bar{v}_H] \subset [\underline{v}, v^*]$ be the solution of

$$(1 - e^{-\lambda_H})\phi(v_H) \geq (1 - e^{-\lambda})\phi(v^*).$$

By construction we have $\underline{v}_H = \underline{v}$; this is so because if $v_H = v_L = \underline{v}$, then $\lambda_H = \lambda$, hence the expected payoff is greater than $(1 - e^{-\lambda})\phi(v^*)$ for seller H .

Now consider the inequality for $j = L$. To derive the upper bound of $\frac{v_H}{v_L}$ given $v_L = \underline{v}$, we claim that we must find the value $v_H = v_H^*$ that satisfies

$$(1 - e^{-\lambda_L})\phi(\underline{v}) = (1 - e^{-\lambda})\phi(v^*).$$

To see why this equality should hold, start by noticing that λ_L is a decreasing function of v_H . This is so, because equilibrium buyers' indifference implies $\frac{1 - e^{-\lambda_H}}{\lambda_H}v_H = \frac{1 - e^{-\lambda_L}}{\lambda_L}v_L$. This expression reveals that λ_L is decreasing in v_H . Hence, the upper bound of $\frac{v_H}{v_L}$, is $\min\left(\frac{\bar{v}_H}{\underline{v}}, \frac{v_H^*}{\underline{v}}\right)$.

Theorem 7 can be immediately extended to show that a continuum of price dispersion equilibria exists also for *any* number $n \leq J$ of sellers' groups.

Corollary 8. *There is a continuum of price dispersion equilibria with an associated promised utility vector $\mathbf{v} = (v_1, \dots, v_J)$, where $v_i < v^*$ for all i , and trade*

frictions arise endogenously.

Once can demonstrate that price dispersion equilibria could leave sellers indifferent to posting prices that are lower than their competitors. To see why, note that the equilibria exist in which sellers vary their prices over time in a pre-specified manner. Therefore, one could construct equilibria in which some sellers front-load their expected profits by posting prices higher than the average, while their competitors back-load their profits by raising their prices at some point in time.

The findings reported in Theorem 7 and Corollary 8 are unique in the literature on directed search. Equilibrium heterogeneity in posted wages emerges in directed search models that assume ex-ante heterogeneity in the value of matches (Montgomery, 1991; Galenianos et al., 2011; Kim and Camera, 2014). The price dispersion reported in Burdett et al. (2001), instead, is inconsistent with the existence of frictions; in equilibrium buyers do not choose sellers at random and, in fact, must coordinate their search strategies—which is why the literature has shied away from studying these equilibria. Price-dispersion equilibria with endogenous frictions are discussed in Camera and Selcuk (2009), but that paper refers to prices that are renegotiated after sellers meet their customers; price dispersion arises in Kircher (2009) when there is ex-post market segmentation because buyers can visit multiple sellers, at a cost.

5.3 Collusion under imperfect monitoring

The results reported above hinge on the assumption of unfettered public monitoring of price deviations. A natural question concerns the robustness of collusion and price dispersion equilibria when price deviations are *imperfectly* observed. This may be especially relevant as markets grow large. This section addresses such question by presenting an extension of our baseline environment, which introduces a form of imperfect monitoring of price deviations.

Ex-ante homogeneous sellers experience production-cost shocks in each period.⁶ Each seller now experiences either a high or a low cost shock in a period. Let the subscript $j = H, L$ denote the type of seller for the period, and let $J_{j,t}$ denote the number of type- j sellers in period t . Assume $\frac{J_{j,t}}{J} = a_j$ for all t , where $a_H + a_L = 1$; i.e., seller types are in fixed proportion and there are no aggregate shocks. Assuming that cost shocks are uniformly distributed across sellers, a_j is a seller's probability of having cost j in a period. Finally, we let $\phi_j(v)$ denote the profit function of a seller who has cost shock $j = H, L$ in that period. It is helpful to consider $\phi_L(v) = \phi_H(v) + C$, where the cost difference is $C > 0$.

