

Shopbot Economics

Jeffrey O. Kephart and Amy R. Greenwald

IBM Institute for Advanced Commerce
IBM Thomas J. Watson Research Center
Yorktown Heights, New York, USA
kephart@watson.ibm.com, amygreen@cs.nyu.edu

Abstract. Shopbots are agents that search the Internet for information pertaining to the price and quality of goods or services. With the advent of shopbots, a dramatic reduction in search costs is imminent, which promises (or threatens) to radically alter market behavior. This research includes the proposal and theoretical analysis of a simple economic model which is intended to capture some of the essence of shopbots, and attempts to shed light on their potential impact on markets. Moreover, experimental simulations of an economy of software agents are described, which are designed to model the dynamic interaction of electronic buyers, sellers, and shopbots. This study forms part of a larger research program that aims to provide new insights on the impact of agent and information technology on the nascent information economy.

1 Introduction

Shopbots, agents that automatically search the Internet for goods and services on behalf of consumers, herald a future in which autonomous agents become an essential component of nearly every facet of electronic commerce [3, 8, 12, 5]. In response to a consumer's expressed interest in a specified item, a typical shopbot can query several dozen web sites, and then collate and sort the available information for the user — all within seconds. For example, `www.shopper.com` claims to compare 1,000,000 prices on 100,000 computer-oriented products! In addition, `www.acses.com` compares the prices and expected delivery times of books offered for sale on-line, while `www.jango.com` and `webmarket.jungle.com` offer everything from apparel to gourmet groceries. Shopbots can out-perform and out-inform even the most patient, determined consumers, for whom it would take hours to obtain far less coverage of available goods and services.

Shopbots deliver on one of the great promises of electronic commerce and the Internet: a radical reduction in the cost of obtaining and distributing information. It is generally recognized that freer flow of information will profoundly affect market efficiency, as economic friction will be reduced significantly [1, 6, 9, 4]. Transportation costs, menu costs — the costs to firms of evaluating, updating, and advertising prices — and search costs — the costs to consumers of seeking out optimal price and quality — will all decrease, as a consequence of the digital nature of information as well as the presence of autonomous agents that find,

process, collate, and disseminate that information at little cost. What are the implications of the widespread use of shopbots? Specifically, do shopbots have the potential to increase social welfare? If so, how can shopbots adequately price their services so as to provide consumers with incentives to subscribe, while retaining profitability? More generally, what is the expected impact of agent technology on the nascent information economy?

Previous work in economics on the impact of search costs on equilibrium prices was oriented towards explaining the phenomenon of price dispersion in social economies; see, for example, [11, 13, 2]. In such work, an attempt is made to approximate human behavior with mathematical functions or algorithms, and under the relevant assumptions, collective behavior and equilibria are studied. In contrast with previous intentions, our mission is to investigate the possible dynamics of the future information economy in which software agents, rather than human constituents, will play the key role. Consequently, we take mathematical functions and algorithms a good deal more seriously, by regarding them as precise specifications of the behavior of economic players. In this paper, we focus on the likely effect that one particular specification of a class of agents, namely shopbots, will have on electronic markets. From this study, we hope to gain insights into the design of adaptive algorithms for economically-motivated, computational agents which successfully maximize utility.

This paper is organized as follows. The next section presents our model of a simple market in which shopbots provide price information, which is analyzed from a game-theoretic point of view in Section 3. In Section 4, we consider the dynamics of interaction among software agents designed to model electronic consumers and producers; moreover, we investigate the effect of non-linear search costs (Section 4.1) and irrational consumers (Section 4.2) via experimental simulations. Finally, Section 5 presents our conclusions and ideas for future work.

2 Model

We consider an economy in which there is a single commodity that is offered for sale by S sellers and of interest to B buyers. Periodically, at a rate ρ_b , a buyer b attempts to purchase a unit of the commodity. Each attempted purchase proceeds as follows. First, buyer b conducts a search of fixed sample size i , which entails requesting $0 \leq i \leq S$ price quotes.¹ A search mechanism (which could be manual or shopbot-assisted) instantly provides price quotes for i randomly chosen sellers. Buyer b then selects a seller s whose quoted price p_s is lowest among the i (ties are broken randomly), and purchases the commodity from seller s if and only if $p_s \leq v_b$, where v_b is buyer b 's valuation of the commodity.

In addition to the purchase price, buyers may incur search costs. The cost c_i of using search strategy i , however, does not enter into the purchasing decision of the buyers, because buyers must commit to conducting a search before the results of that search become available. In other words, search payments are

¹ We permit a search strategy of 0 to allow buyers to opt out of the market entirely, which may be desirable if search costs are prohibitive.

sunk costs. Instead, search costs affect the choice ($0 \leq i \leq S$) of search strategy utilized by buyers. A buyer b is assumed to periodically re-evaluate its strategy at a rate $\sigma_b \leq \rho_b$, where typically, $\sigma_b \ll \rho_b$. Upon re-evaluation, the *rational* buyer estimates a price \hat{p}_i that it expects to pay for the commodity if it uses strategy i , and selects the strategy j that minimizes $\hat{p}_j + c_j$, provided that $\hat{p}_j + c_j \leq v_b$. If this condition is not satisfied, then $j = 0$: *i.e.*, the rational buyer does not search and does not participate in the market at that time.

