Search and herding with capacity constraints (Preliminary draft)

Efraim Berkovich and Robert Tayon Jr.
Department of Economics
University of Pennsylvania
October 6, 2008

Abstract

We study value discovery in discrete-time dynamic markets with imperfect information and capacity constraints. Homogeneous, short-lived buyers with inelastic demand for a unit of a heterogeneous good encounter infinite-lived sellers of fixed types. The number of trades at each seller in the prior period is public information. Buyers do not have information about the type of the good offered at any particular seller, but they can acquire private information about a single seller by sampling. This market segments into areas of known quality and unknown quality. The known quality sellers get "too many" buyers until search ends. In segments of unknown quality, informed traders drive out uninformed traders, reducing trade. This pattern of trade resembles risk-aversion, though buyers are risk-neutral.

We thank Randy Wright, Ken Burdett, Philipp Kircher, Jan Eeckhout, and the attendees of the Search and Matching May 2008 workshop at the University of Pennsylvania. Andrew Clausen provided helpful comments.

1 Introduction

A stranger in town looking for a place to buy some lunch looks for restaurants with a fair number of customers already inside. This instinctive behavior seems optimal in the face of meager private information about local establishments and serves as the intuitive basis for the way many markets work. In this paper, we model a simple market where buyers have limited private information but can see the past actions of other buyers.

In discrete time, short-lived buyers face a finite multitude of long-lived sellers offering goods of heterogeneous quality. Buyer payoffs increase in the quality of the good at the seller and decrease in the number of other buyers at the seller. Each buyer has inelastic demand for one unit of the good. When sellers commit to the same quality forever, buyers try to determine the quality of all sellers in order to distribute themselves optimally in this market.

All buyers see the purchasing decisions of the previous period's buyers. Initially, each buyer has no private information about the quality of the sellers, but he can choose to acquire some private information through costly sampling. With this information, buyers make their sampling and purchase decisions in order to maximize their expected payoff.

Although our model is highly stylized, one story it may describe is the market for physicians. With most medical insurance plans, consumers pay the same fee, usually an insurance co-payment, regardless of the choice of doctor. However, physician quality is heterogeneous and somewhat difficult to determine. The quality of service that patients receive depends on the number of other patients scheduled as well as the doctor's skill. More patients at a doctor impose a cost on each patient since scheduling an appointment becomes more difficult and time with the doctor decreases. A public signal is the doctor's popularity, which we can say is the number of patients. Consumers can determine the quality of the doctor, but at a personal cost of time and perhaps health. Nevertheless, seeing a doctor, even of low quality, is generally better than not seeing one at all. We expect that the patient population eventually discovers higher skill doctors, and that, as time progresses, there is an increasing relationship between the number of patients that a doctor has and her quality.

In our model, the number of previous period buyers is a public signal of the quality of a seller. We find that for, a large market, the market segments into regions of publicly known seller quality and regions where the quality is not known exactly but known to be within a particular range. So long as buyers find it worthwhile to sample, sellers of publicly known quality receive a higher number of buyers in that period than the number they receive in the periods when buyers cease sampling. Therefore, those buyers who buy from known quality sellers get a lower payoff compared to the payoff when no one samples. One can intuit that if some sellers are over-bought others must be under-bought. Our work explains this market inefficiency as a being a result of limited information and the resultant feedback trading rather than being a result of irrational traders, noise traders, or even risk averse traders.

There exists a large literature on search with capacity constraints. For example, Peters (1984) presents a directed search model with rationing of buyers by sellers. Burdett, Shi, and Wright (2001) discuss equilibrium pricing and matching in a similar model. Shimer (2005) presents a directed labor search model with heterogeneous workers and firms and finds assortative matching.

Learning from other agents' actions in markets also has extensive precedent in the literature. Banerjee (1992) describes a model of herding where agents obtain a private signal about a choice of options and also observe the prior actions of a number of other agents who faced the same decision. In an elaboration, Debo, Parlour, and Rajan (2008) describe a model where agents obtain a private signal about a firm's quality and also observe the length of a queue of prior randomly arriving agents. The agents infer the firm's quality from the length of the queue.

This topic has been well-studied in relation to asset markets. For instance, Lang and Nakamura (1990) offer a dynamic model of learning in credit markets where a decrease in the number of traders decreases public information and so increases risk thereby further driving out traders. They find that market prices have greater volatility than the underlying shocks because of this feedback. The effect of traders being driven out of market segments with low public information parallels effects seen in our model. Similarly, Caplin and Leahy (1994) describe a market model with private information where the actions of agents publicly reveal information resulting in periodic market crashes.

The importance of the precision of the public versus private information was examined by Morris and Shin (2002) who model a game with strategic complementarities where traders have private and public information. They find that public information may be welfare decreasing if it is not sufficiently precise. Conversely, high precision public information relative to private information can be welfare increasing. In our analysis, the increasing refinement of the public signal improves market efficiency in time.

When rational traders lack information, herding may be part of a best response strategy and may improve efficiency. Nofsinger and Sias (1999) study the effects of institutional and individual stock ownership in relation to herding. Quoting Nofsinger and Sias (1999): "Herding and feedback trading have the potential to explain a number of financial phenomena, such as excess volatility, momentum, and reversals in stock prices."[1] They find a strong positive relation between annual changes in institutional ownership and returns. In their view, this relation is consistent with either intra-year feedback trading by institutions or a stronger impact of institutional herding (vs. individual) on returns. The latter hypothesis resembles work by Postlewaite and Kircher (2008) while the former may be more closely associated with the model in our paper, though neither hypothesis excludes the other.

2 Model

The model described is close to the one we present in Berkovich and Tayon (2008). That paper studied a class of similar models but for non-rival goods. The lack of capacity constraint led to a sharper herding effect than we see with the sharing constraint in this paper. Much of the following description repeats the setup in the referenced paper.