To model imperfect information about price deviations, assume that sellers cannot directly observe prices and costs in the market. They can only observe the **public** (and truthful) report of an external authority that in each period monitors the prices and cost shocks of $1 \leq s \leq J$ sellers. This monitoring process is subject to frictions, because s is a random variable determined as follows. The monitoring process is sequential, is done in random order and is subject to breakdowns.

⁶We thank an anonymous referee for suggesting an extension along these lines.

The authority starts the process by monitoring an initial seller, chosen at random among all J sellers. The process continues with probability $q \in [0, 1)$, and otherwise it stops. In this manner, the variable q captures the existence of a monitoring friction; the frictions increases as q falls. Fixing a seller, the probability that this seller's price and cost are monitored is

$$p(q, J) := (1 - q)\frac{1}{J} + \dots + q^{J-2}(1 - q)\frac{J-1}{J} + q^{J-1}\frac{J}{J} = \frac{1}{J} \cdot \frac{1 - q^J}{1 - q}.$$

In the formula above, $q^{k-1}(1 - q)$ is the probability of monitoring exactly $1 \leq k \leq J$ sellers; k/J is the probability of monitoring a specific seller, given k total observations. Note that the probability $p(q, J)$ of monitoring any given seller falls as the number of sellers J grows. This captures the notion that monitoring price deviations becomes harder and harder as markets grow large; indeed, $p(q, J)$ approaches zero as J approaches infinity.

Since the proportion of sellers' types is time-invariant, for each type $j = H, L$ we let v_j^* denote the time-invariant static Nash equilibrium promised utility, π_j^* the corresponding demand, while the payoff (using $-j \neq j$) is denoted

$$\begin{aligned} \Pi_j^* &= \mathcal{M}(\pi_j^*)\phi_j(v_j^*) + \beta(a_j\Pi_j^* + a_{-j}\Pi_{-j}^*), \\ \Rightarrow \Pi_j^* &= \frac{(1 - \beta a_{-j})\mathcal{M}(\pi_j^*)\phi_j(v_j^*) + \beta a_{-j}\mathcal{M}(\pi_{-j}^*)\phi_{-j}(v_{-j}^*)}{1 - \beta}. \end{aligned}$$

Now suppose that sellers promise an identical utility level $v_c < \min\{v_H^*, v_L^*\}$, which

is independent of their type in the period. By the first order condition,

$$\begin{aligned} 0 < \mathcal{M}'\left(\frac{1}{j}\right)\pi'(v_c)\phi_H(v_c) + \mathcal{M}\left(\frac{1}{j}\right)\phi'_H(v_c) < \mathcal{M}'\left(\frac{1}{j}\right)\pi'(v_c)[\phi_H(v_c) + C] + \mathcal{M}\left(\frac{1}{j}\right)\phi'_H(v_c) \\ = \mathcal{M}'\left(\frac{1}{j}\right)\pi'(v_c)\phi_L(v_c) + \mathcal{M}\left(\frac{1}{j}\right)\phi'_L(v_c). \end{aligned}$$

The first inequality comes from Proposition 1; the second comes from the fact that $\mathcal{M}'\left(\frac{1}{j}\right)\pi'(v_c)C > 0$. The immediate implication is that, in an outcome in which all sellers promise v_c , low-cost sellers have a greater incentive to deviate (by raising their promised utility) compared to high-cost sellers. Given such differential in incentives, let us consider a collusive strategy according to which sellers' promised utilities depend on their production cost. In this scenario sellers alternate between promising a high or low utility, v_H or v_L , based on their cost for the period; such type of collusion is supported by the threat of reverting to playing the static Nash equilibrium if a deviation is made public by the monitoring authority.