The buyer population at a given moment is characterized by a strategy vector \mathbf{w} , in which the component w_i represents the fraction of buyers employing strategy i and $\sum_{i=0}^S w_i = 1$. A seller s 's expected profit per unit time π_s is a function of the strategy vector \mathbf{w} , the price vector \mathbf{p} describing all sellers' prices, and the cost of production r_s for seller s . In particular, $\pi_s(\mathbf{p}, \mathbf{w}) = D_s(\mathbf{p}, \mathbf{w})(p_s - r_s)$, where $D_s(\mathbf{p}, \mathbf{w})$ is the rate of demand for the good produced by seller s , given the current price and search strategy vectors. The demand $D_s(\mathbf{p}, \mathbf{w})$ is the product of (i) the overall buyer rate of demand $\rho = \sum_b \rho_b$, (ii) the likelihood that seller s is selected as a potential seller, denoted $h_s(\mathbf{p}, \mathbf{w})$, and (iii) the fraction of buyers whose valuations satisfy $v_b \geq p_s$, denoted $g(p_s)$. Specifically, $D_s(\mathbf{p}, \mathbf{w}) = \rho B h_s(\mathbf{p}, \mathbf{w}) g(p_s)$. Without loss of generality, we define the time scale such that $\rho B = 1$. Then we can interpret π_s as seller s 's expected profit per unit sold systemwide.

The likelihood of a given buyer selecting seller s as their potential seller, namely $h_s(\mathbf{p}, \mathbf{w})$, depends on the search strategies of the buyers. In particular, this term is the sum over all buyer types of the fraction of the buyer population of type i times the probability $h_{s,i}(\mathbf{p})$ that seller s is selected by a buyer of type i , namely $h_s(\mathbf{p}, \mathbf{w}) = \sum_{i=0}^S w_i h_{s,i}(\mathbf{p})$. The quantities $h_{s,i}(\mathbf{p})$ are investigated in detail in the following section. Finally, the value $g(p_s) = \int_{p_s}^{\infty} \gamma(x) dx$, where $\gamma(x)$ is the probability density function describing the likelihood that a given buyer has valuation x . For example, if all buyers have the same valuation v , *i.e.*, $v_b = v$, then $\gamma(x)$ is the Dirac delta function $\delta(v - x)$, and the integral yields a step function $g(p_s) = \Theta(v - p_s)$, equal to 1 when $p_s \leq v$ and 0 otherwise. Assuming all buyers have equal valuations v ,² and all sellers share the same cost r , the profit function can now be expressed as follows: $\pi_s(\mathbf{p}, \mathbf{w}) = h_s(\mathbf{p}, \mathbf{w})(p_s - r)$, if $p_s \leq v$, but otherwise, $\pi_s(\mathbf{p}, \mathbf{w}) = 0$.

3 Analysis

In this section, we present a game-theoretic analysis of the prescribed model, assuming sellers are rational (*i.e.*, utility maximizers). Initially, we focus entirely on the strategic decision-making of rational sellers, by assuming the distribution of the buyer population is fixed and exogenously determined. Later, we extend our analysis to rational buyers, thereby permitting \mathbf{w} to vary.

A Nash equilibrium is a vector of prices at which sellers maximize their individual profits and from which they have no incentive to deviate [10]. There

² In this case, \mathbf{w} can be interpreted as representing a mixed search strategy of a single buyer who creates all of the demand in the system.

are no pure strategy Nash equilibria for this model [7]. There does, however, exist a symmetric Nash equilibrium in mixed strategies, which we derive presently. Let $f(p)$ denote the probability density function according to which sellers set their equilibrium prices, and let $F(p)$ be the corresponding cumulative distribution function. In the range for which it is defined, $F(p)$ has no mass points, since otherwise a seller could decrease its price by an arbitrarily small amount and experience a discontinuous increase in profits. Moreover, there are no gaps in the distribution, since otherwise prices would not be optimal — a seller charging a price at the low end of the gap could increase its price to fill the gap while retaining its market share, thereby increasing its profits.