Consider a set $\mathcal{B} = \{1..B\}$ of homogeneous, anonymous, short-lived buyers and a set $\mathcal{S} = \{1..M\}$ of initially homogeneous, long-lived sellers. We assume more buyers than sellers.Sellers produce and sell goods of quality $q \in [0, 1]$. At the start of the game, each seller draws a quality type $q \in [0, 1]$ from a uniform distribution and commits to producing a good of that specific quality forever. Each seller produces one unit of the good each period. We assume that it is difficult to change the production process, and so sellers cannot change the quality of the produced good once the game begins. In this paper, we do not consider strategic sellers and ignore sellers' payoffs.

[1] Quoting further to explain: "Herding is a group of investors trading in the same direction over a period of time; feedback trading involves correlation between herding and lag returns... Most herding models suggest that investors follow some common signal. Feedback trading, a special case of herding, results when lag returns, or variables correlated with lag returns (e.g., earnings momentum, decisions of previous traders, changes in firm characteristics, etc.) act as the common signal."

For the buyer, a higher quality good yields higher utility all else equal. We assume buyers have inelastic demand for one unit of the good from a seller regardless of quality. A single buyer can purchase from only one seller[2]. The buyer's payoff depends on the quality of the good and the number of other buyers at the seller. Following the usual capacity constraint concept in the literature (for example Burdett, Shi, Wright (2001)), we assume a payoff function to the buyer of $u(q,n) = \frac{q}{n}$ where q is the quality of the good and n is the number of buyers at the seller. Although there are a number of justifications for this type of capacity constraint, our preferred interpretation is that it is a proxy for increasing price (decreasing consumer surplus) as demand at a seller increases.

Time proceeds in periods. Each period t, a number of buyers $n_t(i)$ choose to buy from seller i. Sellers vary quality instead of price since in many real-world situations prices are easily visible, whereas quality may be more difficult to determine by the individual buyer. Sellers varying product quality instead of price is somewhat non-standard. However, we think of quality as a stand-in for a consumer surplus which depends on quality and price. Assuming, instead, that sellers have homogenous quality goods and vary the price, we can substitute for price as $p = 1 - q$. If the good is indivisible, then the mathematics is identical. The notion of uncertain quality and known price is similar to Williamson and Wright (1994).

2.1 Public signal and visibility

The information set available to a buyer is critical in determining his actions. One aspect of his information set is what we term *visibility.* We have in mind by *visibility* a characterization of the level of refinement of the information set. A given visibility V, operates on the state of the world θ and yields an object that is the buyer's information set; that is, the visibility operator V determines how the states of the world are partitioned into information sets.

We use V_0 to mean a visibility with no information. This visibility is the one often assumed in standard search literature. All sellers initially look the same to all buyers, that is, buyers see no distinguishing traits for any seller i. We use V_θ to mean "full" visibility of the state of the world where all seller qualities are publicly known. This visibility is the one used in directed search literature (in conjunction with a capacity constraint). Our paper analyzes a "search and herding" visibility, V_N. The V_N visibility means that buyers know the number of buyers who purchased goods at each seller in the previous period. Formally, $V_N(\theta) = \{n_{t-1}(i) : \forall i\}$.

One can consider further refinements to visibility, for example, something between V_N and V_θ. Our interest lies in analyzing the V_N visibility from an initial state of no information (identical to V_0). A different, more refined, initial information state where parts of the sellers space have known quality does not change the basic analysis, nor does it qualitatively affect the end state. If parts of the seller space have known quality at $t = 0$, those sellers receive a known number of buyers and our analysis covers the unknown quality sellers.

2.2 Game

In each period, a buyer's available actions are (a) to sample the quality of the good at a single seller j of their choosing at a cost c, and (b) to choose a seller i from whom to buy. For completeness,

[2] In our model, the binary action of either buying or not buying leads to an over-pricing effect in that searchers cannot "short-sell" a low quality seller. As Nagel (2005) finds, a short-selling restriction leads to an optimism-bias in prices on stocks.

we note that the buyer can choose not to buy from any seller and can choose not to consume the purchased good, but these actions are never optimal.

In the game, time proceeds in discrete periods. The timing of the game is as follows:

1. Sellers are assigned a quality type. They commit to this quality forever.

2. At the start of the period, a population of buyers is born and sees with visibility V_N from whom the previous generation of buyers bought.

3. Buyers optionally sample the quality at one seller at a cost c.

4. Buyers pick a seller from whom to buy, receive their payoff, and die at the end of the period.

6. A new period begins. Repeat from step 2.

Our model differs slightly from many search models in that we relax the requirement that buyers can buy only from a seller they have sampled. This enlarged action space leads to qualitatively different results provided that buyers always buy from a seller (which we have assumed).

Buyers cannot distinguish sellers within a partition set induced by the visibility, that is all sellers with n buyers initially are indistinguishable. We assume a matching technology between buyers and sellers in a partition set is such that buyers are matched with sellers based on a uniform distribution.

3 Buyers' dynamics

The game we would like to consider is of a finite set of buyers of size B and a finite set of sellers of size M. However, for tractability, we study equilibrium of the limit in the number of sellers going to infinity and the ratio of buyers to sellers going to infinity. We suppose that there are $B = N \cdot M$ buyers, so that the average seller receives N buyers. We adjust the payoff function to be $u(q,n) = \frac{q}{n} \cdot N$. Taking the limit as $N \to \infty$, the model approximates the behavior of a continuum of buyers of measure M. When discussing the continuum formulation, we refer to the payoff as $u(q,n) = \frac{q}{n}$ where n is the measure of buyers at the seller.

In each period, buyers decide whether to sample or not. We consider only symmetric equilibria and denote by σ the fraction of buyers who sample[3]. If sampling, the buyer selects a seller $i(V_N(\theta))$ to sample. If not buying from him, the buyer selects another seller $j(q(i), V_N(\theta))$ from whom to buy. After sampling, buyers use a decision function $f(q(i), V_N(\theta))$ to (possibly) mix over buying from the realized i versus buying from a random j. If not sampling, the buyer selects a seller $j(V_N(\theta))$ from whom to buy.