Specifically, let low cost sellers promise higher utility than high-cost sellers with $v_H < v_L < \min\{v_H^*, v_L^*\}$. Let π_j be the corresponding demand for type j .

The payoff is

$$\begin{aligned} \Pi_j &= \mathcal{M}(\pi_j)\phi_j(v_j) + \beta(a_j\Pi_j + a_{-j}\Pi_{-j}), \\ \Rightarrow \Pi_j &= \frac{(1 - \beta a_{-j})\mathcal{M}(\pi_j)\phi_j(v_j) + \beta a_{-j}\mathcal{M}(\pi_{-j})\phi_{-j}(v_{-j})}{1 - \beta}. \end{aligned}$$

Consider a one-time deviation by a seller $j = H, L$. Let v_{dj} and π_{dj} be the best deviation and corresponding demand. This deviation is detected and made public

with probability $p(q, J)$. Hence, the collusive strategy is individually optimal if

$$\begin{aligned} & \mathcal{M}(\pi_{dj})\phi_j(v_{dj}) + p(q, J)\beta(a_H\Pi_H^* + a_L\Pi_L^*) + [1 - p(q, J)]\beta(a_H\Pi_H + a_L\Pi_L) \\ & \leq \Pi_j = \mathcal{M}(\pi_j)\phi_j(v_j) + \beta(a_H\Pi_H + a_L\Pi_L). \end{aligned}$$

Rearrange the inequality as

$$\frac{\mathcal{M}(\pi_{dj})\phi_j(v_{dj}) - \mathcal{M}(\pi_j)\phi_j(v_j)}{a_H(\Pi_H - \Pi_H^*) + a_L(\Pi_L - \Pi_L^*)} \leq p(q, J)\beta, \quad (3)$$

where

$$\begin{aligned} \Pi_j - \Pi_j^* &= \frac{(1 - \beta a_{-j})[\mathcal{M}(\pi_j)\phi_j(v_j) - \mathcal{M}(\pi_j^*)\phi_j(v_j^*)]}{1 - \beta} \\ &+ \frac{\beta a_{-j}[\mathcal{M}(\pi_{-j})\phi_{-j}(v_{-j}) - \mathcal{M}(\pi_{-j}^*)\phi_{-j}(v_{-j}^*)]}{1 - \beta}. \end{aligned}$$

The left hand side of inequality (3) remains bounded away from zero even if β approaches zero. Hence if $p(q, J)\beta$ is small enough, then the above inequality is not satisfied. It follows that the proposed collusive strategy cannot be sustained as an equilibrium when agents are impatient.

Note also that $\Pi_j - \Pi_j^*$ for $j = H, L$ diverge to infinity as β approaches one. Hence, the left hand side of inequality (3) converges to zero as β approaches one. Therefore, for any given probability of detection $p(q, J) > 0$, there always exists $\beta(q, J) < 1$ such that the inequality (3) is satisfied for any $\beta \geq \beta(q, J)$. Note that because the probability $p(q, J) > 0$ falls in J , the lower bound discount factor $\beta(q, J)$ increases with J , converging to one as J approaches infinity.

We conclude that some collusion can still be sustained as an equilibrium even if sellers have heterogeneous and imperfectly observable cost shocks, as long as sellers are sufficiently patient. However, as the number of sellers grows large, it is harder and harder to sustain collusion because the incentive to defect grows stronger, due to the imperfect ability to detect a defection. Supporting collusion is also more difficult as the monitoring process becomes prone to greater frictions (q falls).

6 Final comments

The directed search model was originally conceived as a static game but it has been used to study markets that are inherently dynamic, such as labor markets. This paper breaks new ground in the study of dynamic economies where search can be directed without constraints and costs. We have proved that there generally exists a continuum of equilibria exhibiting trade frictions. Monopoly pricing is supported in small as well as arbitrarily large markets, even if search is costless and unrestricted. Moreover, price dispersion, which is unsustainable in the typical model, can naturally emerge as an equilibrium phenomenon when search is repeatedly directed. Cyclical price movements can be supported, also. These findings can help pushing forward the study of decentralized markets where trade frictions emerge endogenously, as an equilibrium phenomenon.