The cumulative distribution function $F(p)$ is computed in terms of the probability $h_s(\mathbf{p}, \mathbf{w})$ that buyers select seller s as their potential seller. This quantity is the sum of $h_{s,i}(\mathbf{p})$ over $0 \leq i \leq S$. The first component $h_{s,0}(\mathbf{p}) = 0$. Consider the next component $h_{s,1}(\mathbf{p})$. Buyers of type 1 select sellers at random; thus, the probability that seller s is selected by such buyers is simply $h_{s,1}(\mathbf{p}) = 1/S$. Now consider buyers of type 2. In order for seller s to be selected by a buyer of type 2, s must be included within the pair of sellers being sampled — which occurs with probability $(S-1)/\binom{S}{2} = 2/S$ — and s must be lower in price than the other seller in the pair. Since, by the assumption of symmetry, the other seller's price is drawn from the same distribution, this occurs with probability $1 - F(p)$. Therefore $h_{s,2}(\mathbf{p}) = (2/S)[1 - F(p)]$. In general, seller s is selected by a buyer of type i with probability $\binom{S-1}{i-1}/\binom{S}{i} = i/S$, and seller s is the lowest-priced among the i sellers selected with probability $[1 - F(p)]^{i-1}$, since these are $i-1$ independent events. Thus, $h_{s,i}(\mathbf{p}) = (i/S)[1 - F(p)]^{i-1}$, and³ $h_s(p) = \frac{1}{S} \sum_{i=1}^S i w_i [1 - F(p)]^{i-1}$.

The precise value of $F(p)$ is determined by noting that a Nash equilibrium in mixed strategies requires that all pure strategies that are assigned positive probability yield equal payoffs, since otherwise it would not be optimal to randomize. In particular, the expected profits earned by seller s , namely $\pi_s(p) = h_s(p)(p-r)$, are constant for all prices p . The value of this constant can be computed from its value at the boundary $p = v$; note that $F(v) = 1$ because no rational seller charges more than any buyer is willing to pay. This leads to the following relation: $h_s(p)(p-r) = h_s(v)(v-r) = \frac{1}{S} w_1 (v-r)$. Now solving for p in terms of F yields:

$$p(F) = r + \frac{w_1(v-r)}{\sum_{i=1}^S i w_i [1-F]^{i-1}} \quad (1)$$

Eq. 1 has several important implications. First of all, in a population in which there are no buyers of type 1 (*i.e.*, $w_1 = 0$) the sellers charge the production cost c and earn zero profits; this is the traditional Bertrand equilibrium. On the other hand, if the population consists of just two buyer types, 1 and some $i \neq 1$, then it is possible to invert $p(F)$ to obtain:

$$F(p) = 1 - \left[\left(\frac{w_1}{i w_i} \right) \left(\frac{v-p}{p-r} \right) \right]^{\frac{1}{i-1}} \quad (2)$$

³ In the final equation, $h_s(p)$ is expressed as a function of seller s 's scalar price p , since we average over all other components of the price vector.

The case in which $i = S$ was studied previously by Varian [13]; in this model, buyers either choose a single seller at random (type 1) or search all sellers and choose the lowest-priced among all sellers (type S).

Since $F(p)$ is a cumulative probability distribution, it is only valid in the domain for which its valuation is between 0 and 1. As noted previously, the upper boundary is $p = v$; the lower boundary p^* can be computed by setting $F(p^*) = 0$ in Eq. 1, which yields:

$$p^* = r + \frac{w_1(v - r)}{\sum_{i=1}^S iw_i} \quad (3)$$

In general, Eq. 1 cannot be inverted to obtain an analytic expression for $F(p)$. It is possible, however, to plot $F(p)$ without resorting to numerical root finding techniques. We use Eq. 1 to evaluate p at equally spaced intervals in $F \in [0, 1]$; this produces unequally spaced values of p ranging from p^* to v .

We now consider the probability density function $f(p)$. Differentiating both sides of the equation $h_s(p)(p - r) = \frac{1}{S}w_1(v - r)$, we obtain an expression for $f(p)$ in terms of $F(p)$ and p that is conducive to numerical evaluation:

$$f(p) = \frac{w_1(v - r)}{(p - r)^2 \sum_{i=2}^S i(i - 1)w_i [1 - F(p)]^{i-2}} \quad (4)$$

The values of $f(p)$ at the boundaries p^* and v are as follows:

$$f(p^*) = \frac{\left[\sum_{i=1}^S iw_i \right]^2}{w_1(v - r) \left[\sum_{i=2}^S i(i - 1)w_i \right]} \quad \text{and} \quad f(v) = \frac{w_1}{2w_2(v - r)} \quad (5)$$

Fig. 1(a) and 1(b) depict the PDFs in the prescribed model under varying distributions of buyer strategies — in particular, $w_1 = 0.2$ and $w_2 + w_S = 0.8$ — when $S = 5$ and $S = 20$, respectively. In both figures, $f(p)$ is bimodal when $w_2 = 0$, as is derived in Eq. 5. Most of the probability density is concentrated either just above p^* , where sellers expect low margins but high volume, or just below v , where they expect high margins but low volume. In addition, moving from $S = 5$ to $S = 20$, the boundary p^* decreases, and the area of the no-man's land between these extremes diminishes. In contrast, when $w_2, w_S > 0$, a peak appears in the distribution. If a seller does not charge the absolute lowest price when $w_2 = 0$, then it fails to obtain sales from any buyers of type S . In the presence of buyers of type 2, however, sellers can obtain increased sales even when they are priced moderately. Thus, there is an incentive to price in this manner, as is depicted by the peak in the distribution. The case in which $w_S = 0$: *i.e.*, $w_1 + w_2 = 1$ is explored in more detail in the next section.