The expected payoff to sampling is

$$W_S(t) = E\left[u(q(i), n_t(i)) f(q(i), V_N(\theta)) + (1 - f(q(i), V_N(\theta))) u(q(j), n_t(j))\right]$$

and the expected payoff to not sampling is

$$W_N(t) = E\left[u(q(j), n_t(j))\right]$$

If buyers had full knowledge of seller quality (via, for instance, the V_θ visibility), then each buyer receives a payoff of $\frac{1}{2}$ almost surely because of the assumption of the uniform quality distribution over $[0,1]$.

In period $t = 0$, buyers have no information about the quality of any seller, so they sample uniformly across all sellers. The decision function, $f(q)$ of whether to buy from the sampled seller

[3]When dealing with a finite number of buyers, σ represents the probability of a buyer sampling.

depends only on the quality of the seller in period $t = 0$. Define by $L(q)$, the expected payoff to leaving the sampled seller and buying from a random other seller. Define by $S(q)$, the expected payoff to buying from the sampled seller.

We use the following definition to describe the shape of the decision function $f(q)$.

Definition 1 *"Flats" are positive measure regions $f^{-1}(0)$ and $f^{-1}(1)$.*

In the first period of the game, if buyers sample, the decision function is increasing and has the following shape: flat-gradient-flat.

Proposition 2 *In $t = 0$, if $\sigma > 0$,*
(a) $f(q)$ is increasing and strictly increasing when $f(q) \in (0,1)$,
(b) a flat exists $\Leftrightarrow W_S > W_N$, and
(c) two flats exist, $f^{-1}(0) = [0, q_{lo}]$ and $f^{-1}(1) = [q_{hi}, 1]$, $0 < q_{lo} < q_{hi} < 1$.

Proof. Part (a): Non-samplers (fraction $1 - \sigma$ of the buyers) and "unhappy" samplers (that is, samplers who do not buy from their sampled seller) buy with a uniform distribution across all sellers. We look at three cases:

Case 1: If $f(q) = 1$, then $f(q') = 1 \; \forall q' > q$.

Case 2: If $f(q) = 0$, then $f(q') = 0 \; \forall q' < q$.

Case 3: Otherwise, $f(q)$ is strictly increasing for q such that $f(q) \in (0,1)$.

Proof for cases 1 and 2 is straightforward and omitted. For case 3, the payoffs to staying and leaving must be the same. Suppose, toward contradiction, that $f(q) \geq f(q')$ for some $q < q'$.All else being equal, the expected payoff to buyers who leave depends on f. It must be that $L(q) \geq L(q')$ as more buyers leaving the sampled seller lowers payoffs at the other sellers. We have that $L(q) = S(q)$ and $L(q') = S(q')$ since buyers are mixing between leaving and staying. So $S(q) \geq S(q')$, but this is impossible since $u(q,n)$ strictly increases in q and strictly decreases in n and the expected number of buyers is lower at sellers of quality q'.

Part (b): Suppose no flat exists, then all samplers get the same expected payoff because $L(q) = S(q)$ everywhere so sampling and then purchasing from a random seller gives the same expected payoff as purchasing from a random seller—the expected payoff of a non-sampler. Conversely, payoffs on a flat are increasing linearly in q; so samplers get a surplus. Note, since $W_S - c \geq W_N$ in equilibrium, then $c > 0$ implies flats.

Part (c): In the finite-sized game, the payoff to staying at a very low quality seller is decreasing to zero in the quality (regardless of the number of buyers). Therefore, there must be some positive measure region of low quality where samplers leave for sure. Define by $D(q)$ the expected payoff to buying uniformly from a random seller (including the sampled seller) after sampling $q(i) = q$. We have that $D(q) = \frac{1}{M} S(q) + \frac{M-1}{M} L(q)$. Because buyers sample all sellers uniformly, $\int D(q)\,dq = \int L(q)\,dq$. We have that $D(q) < L(q)$ for q in the low flat. Furthermore, $D(q) = L(q)$ for q in the gradient Therefore, it must be that $D(q) > L(q)$ in the high flat and since the low flat is positive measure, the high flat must be positive measure.

We turn to the question of whether both flats exist in the limit of sellers going to infinity and buyer to seller ratio N going to infinity. It is sufficient to consider only when all buyers sample.

Two flats do not exist when taking the limit in N going to infinity but keep the number of sellers fixed at some $M < \infty$. Suppose that there exists a positive measure high flat and that there exists a low flat. A buyer on the low flat knows there is a positive probability that all other sellers are high flat and thus a positive probability that he would be the only buyer at his low flat seller if he stays. Since the payoff to being the sole buyer goes to infinity, he stays. This effect is due to the possibility of unbounded payoffs.

Two flats exist when we fix the ratio of buyers to sellers at $N < \infty$ and take the number of sellers to infinity. At any seller, the arrival rate of "unhappy" samplers approaches a Poisson distribution with some rate $\lambda(N) > 0$. Therefore, the expected payoff at a low flat seller is

$$E\left[u(q,n)\,|q \in \text{ low flat}\right] = qNe^{-\lambda}\sum_{k=0}^{\infty}\frac{\lambda^k}{k!}\frac{1}{k+1}$$

Since the summation converges to a finite number (less than e^{λ}), then expected payoff is decreasing in q to zero and there exists a low flat.

With the number of sellers at infinity, two flats exist in the limit of $N \to \infty$. It must be that $0 <; \lim_{N\to\infty}\frac{N}{\lambda} < K$, for some $0 < K < \infty$. Clearly, the ratio cannot go to zero. If the ratio goes to infinity, that implies nearly all buyers stay with their sampled seller, which is not individually rational. So,

$$E\left[u(q,n)\,|q \in \text{ low flat}\right] \leq \lim_{\lambda\to\infty} qK\lambda e^{-\lambda}\sum_{k=0}^{\infty}\frac{\lambda^k}{k!}\frac{1}{k+1}$$

But $\lim_{\lambda\to\infty} e^{-\lambda}\sum_{k=0}^{\infty}\frac{\lambda^{k+1}}{(k+1)!} = 1$, so there exists a low flat. A high flat then necessarily exists because expected payoffs are bounded. ■

Having a finite number of buyers implies that to solve for equilibrium we must calculate the payoffs to a variety of possible (and low probability) distributions of buyers over sellers. But when taking the limit, the expectation of the payoff simplifies..We therefore proceed with the analysis as if we have a continuum of buyers of measure $B \equiv M$, that is, there are measure one buyers for every seller and then take M to infinity. The proposition above suggests that since the buyers' decision function is well-behaved in the limit, our analysis converges to the dynamics of the model taking the limit in M and N. The existence of possible unbounded payoffs in the continuum formulation does not occur in our equilibrium limit formulation, so we ignore them.