Allowing free-entry of sellers would not make collusion impossible. Let sellers pay a fixed cost to enter the market and suppose that every entrant promises

$v_c < \bar{v}$. The number of sellers who enter corresponds to the value J that supports a zero payoff net of entry costs. There is an associated static Nash promised utility $v^*(J)$, which increases in J , reaching the value \bar{v} as J grows large. By continuity we can always find parameters such that $v_c < v^*(J)$. Free-entry does not rule out the possibility of collusion because sellers can sustain v_c by resorting to the implicit threat of playing $v^*(J)$ as a response to any defection. However, free-entry bounds v_c away from zero because of the wedge presented by the entry cost.

Note also that a threat of reversion to the static Nash equilibrium price cannot generally prevent free entry. To see this, consider a second scenario. Let there be $J \geq 2$ incumbent sellers in the market who promise utilities $v_c < v^*(J)$ and then open up the market to free entry. Suppose it is profitable to enter, given $v^*(J)$. Note that the static Nash equilibrium utility $v^*(J + x) > v^*(J)$ for all $x \geq 1$. Consider a strategy such that the incumbent sellers threaten to play $v^*(J + x)$ if $x \geq 1$ new sellers enter the market. Such an action cannot keep out potential entrants because it does not represent a threat for them, as the new entrants would make zero profits by staying out of the market.

References

- Acemoglu, D., and R. Shimer (1999). Efficient Unemployment Insurance. *Journal of Political Economy*, 107, 893-928.
- Albrecht J., P. A. Gautier, and S. Vroman (2006). Equilibrium directed search with multiple applications. *Review of Economic Studies*, 73(4), 869-891.

- Anbarci, N., and N. Feltovich (2013). Directed search, coordination failure and seller profits: An experimental comparison of posted pricing with single and multiple prices. *International Economic Review* 54, 873-884.
- Burdett K., S. Shi, and R. Wright (2001). Pricing and matching with frictions. *Journal of Political Economy* 109 (5), 1060-1085.
- Camera, G., and J. Kim (2013). Buyer's equilibrium with capacity constraints and restricted mobility: A recursive approach. *Economics Letters* 118, 321-323.
- Camera, G., and C. Selcuk (2009). Price Dispersion with Directed Search. *Journal of the European Economic Association* 7(6), 1193-1224.
- Cason, T., and C. Noussair (2007). A market with frictions in the matching process: an experimental study. *International Economic Review* 48(2), 665-691.
- Diamond, P. (1971). A Model of Price Adjustment. *Journal of Economic Theory*, 3, 156-168.
- Galenianos, M., and P. Kircher (2009). Directed Search with Multiple Job Applications. *Journal of Economic Theory*, 133, 445-471.
- Galenianos, M., P. Kircher, and G. Virág (2011). Market power and efficiency in a search model. *International Economic Review* 52 (1), 85-103.
- Julien, B., Kennes J., and King I. (2000). Bidding for Labor. *Review of Economic Dynamics*, 3, 619-649.
- Kim, J., and G. Camera (2014). Uniqueness of equilibrium in directed search models. *Journal of Economic Theory* 151, 248-267.

- Kircher, P. (2009). Efficiency of Simultaneous Search, *Journal of Political Economy*, 117 (5), 861-913.
- Michelacci, C., and J. Suarez (2006). Incomplete Wage Posting. *Journal of Political Economy*, 114, 1098-1123.
- Montgomery, J. D. (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.
- Rauh, M. T. (2001). Heterogeneous beliefs, price dispersion, and welfare-improving price controls. *Economic Theory* 18(3), 577-603.
- Slade, M. (1992). Vancouver's Gasoline Price Wars: An Empirical Exercise in Uncovering Supergame Strategies. *Review of Economic Studies* 59, 257-276.
- Virág, G. (2010). High profit equilibria in directed search models. *Games and Economic Behavior* 71, 224-234.