Recall that the profit earned by each seller is $(1/S)w_1(v - r)$, which is strictly positive so long as $w_1 > 0$. It is as though only buyers of type 1 are contributing to sellers' profits, although the actual distribution of contributions from buyers of type 1 vs. buyers of type $i > 1$ is not as one-sided as it appears. In reality, buyers

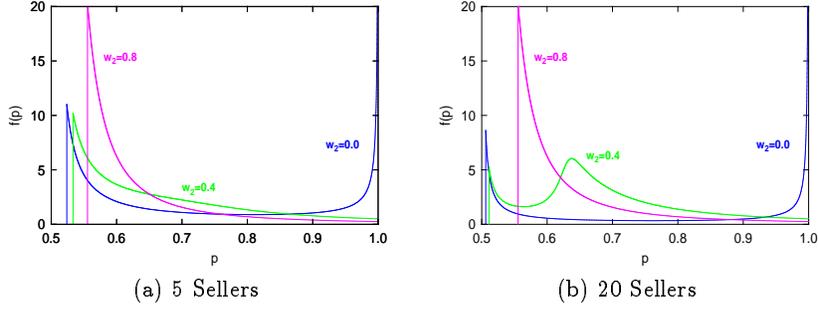


Fig. 1. PDFs for $w_1 = 0.2$ and $w_2 + w_{20} = 0.8$.

of type 1 are charged less than v on average, and buyers of type $i > 1$ are charged more than r on average, although total profits are equivalent to what they would be if the sellers practiced perfect price discrimination. In effect, buyers of type 1 exert negative externalities on buyers of type $i > 1$, by creating surplus profits for sellers.

3.1 Endogenous Buyer Decisions

Heretofore in our analysis, we have assumed rational decision-making on the part of the sellers, but an exogenous distribution of buyer types. It is also of interest to consider buyers as rational decision-makers, with the cost c_i of comparing the prices of i sellers defined explicitly, thereby giving rise to endogenous search behavior. As mentioned previously, rational buyers estimate the commodity's price \hat{p}_i that would be obtained by searching among i sellers, and select the strategy i^* that minimizes $\hat{p}_i + c_i$, provided that $\hat{p}_i + c_i \leq v_b$; otherwise, the buyer does not search and does not participate in the marketplace.

Before studying the decision-making processes of individual buyers, it is useful to analyze the distributions of prices paid by buyers of various types and their corresponding averages at equilibrium. Recall that a buyer who obtains i price quotes pays the lowest of the i prices. (At equilibrium, the sellers' prices never exceed v since $F(v) = 1$, so a buyer is *always* willing to pay the lowest price.) The cumulative distribution for the minimal values of i independent samples taken from the distribution $f(p)$ is given by $Y_i(p) = 1 - [1 - F(p)]^i$. Differentiation with respect to p yields the probability distribution: $y_i(p) = i f(p) [1 - F(p)]^{i-1}$. The average price for the distribution $y_i(p)$ can be expressed as follows:

$$\bar{p}_i = \int_{p^*}^v dp p y_i(p) = v - \int_{p^*}^v dp Y_i(p) = p^* + \int_0^1 dF \frac{(1-F)^i}{f} \quad (6)$$

where the first equality is obtained via integration by parts, and the second depends on the observation that $dp/dF = [dF/dp]^{-1} = \frac{1}{f}$. Combining Eqs. 1, 4, and 6 would lead to an integrand expressed purely in terms of F . Integration

over the variable F (as opposed to p) is advantageous because F can be chosen to be equispaced, as standard numerical integration techniques require.

Fig. 2(a) depicts sample price distributions for buyers of various types: $y_1(p)$, $y_2(p)$, and $y_{20}(p)$, when $S = 20$ and $(w_1, w_2, w_{20}) = (0.2, 0.4, 0.4)$. The dashed lines represent the average prices \bar{p}_i for $i \in \{1, 2, 20\}$ as computed by Eq. 6. The blue line labeled *Search-1*, which depicts the distribution $y_1(p)$, is identical to the green line labeled $w_2 = 0.4$ in Fig. 1(b), since $y_1(p) = f(p)$. In addition, the distributions shift toward lower values of p for those buyers who base their buying decisions on information pertaining to more sellers.

Fig. 2(b) depicts the average buyer prices obtained by buyers of various types, when w_1 is fixed at 0.2 and $w_2 + w_{20} = 0.8$. The various values of i (*i.e.*, buyer types) are listed to the right of the curves. Notice that as w_{20} increases, the average prices paid by those buyers who perform relatively few searches increases rather dramatically for larger values of w_{20} . This is because w_1 is fixed, which implies that the sellers' profit surplus is similarly fixed; thus, as more and more buyers perform extensive searches, the average prices paid by those buyers decreases, which causes the average prices paid by the less diligent searchers to increase. The situation is slightly different for those buyers who perform larger searches but do not search the entire space of sellers: *e.g.*, $i = 10$ and $i = 15$. These buyers initially reap the benefits of increasing the number of buyers of type 20, but eventually their average prices increase as well. Given a fixed portion of the population designated as buyers of type 1, Fig. 2(b) demonstrates that searching S sellers is a superior buyer strategy to searching $1 < i < S$ sellers. Thus, there is value in performing price searches: *shopbots offer added value* in markets in which there exist buyers who shop at random. This observation leads us directly into a discussion of explicit buyer search costs.