Assuming a finite M number of sellers and a continuum of measure M buyers, in $t = 0$,the payoff to leaving a seller of quality $\tilde{q}$ after sampling is

$$L(\tilde{q}) = \int_0^1 \frac{q}{1-\sigma+\sigma f(q)+\gamma(\tilde{q})}dq$$

where $\gamma(\tilde{q})$ is the expected measure of "unhappy" samplers who arrive to buy from every other seller other than the $\tilde{q}$ seller. It is

$$\gamma(\tilde{q}) = \frac{\sigma}{M-1}\left[(1-f(\tilde{q})) + (M-2)\int_0^1 (1-f(q))\,dq\right]$$

However, as $M \to \infty$, $\gamma(\tilde{q}) \to \gamma$ where

$$\gamma = \sigma\int_0^1 (1-f(q))\,dq$$

The payoff to buying from the sampled seller is

$$S(\tilde{q}) = \frac{\tilde{q}}{1-\sigma+\gamma+\sigma f(\tilde{q})}$$

where γ is the measure of "unhappy" samplers who arrive to buy from this seller as before.

For buyers to mix between leaving and staying, it must be that $S(q) = L(q)$ for some q. In the limit of M, the expected payoff to leaving, $L(\tilde{q})$, goes to the expected payoff to buying randomly from a seller, that is, $L(\tilde{q}) \to W_N$, the payoff to non-samplers. Looking at the ends of the gradient, that is q_{lo} and q_{hi}, a simplified form for W_N arises. Since in the limit $S(q_{lo}) = S(q_{hi}) = W_N$, then $W_N = \frac{q_{lo}}{1-\sigma+\gamma} = \frac{q_{hi}}{1+\gamma}$.

The value of the sample or "search again" option can be described in two ways. The option can be described as the right to reject the sampled seller, a "put" option, or it can be described as the right to buy from the sampled seller, a "call" option. Implicit in both definitions is that the outside option of buying from a random seller exists.

Proposition 3 *In the limit of M and N, the value of a put option is c for $\sigma \in (0,1)$ and is greater than c for $\sigma = 1$. Furthermore, the values of a put and call option are equal.*

Proof. For $\sigma = 1$, define W_N as the expected payoff to buying from a random seller. We have that $W_S - c \geq W_N$ with equality for $\sigma \in (0,1)$. Given the shape of $f(q)$,

$$\begin{aligned} W_N &= \int_0^{q_{lo}} \frac{q}{1-\sigma+\gamma} dq + \int_{q_{lo}}^{q_{hi}} W_N dq + \int_{q_{hi}}^{1} \frac{q}{1+\gamma} dq \\ W_S &= \int_0^{q_{hi}} W_N dq + \int_{q_{hi}}^{1} \frac{q}{1+\gamma} dq \end{aligned}$$

Subtracting, we get that $c \leq \int_0^{q_{lo}} \frac{q_{lo}-q}{1-\sigma+\gamma} dq$ where the r.h.s. is the value of the put option. Similarly, $c \leq \int_{q_{hi}}^{1} \frac{q-q_{hi}}{1+\gamma} dq$ where the r.h.s. is the value of the call option. ■

The $\frac{q}{n}$ capacity constraint is zero-sum, meaning total consumer surplus is constant regardless of the distribution of the buyers, implying the equation

$$\sigma W_S + (1-\sigma) W_N = \frac{1}{2} \tag{1}$$

For interior σ, the indifference $W_S - c = W_N$ implies that the value of the search option (for example the put option) can be written as

$$c = \frac{1}{2} \cdot \frac{q_{lo}}{1-\sigma+\gamma} \tag{2}$$

To derive the decision function $f(q)$, we note that the buyers are indifferent between staying and leaving on the gradient, so $W_N = \frac{q(i)}{1-\sigma+\gamma+\sigma f(q(i))}$. Differentiating with respect to $q(i)$ gives $f'(q) = \frac{1-\sigma+\gamma+\sigma f(q)}{\sigma q}$ and with the condition $f(q_{lo}) = 0$, we get

$$f(q) = \frac{\sigma - 1 - \gamma}{\sigma} + q \frac{1-\sigma+\gamma}{\sigma q_{low}}$$

So, the decision function f is linear on the gradient and so $\gamma = \sigma \int_0^1 (1 - f(q)) dq = \sigma \left(q_{lo} + \frac{q_{hi}-q_{lo}}{2} \right)$. Since $W_N = \frac{q_{lo}}{1-\sigma+\gamma} = \frac{q_{hi}}{1+\gamma}$, one can write q_{hi} in terms of the only σ and q_{lo}. Substituting for W_N in (1) yields $\sigma c + \frac{q_{lo}}{1-\sigma+\gamma} = \frac{1}{2}$. Solving with (2) gives a unique equilibrium for interior σ, graphed in Figure 1 . When the value of the search option is strictly greater than c, then $\sigma = 1$, and the equation 1 defines q_{lo} uniquely.