Appendix

Proof of Proposition 1. Let $v_{-j} = x > 0$.⁷ A seller j maximizes the payoff

$$\Phi(v_j, x) = \mathcal{M}(\pi_j(v_j, x))\phi(v_j).$$

Note that $\pi_j(v_j, x)$ and $\phi(v_j)$ are continuous in $v_j \in [\underline{v}, \bar{v}]$, and $\mathcal{M}(\pi)$ is continuous in $\pi \in [0, 1]$. In addition Φ is strictly concave in v_j (Camera and Kim, 2013). Hence a unique maximizer $v_j(x)$ exists in $[\underline{v}, \bar{v}]$. We also have

$$\frac{\partial \Phi(v_j, x)}{\partial v_{-j}} = \mathcal{M}'(\pi_j) \frac{\partial \pi_j(v_j, x)}{\partial v_{-j}} \phi(v_j).$$

We have $\mathcal{M}'(\pi_j) > 0$ (Camera and Kim, 2013). Defining $\pi_{-j} := \frac{1 - \pi_j}{J - 1}$, we get

$$\frac{\partial \pi_j(v_j, x)}{\partial v_{-j}} = \frac{\mathcal{H}(\pi_{-j})}{\mathcal{H}'(\pi_j)v_j + \frac{1}{J-1}\mathcal{H}'(\pi_{-j})x} < 0.$$

Hence $\frac{\partial \Phi(v_j, x)}{\partial v_{-j}} < 0$.

To prove that $\Phi(v_j(x), x)$ decreases in x , we use the envelope theorem:

$$\left. \frac{d\Phi(v_j(v_{-j}), v_{-j})}{dv_{-j}} \right|_{v_{-j}=x} = \mathcal{M}'(\pi_j) \frac{\partial \pi_j(v_j(x), x)}{\partial v_{-j}} \phi(v_j(x)) < 0,$$

because we can treat v_j as a constant when $v_j(x)$ is a maximizer.

The arg max function $v_j(x)$ is continuous from Berge's theorem. Thus define $f(x) = v_j(x) - x$. Note that $f(x) = 0$ only at $x = v^* \in (0, \bar{v})$ (Galenianos and Kircher, 2009). That is, if $v_{-j} = x = v^*$ then the maximizer is $v_j(v^*) = v^*$.

Suppose $v_{-j} = x < v^*$. We prove by contradiction that $v_j(x) > x$. Suppose $v_j(x) \leq x$ for $x < v^*$. Consider the lower bound \underline{v} . We have two cases, $\underline{v} > 0$ and $\underline{v} = 0$.

- Let $\underline{v} > 0$. If $x = \underline{v}$, then $v_j(\underline{v}) > \underline{v}$ maximizes the profit of seller j . To see this notice that there is a unique maximizer but it is not $v_j = x = \underline{v}$, because $v_j = x$ is the maximizer only when $x = v^*$. It follows that $f(\underline{v}) > 0$. Continuity of f and the intermediate value theorem imply that $f(x) > 0$ for all $x < v^*$ which gives us the desired contradiction.
- Let $\underline{v} = 0$. If $x = \epsilon > 0$ small, then seller j can improve the payoff by capturing the entire demand setting $v_j = \tilde{v}_j(\epsilon) > \epsilon$. We claim that the maximizer is $v_j(\epsilon) > \epsilon$. By means of contradiction, suppose $v_j(\epsilon) = \eta \leq \epsilon$.