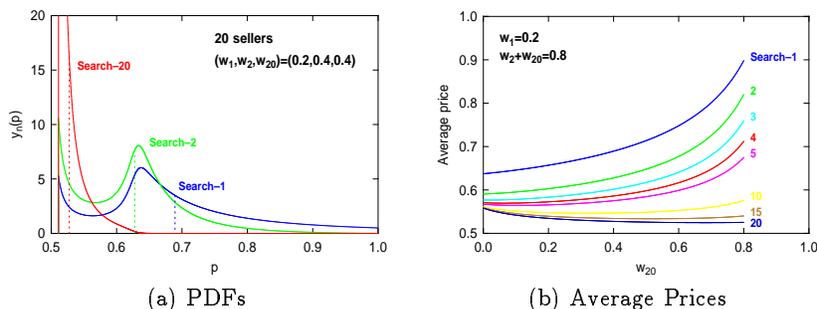


Fig. 2. (a) Buyer price distributions for 20 sellers, with $w_1 = 0.2$, $w_2 = 0.4$, $w_{20} = 0.4$. (b) Average buyer prices for various buyer types; 20 sellers, $w_1 = 0.2$, $w_2 + w_{20} = 0.8$.

Initially, we model buyer search costs following Burdett and Judd [2], who assume costs are linear in the number of searches; in particular, $c_i = c_1 + \delta(i-1)$, where $c_1, \delta > 0$ are, respectively, fixed and variable costs of obtaining price

quotes. Moreover, we assume buyers are rational decision-makers who strive to minimize overall expenditure, and who use \bar{p}_i (as in Eq. 6) as an estimate of \hat{p}_i . Thus, an optimal buyer strategy i^* satisfies: $i^* \in \arg \min_{0 \leq i \leq S} \bar{p}_i + c_i$. At equilibrium, $0 < w_1 \leq 1$, since if $w_1 = 0$, then all buyers perform some degree of search, in which case all sellers charge the competitive price r (see Eqs. 2 and 3), from which it follows that it is in fact not rational for buyers to search at all, leading to the contradiction that $w_1 = 1$. Now since the buyer cost function $\bar{p}_i + c_i$ is convex, it is minimized at either a single integer value i^* , or two consecutive integer values i^* and $i^* + 1$. Thus, at equilibrium, either $w_1 = 1$, in which case all sellers charge the monopolistic price v , or $w_1 + w_2 = 1$ and the sellers' prices are given by the distribution $f(p)$.⁴

In the case where $w_1 + w_2 = 1$, we can obtain analytic expressions for the average prices seen by buyers of types 1 and 2:

$$\bar{p}_1 = p^* + \frac{(-1 + w_2) \left(\frac{2w_2}{1+w_2} + \log \left(\frac{1-w_2}{1+w_2} \right) \right)}{2w_2} (v - r) \quad (7)$$

$$\bar{p}_2 = p^* + \frac{(1 - w_2) \left(2w_2 + (1 - w_2^2) \log \left(\frac{1-w_2}{1+w_2} \right) \right)}{2w_2^2 (1 + w_2)} (v - r) \quad (8)$$

Fig. 3(a) plots \bar{p}_1 (*i.e.*, *Search-1*) and \bar{p}_2 (*i.e.*, *Search-2*) as a function of w_2 . Not surprisingly, these curves are downward sloping, which reflects the fact that price decreases on average as the degree of search increases.

Fig. 3(b) plots the marginal cost of obtaining only one price quote rather than searching for two. More specifically, this figure displays $\bar{p}_1 - \bar{p}_2$ as a function of w_2 . Notice that there exist $\delta > 0$ such that $\bar{p}_1 = \bar{p}_2 + \delta$. In the diagram, δ is arbitrarily set at 0.02. The points of intersection between the marginal cost curve and $\delta = 0.02$ represent the points at which buyers are indifferent between obtaining a single price quote and obtaining two price quotes at price δ , but purchasing the commodity at the lower price of the two. In other words, there are two equilibria on the curve, indicated by the colored circles. Above the dotted line, the marginal cost is greater than δ ; thus, it is advantageous to search and there is momentum in the rightward direction. On the other hand, below the dotted line, the marginal cost is less than δ , and it is therefore more desirable not to search; hence, there is momentum in the leftward direction. Following the direction of the arrows, we observe that the open circle represents an unstable equilibrium, while the filled-in circle that falls on the curve is a stable equilibrium. In addition, there is a second stable equilibrium in the lower left-hand corner of the graph (indicated by a second filled-in circle) where $w_1 = 1$ and the equilibrium price is the monopolistic price v . The unstable equilibrium represents a boundary between two basins of attraction: initial values of w_2 greater than this will migrate towards the equilibrium near $w_2 = 1$, while those less than this will migrate towards $w_1 = 1$.