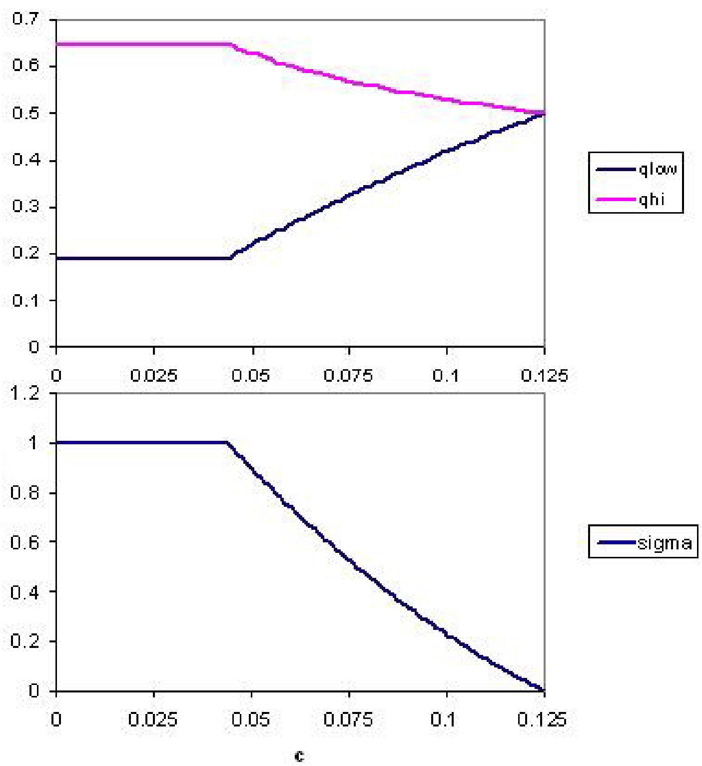

Figure 1: At $t = 0$, equilibrium values for σ, q_{low}, and q_{hi} given c.

The fraction of samplers, σ, is determined by c. When $\sigma = 0$, no one samples and the marginal buyer's reservation quality is $\frac{1}{2}$. The maximum search option value is $\frac{1}{8}$. For $c \geq \frac{1}{8}$, $\sigma = 0$. As c decreases from $\frac{1}{8}$ and σ increases, q_{lo} and q_{hi} diverge from $\frac{1}{2}$. For low enough c, all buyers sample. There exists a $\bar{c} \approx 0.043689$ such that $\sigma = 1\ \forall c < \bar{c}$. This result differs from the endogenous directed search model in Lester (2007) In that paper, the endogenously determined measure of searchers approached one only if search was costless whereas, in our model, all buyers choose to search even for low but positive costs to searching.

At $t = 1$, buyers see three regions: low flat, gradient, and high flat. Buyers know the quality of any seller on the gradient because $f(q)$ is strictly increasing. Therefore, searchers only sample sellers in the flats. Searchers use regional decision functions f_{low} and f_{high}, that have properties as in $t = 0$.

Consider a parameterization of $t = 0$, the flat i is $[a_i, b_i]$ and has μ_i fraction of the buyers. It suffices to consider equilibrium strategies where the problem is "regionalized." A multiplicity of equilibrium strategies exist where "unhappy" and non-sampling buyers buy from different regions in varying measures. Since we are interested in equilibrium outcomes, we do not distinguish between these strategies. We use the notation σ_i is the fraction of the buyers in the region i who sample and γ_{ii} is the fraction who are "unhappy" samplers, so, for instance, the measure of buyers who search in region i is $\sigma_i \mu_i$ Rewriting the equation (1),

$$\sigma_i c + \frac{q_{lo} - a_i}{\frac{\mu_i}{b_i - a_i}(1 - \sigma_i + \gamma_i)} = \frac{\frac{a_i + b_i}{2}}{\frac{\mu_i}{b_i - a_i}}$$

and substituting $k_i = \frac{\mu_i}{b_i - a_i} c$ yields,

$$\sigma_i k_i + \frac{q_{lo} - a_i}{1 - \sigma_i + \gamma_i} = \frac{a_i + b_i}{2} \tag{3}$$

For interior σ_i, the value of the search option equation (2), similarly yields

$$k_i = \frac{1}{2} \frac{(q_{lo} - a_i)^2}{1 - \sigma_i + \gamma_i} \tag{4}$$

Additionally, these market segments are linked, so $W_N = W_{N,i}\ \forall i$ and

$$W_{N,i} = \frac{q_{lo,i} - a_i}{\frac{\mu_i}{b_i - a_i}(1 - \sigma_i + \gamma_i)} = \frac{b_j^2 - a_j^2}{2\mu_j} = W_{N,j} \tag{5}$$

for flat region i and gradient region j. Since the total measure of buyers is fixed, $\sum \mu_k = 1$.

Sampling continues in subsequent periods until such time as the value of the search option is less than c. The market continues to segment into gradient and flat regions as each former flat region becomes transformed into a flat-gradient-flat. Figure 2 provides an illustration of buyer distribution in periods $t = 0$ and $t = 1$.

By our analysis of $t = 0$, we have that the regionalized solution is unique given the endpoints a and b, the search cost c, and the total measure of buyers in the segment μ. Since $W_{N,k}$ is the same for all regions k in equilibrium, the solution requires finding μ_k for all regions.

Proposition 4 *In the limit of M and N, there is a unique equilibrium outcome for $c = 0$ and a finite number of equilibrium outcomes for $c > 0$.*

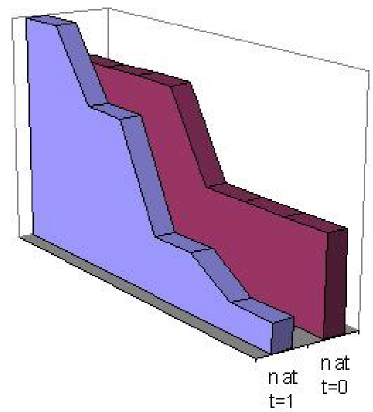

Figure 2: Illustrative graphic of buyer distribution in $t = 0$ and $t = 1$.

Proof. Consider period $t-1$ where there are multiple regions indexed by k. We know that each region has a unique equilibrium outcome given the measure of buyers there, μ_k. The expected payoff to a buyer in region k is defined by $W_k = \max\{W_{N,k}(\mu_k; c, a_k, b_k), W_{S,k}(\mu_k; c, a_k, b_k) - c\}$.