⁷If $x = 0$, for example because $x = \underline{v} = 0$ then $\pi_j(v_j, 0)$ is not continuous in v_j . Seller j can capture all demand by promising utility $v_j > 0$. That is, for small $v > 0$, $\Phi(0, x) < \Phi(v, x) < \Phi(\frac{v}{2}, x)$. So, there is no maximizer for $x = 0$. This is why we consider $x > 0$.

Then we have

$$\Phi(\tilde{v}_j(\epsilon), \epsilon) > \Phi(0, 0) > \Phi(\eta, \eta) \geq \Phi(\eta, \epsilon),$$

which contradicts that $v_j(\epsilon) = \eta$ is a maximizer. It follows that since $v_j(\epsilon) > \epsilon$ is the maximizer, then $f(\epsilon) > 0$, and with the same argument as the previous case, we have $f(x) > 0$ for all $\epsilon < x < v^*$ for all small $\epsilon > 0$. Hence $v_j(x) > x$ for all $x \in (0, v^*)$.

The proof that if $v_{-j} = x > v^*$ then $v_j(x) < x$ is similar. \square

Proof of Proposition 5. Consider the collusive strategy in Definition 4. Consider seller j . Define μ_j the queue at this seller, hence the period payoff is $(1 - e^{-\mu_j})\phi(v_j)$. In equilibrium $v_j = v_c$, $\mu_j = r$.

We start by showing that $0 < v_c \neq v^*$ cannot be an equilibrium. By means of contradiction, suppose $0 < v_j = v_c \neq v^*$ is equilibrium. In this case the payoff to the seller is

$$\frac{(1 - e^{-r})\phi(v_c)}{1 - \beta} < (1 - e^{-\lambda^*})\phi(v(v_c)) + \frac{\beta}{1 - \beta}(1 - e^{-r})\phi(v_c),$$

where $v(v_c) \neq v_c$ is the best one-time deviation. Note that under the deviation $v(v_c)$, we have $\mu_j = \lambda^* > 0$ and $\mu_{-j} = r$ for all sellers other than j ; this follows from buyers' indifference. This means that the punishment phase is not triggered. Therefore $\mu_j = r$ after the one-time deviation.

If $v_c = 0$, then any small deviation $v_j = v_d > 0$ improves seller j 's period payoff. In this case, buyers' indifference is satisfied for $\lambda^* = \infty$ and any $\mu_{-j} \leq r$. Note that in a finite market a seller who deviates from $v_c = 0$ captures entire demand. Hence we assume this is also true in the limiting case of infinite market, so non-deviant sellers have $\mu_{-j} = 0 < r$. The modified collusive strategy implies that all sellers switch to punishing by promising v^* . Hence $v_c = 0$ is a sequential equilibrium if

$$\lim_{v_d \rightarrow 0} \phi(v_d) + \frac{\beta}{1 - \beta}(1 - e^{-r})\phi(v^*) \leq \frac{1}{1 - \beta}(1 - e^{-r})\phi(0)$$

Equivalently, if

$$\beta \geq \beta(0) := \frac{e^{-r}\phi(0)}{\phi(0) - \phi(v^*) + e^{-r}\phi(v^*)}.$$

Therefore with the modified strategy in Definition 4, there is no continuum of equilibria, but only two possible symmetric equilibria, $v_c = 0, v^*$. \square

Proposition 9. Consider a static game. Fix a seller j and a promised utility $x \in (0, v^*)$. There exists an exogenous noise factor $\epsilon_x > 0$ such that if each seller

$i \neq j$ promises utility $v_i \in (x - \epsilon_x, x + \epsilon_x)$, then the unique maximizer for seller j is to promise something greater than everyone else, i.e., $v_j > x + \epsilon_x$.