⁴ This depends on the assumption that c_1 is sufficiently small such that $w_0 = 0$. Otherwise, the equilibria which arise are such that $w_1 = 1 - w_0$ or $w_1 + w_2 = 1 - w_0$.

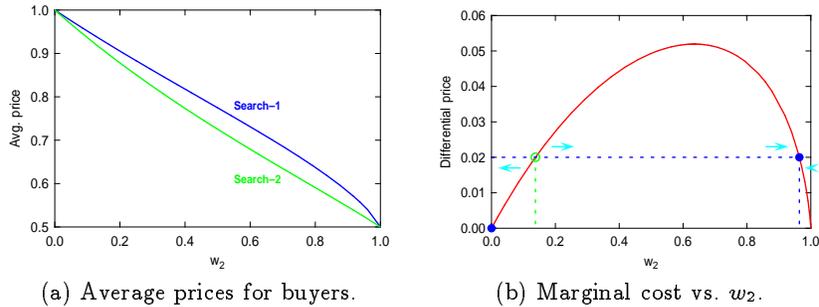


Fig. 3. An economy of buyers of type 1 and 2.

4 Shopbot Experiments

In order to explore the likely effect of shopbots on market behavior, we consider three distinctive characteristics of shopbots in turn, focusing on how they affect search costs and buyers' strategies.

First of all, a typical shopbot such as the one residing at `www.acses.com` permits users to choose the number of sellers among whom to search. Since the service is free to buyers at present, and since the search is very fast (`acsess` searches prices at 25 book retailers within about 20 seconds), there is only a very mild disincentive to requesting a large number of price quotations. Thus, the effective search cost is only weakly dependent on the number of searches. One way to model weak dependence on the number of searches is via a nonlinear search cost schedule: $c_j = c_1 + \delta(j - 1)^\alpha$, where the exponent α is in the range $0 \leq \alpha \leq 1$. Note that $\alpha = 1$ yields the linear search cost model, while $\alpha = 0$ yields a search cost that is independent of the number of searches for $j > 1$.

Second, today's shopbots are used by only a small fraction of shoppers. This is due at least in part to the fact that many potential users are unaware of the existence of shopbots, and others do not know where to find them or how to use them. One way of modeling buyers who do not use shopbots is to assume that such uninformed or "irrational" users are buyers of type 1, for which they incur fixed cost c_1 . This establishes a lower limit on the fraction w_1 , which we denote $\lfloor w_1 \rfloor$. In particular, $\lfloor w_1 \rfloor$ represents the fraction of uninformed buyers who guarantee the sellers a strictly positive profit surplus. In the following two subsections, we explore these issues in greater detail.

4.1 Nonlinear search costs

Suppose that buyers periodically (at random times) re-evaluate their search strategies and choose the strategy j that minimizes $\hat{p}_j + c_j$, where \hat{p}_j is their estimate of the average price they are likely to get by using search strategy j . One possibility is that the buyer (or an agent acting on the buyer's behalf) could use historical data on sellers' prices to estimate \hat{p}_j . However, we shall assume

here that the buyers are perfectly knowledgeable about the sellers' marginal production cost r and the current state of the strategy vector \mathbf{w} , and thus they can integrate Eq. 6 numerically to compute $\hat{p}_j = \bar{p}_j$. As the buyers modify their strategies in this manner, we assume further that the sellers monitor \mathbf{w} and instantaneously re-compute the symmetric price distribution $f(p)$ and choose their prices according to this distribution.

We can approximate this evolutionary process by a discrete time process in which, at each time step, a fraction ϵ of the buyer population is given the opportunity to switch to the optimal strategy. Then the strategy vector evolves according to: $w_i(t+1) = w_i(t) + \epsilon(\delta_{ij} - w_i(t))$, where j is the strategy that minimizes $\bar{p}_j + c_j$ and δ_{ij} represents the Kronecker delta function, equal to 1 when $i = j$ and 0 otherwise.

Fig. 4(a) illustrates the evolution of the components of \mathbf{w} in a 5-seller system when w_1 is completely endogenous ($\lfloor w_1 \rfloor = 0$, and the search costs are linear ($\alpha = 1$, $c_1 = 0.05$, and $\delta = 0.02$). The value of ϵ is 0.005. Recall that according to Burdett and Judd [2], \mathbf{w} must evolve toward an equilibrium consisting of a finite number of type 1 and type 2 buyers. Indeed, this does occur, but what is most interesting is the trajectory of the \mathbf{w} on its route toward equilibrium.