We show that, for two flats i and j, there are a finite number of crossings for $W_i(\mu)$ and $W_j(d-\mu)$ given μ, where d is the total mass to be split between them. Since W_k strictly decreases when $\sigma = 0$ or 1, we only need to check that $W_i(\mu) = W_j(d-\mu)$ for a finite number of points when σ is interior. Since $W_i(\mu)$ is smooth when σ is interior and has a limit as $\sigma = 0$ and 1, then if there are an infinite number of intersections it must be that $W_i(\mu) = W_j(d-\mu)$ over the entire region $\sigma \in (0,1)$ Define $f(\mu; c, a_i, b_i) = \frac{q_{lo}-a_i}{1-\sigma_i+\gamma_i}$, so $W_i(\mu) = f(\mu)\frac{b_i-a_i}{\mu}$, so that means

$$\frac{f(\mu; c, a_i, b_i)}{f(d-\mu; c, a_j, b_j)} = \frac{\mu}{d-\mu}\frac{b_j-a_j}{b_i-a_i}$$

We have that $f(0)$ is zero only if $q_{lo} = a$, but we calculate that there is no real solution for σ when $q_{l0} = a$. Therefore, these two smooth functions are coincident over a finite number of intersecting points. There is at least one intersection since $W_i \to \infty$ as $\mu \to 0$.

Where g is a gradient region, the payoff $W_g = \frac{b_g^2-a_g^2}{2\mu_g}$ strictly decreases in μ_g. Therefore, in equilibrium, the payoff in all gradients strictly decreases in μ_g (any measure added to any gradient). Since $W_g(\mu)$ and $W_i(1-\mu)$, where i is a flat region, are both (piece-wise) smooth functions, they can only be coincident if they have the same derivative. But they do not, so $W_g(\mu) = W_i(1-\mu)$ in only a finite number of points (and again there is at least one intersection). Therefore, since any two regions have a finite number of intersections, the whole market has a finite number of equilibrium outcomes in any period.

Since for $c = 0$, $\sigma = 1$ and W_g is strictly decreasing always, there is a unique equilibrium oucome in every period and hence a unique equilibrium outcome for the whole game. For $c > 0$, (we show later) there is a period $T_c < \infty$ such that there is no sampling for $t \geq T_c$ so there are a finite number of equilibrium outcomes for the whole game. ■

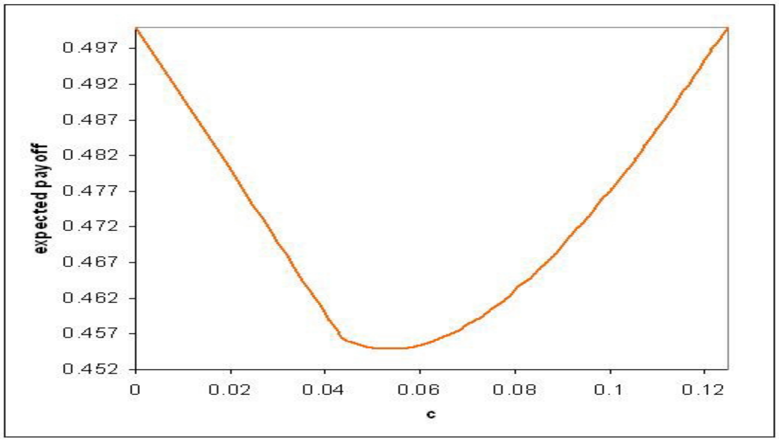

Figure 3: Total welfare (or expected payoff) in $t = 0$ as a function of c.

In the formulation with a continuum of buyers, buyers know the entire history of the market from the public signal of the current period in the sense that they know how many buyers each seller had in every prior period. Since each gradient seller's quality is known and the equilibrium strategy of past periods is known, buyers can determine which flat (or gradient) that seller was in during every prior period. Current flats sellers are known to be within a quality range which must be within a flat in every prior period. The public signal serves as a perfect substitute for history much as the state variable of money distribution does in Corbae, Temzelides, and Wright (2003) for the directed matching economies (but not the random matching economies) described in that paper. When N is finite, however, history cannot be perfectly known, rather buyers have a probability distribution over possible histories.

4 Welfare

Because of the payoff function used, the total amount of consumer surplus is fixed at $\frac{1}{2}$. Social welfare is total surplus minus total search costs. In $t = 0$, social welfare is then $[\sigma W_S + (1-\sigma) W_N] - \sigma c = \frac{1}{2} - \sigma c$ For interior σ, $W_S = W_N + c$ and so social welfare is W_N. W_N is non-monotonic in c. When adding on the solutions where $\sigma = 1$, this equation for social welfare is, in fact, the expected payoff to buyers as a function of c.

For interior equilibria, all buyers receive payoff of W_N. Given the risk-neutrality of buyers, the centralized assignment problem, or the social planner's problem, is solved by maximizing total welfare, assuming transferable utility. From the standpoint of total welfare, the first-best outcome is when no one pays to sample (or sampling is free). The second-best outcome, the equilibrium we describe, is therefore welfare reducing.

Therefore, rather than looking at allocative efficiency, we believe it makes more sense to consider informational efficiency. We mean by informational efficiency that the public signal about a seller shows that seller's quality (and hence the payoff to buyers from buying there in equilibrium). An analogous formulation is that informational efficiency is achieved when the variance in buyer payoffs is zero. In this decentralized market, if the market becomes informationally efficient, then it is not individually rational to pay to sample, and so the equilibrium and social planner's outcomes coincide.

So long as there are flats, the market is not informationally efficient. Each period when buyers search brings the market closer to informational efficiency.

5 Finite variations

Consider the model with finite buyers and finite sellers. If buyers search in $t = 0$, they search in an infinite number of periods. This observation results from the fact that a finite number of buyers cannot faithfully reproduce any distribution with probability one. Even if buyers achieve informational efficiency, in the following periods buyers cannot replicate the configuration. Without further search, the public signal becomes completely uninformative as $t \to \infty$.

When N, the ratio of buyers to sellers, goes to infinity, the public signal of the number of buyers at a seller becomes increasingly less noisy for next period buyers who try to invert from the signal to the quality of the seller. At the limit of N, even with finite sellers, the quality of gradient seller can be known from the public signal in $t = 1$. We note that, there exists only one flat in this setup (as explained in a proposition above), but that the decision function is still increasing otherwise.