Proof of Proposition 9. We present a proof by recursion. Without loss of generality, let $j = J$. By Proposition 1, if every seller $i < J$ promises utility $0 < x < v^*$, then a unique maximizer is $v_J > x + \delta_1$ for some $\delta_1 > 0$. Now let all sellers $2 \leq i < s$ promises x and seller 1 promises $v_1 \in (x - \epsilon_2, x + \epsilon_2)$. By continuity of best response v_J , there is a $\epsilon_2 > 0$, such that seller J has a unique maximizer $v_J > x + \delta_2$ for some $\delta_2 \in (0, \delta_1]$. Now let all sellers $3 \leq i < J$ promise x and let seller 1 promise $v_1, v_2 \in (x - \epsilon_3, x + \epsilon_3)$. By continuity, there is a $\epsilon_3 \in (0, \epsilon_2]$ such that seller J has a unique maximizer $v_J > x + \delta_3$ for some $\delta_3 \in (0, \delta_2]$. Repeating this argument for all $i < J$ and denoting $\epsilon_x := \min\{\delta_J, \epsilon_J\}$, we prove this claim. \square

Proof of Theorem 6. Start by noticing that, for any $x \in (0, v^*)$, we can find $\epsilon > 0$ such that Proposition 9 is satisfied. In addition continuity of demand $\pi(v_j, v_{-j})$ implies that if $v_j = x - \epsilon$ and $v_{-j} = x + \epsilon$, then $\pi_j > 0$. See the results in Camera and Kim (2013)

Now consider the worst-case scenario for seller j , i.e., $v_j = x - \epsilon$, and $v_{-j} = x + \epsilon$ for other sellers in every period. Given that all sellers adopt the strategy defined in Definition 5, then the expected profit for seller j in equilibrium is $\Pi_j := \frac{1}{1 - \beta} \mathcal{M}(\pi_j) \phi(v_j)$ with $\mathcal{H}(\pi_j) v_j = \mathcal{H}(\pi_k) v_k$ for all $k = 1, \dots, J$ and $\sum_{k=1}^J \pi_k = 1$ in all $t \geq 1$. That is no seller is out of the market.

Here we have $\pi_j(\epsilon) := \pi_j(v_j, v_{-j}) = \pi_j(x - \epsilon, x + \epsilon)$. Deviating by choosing to promise something other than the publicly randomized promised utility value, i.e., promising $v_d \neq v_j^t \in (x - \epsilon, x + \epsilon)$, is not optimal for seller j as long as the degree of dispersion in promised utilities is not too large. That is, if

$$\frac{1}{1 - \beta} \mathcal{M}(\pi_j(\epsilon)) \phi(x - \epsilon) \geq \mathcal{M}(\pi_d) \phi(v_d) + \frac{\beta}{1 - \beta} H\left(\frac{1}{J}\right) \phi(v^*).$$

From Proposition 9, we know that the optimal deviation is $v_d \geq x + \epsilon$. The above inequality holds if

$$\beta \geq \beta_\epsilon := \frac{\mathcal{M}(\pi_d) \phi(v_d) - \mathcal{M}(\pi_j(\epsilon)) \phi(x - \epsilon)}{\mathcal{M}(\pi_d) \phi(v_d) - \mathcal{M}\left(\frac{1}{J}\right) \phi(v^*)}.$$

If $\epsilon > 0$ is small enough, then $\beta_\epsilon < 1$. Finally, punishing following a deviation is optimal because $v_j = v^*$ is a best response to $v_{-j} = v^*$. \square




Figure 1: The lower-bound discount factor $\beta(v_c)$ for $v_c = 0$.

Notes: The figure reports values for markets with $J = 2, 10, 25, 50, 75, 100$ sellers, $I = 2, 20, 100$ buyers, when sellers collude on monopoly prices, $v_c = 0$, given the function $\phi(v) = 1 - v$ for $v \in [0, 1]$.




Figure 2: The static Nash equilibrium promised utility v^* .

Notes: The figure reports values for markets with $J = 2, 10, 25, 50, 75, 100$ sellers and $I = 2, 20, 100$ buyers, given the function $\phi(v) = 1 - v$ for $v \in [0, 1]$.