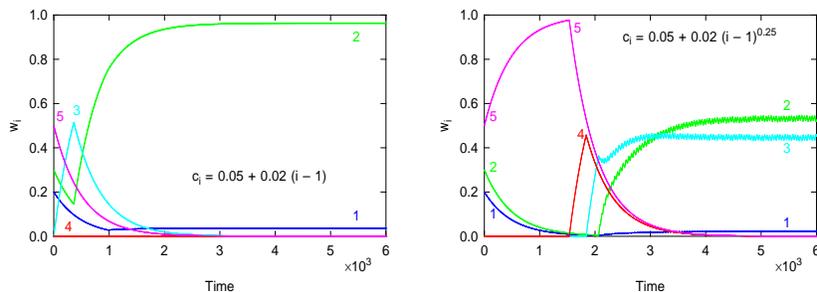


Fig. 4. (a) Evolution of indicated components of buyer strategy vector \mathbf{w} for 5 sellers, with linear search costs $c_i = 0.05 + 0.02(i - 1)$. Final equilibrium oscillates with small amplitude around theoretical solution involving a mixture of strategy types 1 and 2. (b) Evolution of indicated components of buyer strategy vector \mathbf{w} for 5 sellers, with nonlinear search costs $c_i = 0.05 + 0.02(i - 1)^{0.25}$. Final equilibrium oscillates chaotically around a mixture of strategy types 1, 2, and 3.

Initially, $\mathbf{w}_0 = (0.2, 0.3, 0.0, 0.0, 0.5)$. In this situation, the favored strategy is type 3, and so w_3 begins to grow at the expense of w_1 , w_2 and w_5 . However, as w_5 diminishes, the total amount of search in the system diminishes, and $f(p)$ flattens and shifts in such a way that eventually the favored strategy shifts from 3 to 2. Thereafter, w_2 grows at the expense of w_3 and the other components. In this simulation, near but imperfect equilibrium is achieved: due to the finite size of ϵ (equal to 0.005), there are small oscillations in w_2 around an average value that is close to the theoretical value of 0.9641721. This value can be derived by

identifying the value of w_2 corresponding to $\delta = 0.02$ in Fig. 3(b). In Fig. 3(b), there is a second value of w_2 satisfying $\delta = 0.02$, near $w_2 = 0.1375564$. However, this is the unstable equilibrium, and as discussed in the previous section it marks the boundary between two basins of attraction, one in which the final equilibrium is $(w_1, w_2) = (0.0358279, 0.9641721)$, and the other in which $(w_1, w_2) = (1, 0)$.

The derivation of an equilibrium in which only type 1 and type 2 strategies could co-exist was founded on the assumption that search costs are linear in the amount of search. In order to investigate the effect of nonlinear search costs that grow only weakly with the amount of search, we run the same experiment, in which all parameters are identical except for the exponent α , which is reduced from 1.0 to 0.25. Fig. 4(b) depicts the result. Interestingly, in this case the system evolves to an equilibrium in which types 1, 2 and 3 co-exist: w oscillates around the value $(0.0217, 0.5357, 0.4426, 0.0000, 0.0000)$ in a way that appears to be chaotic, but it remains to conduct further tests of this phenomenon. While the chaotic oscillations are an artifact of the finite size of ϵ , and would disappear in the limit $\epsilon \rightarrow 0$, they hint that the system would undergo large-scale nonlinear and possibly chaotic oscillations if the buyers were to revise their strategies synchronously rather than asynchronously.

4.2 Lower limit on w_1

In order to explore the consequences of some proportion of users failing to adopt low-cost search methods (perhaps due to ignorance about their existence or about how to use them), we now impose a lower limit on w_1 , denoted $\lfloor w_1 \rfloor$. Fig. 5(a) depicts the result of imposing $\lfloor w_1 \rfloor = 0.04$, with linear search costs $c_i = 0.05 + 0.005(i - 1)$. Starting from an initial strategy vector $w_0 = (0.04, 0.20, 0.00, 0.00, 0.76)$, the system evolves to an equilibrium in which only types 1 and 4 co-exist, with $w_1 = 0.04$ and $w_4 = 0.96$.

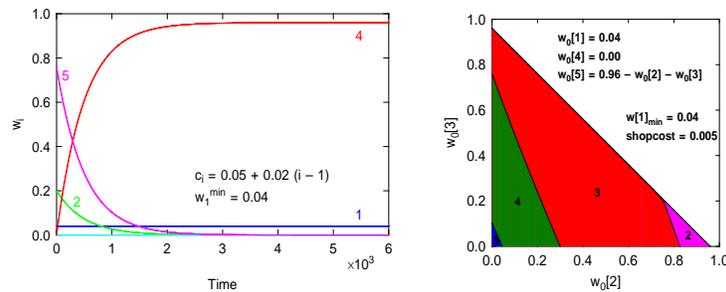


Fig. 5. (a) Evolution of indicated components of buyer strategy vector w for 5 sellers, with linear search costs $c_i = 0.05 + 0.005(i - 1)$ and $\lfloor w_1 \rfloor = 0.04$. Starting from the initial w indicated in the text, the strategy vector evolves towards an equilibrium in which only types 1 and 4 are present. (b) Two-dimensional cross-section of basin of attraction for $(\lfloor w_1 \rfloor, \delta) = (0.04, 0.005)$.