Proposition 5 *For finite M sellers and N going to infinity, in $t = 1$, the signal $n_t(i)$ is invertible to $q(k)$ for $q(k)$ in the gradient, publicly if a flat exists and privately by all samplers otherwise.*

Proof. Only one flat exists in the limit of N for finite M, but the decision function is still increasing so there is a gradient. It is sufficient to consider the case of $\sigma = 1$, since buyers can subtract off the measure of non-samplers from every seller.

In $t = 1$, if a flat is publicly visible (that is, two or more sellers have the same number of buyers), then buyers know that $f(q(i)) = 1$ (say for the high flat) for all i in the flat. If no flat is publicly visible, then at most one seller is in the flat. In that case, buyers sample from an extreme seller (analogous to searching the flats). In that case, the quality of this extreme seller i (highest or lowest depending on the shape of f) is known to the samplers and so is $f(q(i))$.

Suppose seller i is this seller so that $f(q(i))$ is known. The total number of buyers is $n_0(i) = f(q(i)) + \frac{1}{M-1}\sum_{j\neq i}(1 - f(q(j)))$. Pick any seller k on the gradient. Then,

$$\begin{aligned} n_0(i) - n_0(k) &= f(q(i)) + \frac{1}{M-1}(1 - f(q(k))) \\ &\quad - f(q(k)) - \frac{1}{M-1}(1 - f(q(k))) \end{aligned}$$

Since $n_0(i)$ and $n_0(k)$ are known publicly and $f(q(i))$ is known, this equation is solved for $f(q(k))$. And since $f(q)$ is strictly increasing on the gradient, it is solved for $q(k)$. ■

With finite sellers, it may be the case that all sellers are in the gradient in $t = 0$. If that event occurs, in $t = 1$ samplers know the quality of every seller by the result above and distribute themselves so that every sampler gets the same payoff. In future periods, the public signal becomes less informative because multiple histories can yield the same buyer distribution. For example, a state of the world where sellers are in a flat in $t = 0$ and are in the gradient in $t = 1$, yielding a market distribution of buyers which can look identical to the distribution resulting from the state of all sellers being in the gradient in $t = 0$. Therefore, after $t = 1$, an informationally efficient distribution is not possible for finite sellers.

6 Excess trade at the gradient and end of search

Consider $t = 1$, there are three sub-markets: the flats and the informationally efficient gradient. Search in the flats reduces payoff to non-searchers in the flats causing an increased number of non-searchers on the gradient. So, relative to the informationally efficient amount of buyers, the gradient has too many buyers and the flats have too few buyers. The presence of search reduces payoffs to uninformed buyers and reduces trade when there are linked markets.

Proposition 6 *If $\sigma > 0$, then, after search, all flats become smaller by a fraction bounded away from one. All flats go to length zero if search continues forever.*

Proof. Consider a distribution at some period t that has flats. Given a positive measure flat $[a, b]$ that is searched by a positive measure of buyers who all sample, the following period distribution is flat-gradient-flat. The gradient is informationally efficient, in the sense that seller qualities are publicly known. The flats have sellers of quality that are within a known range. It suffices to show that the gradient in $[a, b]$ is always of measure greater than $\eta (b - a)$ for some $\eta > 0$. To simplify the presentation, we normalize the average measure of sampling buyers at a seller to one. First, note that the shape of the decision function cannot approach all flat. If the decision function goes to all high flat, then a buyer who samples a seller with quality $\frac{a+b}{4}$, and hence an expected payoff to staying of $\frac{a+b}{4}$, can, by leaving, get a higher expected payoff that approaches $\frac{a+b}{2}$. (Similarly, if the decision function goes to all low flat a buyer who samples a seller with quality $\frac{3(a+b)}{4}$ is better off staying). Therefore, there is a positive measure γ bounded away from zero of unhappy samplers. Consider the expected payoffs at the low and high ends of the gradient; they must be equal so $\frac{q_{lo}}{\gamma} = \frac{q_{hi}}{1+\gamma}$. Since γ is bounded away from zero, it cannot be that $q_{lo} \to q_{hi}$ and the gradient must be a positive fraction bounded away from zero of the flat $[a, b]$. ∎

When c is greater than the search option value, search ends—that is, buyers stop sampling. As a consequence to the above proposition, the search option goes to zero if search continues forever. So for any $c > 0$, $\exists T_c$ such that there is no search for all $t \geq T_c$. When search ends, buyers know the quality of each seller in the gradient and they know the average quality in the flats. Since the gradient is informationally efficient, buyers buy from gradient sellers in such measure so that a buyer's payoff at each gradient seller is $\frac{1}{2}$ Each seller in a flat region receives the same number of buyers, with the expected payoff being $\frac{1}{2}$ for a buyer in a flat. If $c = 0$, then search never ends, and the whole market approaches informational efficiency in the limit of time.

Proposition 7 *When sampling is costless and buyers sample forever, for each seller i, $\lim_{t\to\infty} n_t(i) = 2q(i)$ and the payoff to every buyer approaches $\frac{1}{2}$ as $t \to \infty$.*

Proof. By the proposition above, we have that all flats converge to zero length. Given any $\varepsilon > 0$, all sellers qualities are publicly known within ε at time $t > T_\varepsilon$ for some finite T_e. Since all buyers sample, all buyers have an expected payoff of $\frac{1}{2}$ (by the zero-sum condition on consumer surplus). So, for a flat $[\hat{q} - \delta, \hat{q} + \delta]$ at time $t - 1$, the average measure of buyers that buy from a seller in that flat at time t must be $2\hat{q}$. As $\delta \to 0$, $\lim_{t\to\infty} n_t(i) = 2q(i)$ for buyers on the flats. Since all buyers must get the same expected payoff, it must be that $\lim_{t\to\infty} n_t(i) = 2q(i)$ for buyers on the gradient as well. The payoff to every buyer approaches $\frac{1}{2}$. ∎

When search goes forever, the limiting distribution of measure $2q(i)$ buyers at seller i is informationally efficient.

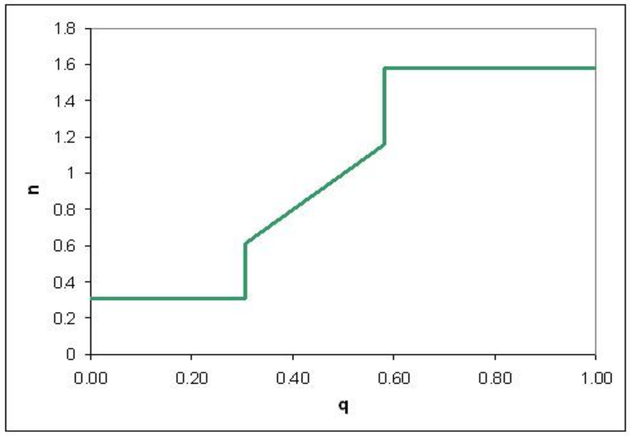

Figure 4: Buyer distribution for $t \geq 1$, for $c = 0.07$ (when search ends at $t = 1$).

7 Conclusion

We describe the dynamic evolution of a very simplified market with public and private information. In our model, buyers see the purchasing actions of the previous period's buyers leading to social learning through time. In Berkovich and Tayon (2008) with a similar model but with non-rival goods, most sellers are repeatedly abandoned and public information about them is lost. However, due to the capacity constraint imposed in the model presented in this paper, buyers spread out across all sellers in every period. The dynamics generate more refined information about the quality of every seller.

Public information about each seller does not grow uniformly for all types of seller. Buyer activity segments the market into regions of known value and regions of unknown, though range bound, value. Buyers pay a cost in order to get private information about the unknown regions. The existence of these informed buyers inhibits uninformed buyers from buying in the market segments of unknown value and drives them to buy from segments with known quality sellers. This excess trade with known quality sellers continues until the cost of searching exceeds the expected return, at which point search ends. Although all traders are risk-neutral, the resulting market dynamics create the impression that uninformed traders are risk-averse because they engage in excess trade with known value sellers.

Thinking of this model as an asset market with prices being related to the number of buyers, the initial period of no asset specific information may correspond to a shock to an existing market so that prices of individual assets vary from the informationally efficient value. As time proceeds and the public signal contains more information, prices come closer to the value of the underlying asset. Initially, known value assets are over-bought because uninformed traders have nowhere else to go, having been forced out of unknown value market segments by lower returns. Meanwhile, informed traders get higher returns by sifting through the market segments of unknown value assets, thus lowering returns to uninformed traders. This trading pattern looks like a "flight to safety" which eventually dissipates as the public information about the market becomes more refined. In time, prices converge to the underlying value as traders learn about asset values by observing past trades.

This model offers caution about comparing returns in new asset markets. The return achieved by the average investor in a particular market may be higher, in general, than the return to the average asset in the market because of the higher returns achieved by informed investors. This error of comparison may explain effects such as described by Dimson and Marsh (1999) who find that stock market returns for small UK firms had a premium return of 6% until the effect was publicized after which point they had a discount return of 6%. Securities based on a broad index of a market segment may underperform relative to the return of prior investors in that market segment. One may speculate that the error was also committed recently by investors in securitized mortgage bonds who had the expectation of comparable returns to prior mortgage bond investors. The investors in these securities presumably aimed for the average investor return instead of the average asset return they in fact earned.

8 References

Banerjee, Abhijit. "A simple model of herd behavior," *Quarterly Journal of Economics,* (1992) v.107(3), pp.797-817.

Berkovich, Efraim, and Robert Tayon, Jr. "Search, herding, and limits to efficiency" *manuscript* (2008) .

Burdett, Kenneth, Shouyong Shi, and Randall Wright. "Pricing and Matching with Frictions," *Journal of Political Economy* (2001), v.109, pp. 1060–1085 .

Caplin, Andrew and John Leahy. "Business as Usual, Market Crashes, and Wisdom After the Fact," *American Economic Review* (1994) v.84(3), pp. 548-565

Corbae, Dean, Ted Temzelides, and Randall Wright. "Directed Matching and Monetary Exchange," *Econometrica* (2003) v. 71(3), pp. 731-756

Debo, Laurens. G., Christine A. Parlour, and Uday Rajan. "Inferring Quality from a Queue," *manuscript* (2008)

Dimson, Elroy and Paul Marsh. "Murphy's Law and Market Anomalies," *Journal of Portfolio Management* (1999) v.23 pp.53-69

Kircher, Philipp and Andrew Postlewaite, "Strategic firms and endogenous consumer emulation," *Quarterly Journal of Economics*, (2008) v. 123(2), pp. 621-661

Lang, William W. and Leonard I. Nakamura. "The dynamics of credit markets in a model with learning," *Journal of Monetary Economics,* (1990) v. 26(2), pp. 305-318

Lester, Benjamin, "Differential information, limited availability, and prices" *manuscript* (2007).

Morris, Stephen and Hyun Song Shin. "The Social Value of Public Information," *American Economic Review* (2002), v.92, pp.1521-1534.

Nagel, Stefan. "Short sales, institutional investors and the cross-section of stock returns," *Journal of Financial Economics* (2005) v. 78(2), pp. 277-309

Nofsinger, John R. and Richard W. Sias. "Herding and Feedback Trading by Institutional and Individual Investors," The Journal of Finance (1999) v.54(6) pp. 2263-2295

Peters, Michael. "Bertrand Equilibrium with Capacity Constraints and Restricted Mobility," *Econometrica* (1984), v.52,(5) pp. 1117-1127

Shimer, Robert. "The Assignment of Workers to Jobs in an Economy with Coordination Frictions," *Journal of Political Economy*, (2005) v. 113(5), pp. 996-1025

Williamson, Steve and Randall Wright. "Barter and Monetary Exchange Under Private Information," *American Economic Review* (1994) v. 84(1), pp. 104-123