In numerous experiments with linear search costs, we have observed that the final equilibrium always consists of a mixture of types 1 and i , where i is not necessarily 2, as it must be when w_1 is determined in an entirely endogenous fashion. The strategy i depends on the values of $\lfloor w_1 \rfloor$ and δ . Table 1 illustrates the dependence of the strategy i that mixes with strategy 1 upon $\lfloor w_1 \rfloor$ and the incremental cost δ . Higher values of $\lfloor w_1 \rfloor$ lead to higher equilibrium strategies i (more extensive search) while higher incremental costs δ lead to lower equilibrium strategies i (less extensive search). For the table entries $(\lfloor w_1 \rfloor, \delta) = (0.04, 0.005)$ and $(\lfloor w_1 \rfloor, \delta) = (0.20, 0.020)$, multiple equilibria are obtained. In these cases, the initial setting of the strategy vector determines which equilibrium is obtained.

The effect of initial conditions on equilibrium selection in the case $(\lfloor w_1 \rfloor, \delta) = (0.04, 0.005)$ is illustrated in Fig. 5(b). Four equilibria are possible, all of the form $w_1 + w_i = 1$, for $i = 2, 3, 4, 5$. The set of initial conditions leading to equilibrium i — its “basin of attraction” — forms a contiguous, smoothly bounded region, a two-dimensional cross-section of which is depicted in Fig. 5(b).

$\lfloor w_1 \rfloor$	$\delta = 0.001$	$\delta = 0.005$	$\delta = 0.020$
0.01	5	2	2
0.04	5	2-5	2
0.20	5	5	2-3

Table 1. Search strategy or strategies that co-exist with type 1 search strategy, as a function of $\lfloor w_1 \rfloor$ and incremental cost δ .

5 Conclusions and Future Work

Our desire to explore the economic impact of shopbots in obtaining price and product information has led us to a model that is similar in spirit to those that have been investigated by economists interested in understanding the phenomenon of price dispersion. Our goals, however, are *prescriptive*, rather than *descriptive*, leading us to consider somewhat different causes and effects than are typical of price dispersion studies. Ultimately, we are interested in designing economically-motivated software agents, as well as an infrastructure that will support their interactions; thus, we have emphasized the constructive computation of price distributions and averages, rather than merely providing classical proofs of existence and other properties of equilibria.

Arguing that nonlinear search cost schedules are likely to exist naturally, or might even be adopted intentionally by shopbots, we studied their effect within the context of our model; our findings reveal that nonlinear search costs can lead to more complicated mixtures of buyer strategies and more extensive search than occur with linear costs. Another practical assumption, namely the existence of a positive number of uninformed buyers who do not use search mechanisms,

can lead to similar outcomes. Taking evolutionary dynamics of buyer strategies into account, we found that the final equilibrium strategy vector depends on its initial value, and the route toward equilibrium can be surprisingly complicated.

In closing, we briefly mention two promising areas for future work. First, combining the evolutionary dynamics of buyers with more interesting and realistic models for seller pricing behavior such as those described in [7, 8] would be of practical importance, and are certain to lead to interesting dynamics. Secondly, since shopbots are beginning to provide additional information about product attributes, it would also be of interest to analyze and simulate a model that accounts for both horizontal [1] and vertical differentiation.

References

1. J. Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 43(12), December 1997.
2. K. Burdett and K. L. Judd. Equilibrium price dispersion. *Econometrica*, 51(4):955–969, July 1983.
3. A. Chavez and P. Maes. Kasbah: an agent marketplace for buying and selling goods. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London, U.K., April 1996.
4. J. B. DeLong and A. M. Froomkin. The next economy? In D. Hurley, B. Kahin, and H. Varian, editors, *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*. MIT Press, Cambridge, Massachusetts, 1998.
5. J. Eriksson, N. Finne, and S. Janson. Information and interaction in MarketSpace — towards an open agent-based market infrastructure. In *Proceedings of the Second USENIX Workshop on Electronic Commerce*, November 1996.
6. H. Green. Good-bye to fixed pricing? *Business Week*, pages 71–84, May 4 1998.
7. A. Greenwald. *Learning to Play Network Games*. Ph.D. Dissertation, New York University, New York, Expected May 1999.
8. J. O. Kephart, J. E. Hanson, D. W. Levine, B. N. Grosz, J. Sairamesh, R. B. Segal, and S. R. White. Dynamics of an information filtering economy. In *Proceedings of the Second International Workshop on Cooperative Information Agents*, 1998.
9. T. G. Lewis. *The Friction-Free Economy: Marketing Strategies for a Wired World*. Hardcover HarperBusiness, 1997.
10. J. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
11. S. Salop and J. Stiglitz. Bargains and rip-offs: a model of monopolistically competitive price dispersion. *Review of Economic Studies*, 44:493–510, 1977.
12. M. Tsvetovatyy, M. Gini, B. Mobasher, and Z. Wieckowski. MAGMA: an agent-based virtual market for electronic commerce. *Applied Artificial Intelligence*, 1997.
13. H. Varian. A model of sales. *American Economic Review, Papers and Proceedings*, 70(4):651–659, September 1980.