Shopbot Economics

Jeffrey O. Kephart (kephart@us.ibm.com)
IBM Institute for Advanced Commerce
IBM Thomas J. Watson Research Center
Yorktown Heights, NY 10598 USA

Amy R. Greenwald (amygreen@cs.brown.edu)
Department of Computer Science
Brown University, Box 1910
Providence, RI 02912 USA

Abstract. Shopbots are Internet agents that automatically search for information pertaining to the price and quality of goods and services. As the prevalence and usage of shopbots continues to increase, one might expect the resultant reduction in search costs to alter market behavior significantly. We explore the potential impact of shopbots upon market dynamics by proposing, analyzing, and simulating a model that is similar in form to some that have been studied by economists investigating the phenomenon of price dispersion. However, the underlying assumptions and methodology of our approach are different, since our ultimate goal is not to explain human economic behavior, but rather to design economic software agents and study their behavior. We study markets consisting of shopbots and other agents representing buyers and sellers in which (i) search costs are nonlinear, (ii) some portion of the buyer population makes no use of search mechanisms, and (iii) shopbots are economically motivated, strategically pricing their information services so as to maximize their own profits. Under these conditions, we find that the market can exhibit a variety of hitherto unobserved dynamical behaviors, including complex limit cycles and the co-existence of several buyer search strategies. We also demonstrate that a shopbot that charges buyers for price information can manipulate markets to its own advantage, sometimes inadvertently benefitting buyers and sellers.

Keywords: Shopbots, Economic software agents, Price dispersion, Search costs

1. Introduction

Shopbots—software agents that automatically query multiple on-line vendors to gather information about prices and other attributes of consumer goods and services—herald a future in which autonomous agents are an essential component of nearly every facet of electronic commerce [5, 16, 26, 8]. In response to a buyer's expressed interest in an item, a shopbot can query several dozen sellers' web sites, and return sorted information within seconds. Shopbots outperform and out-inform humans by providing extensive product coverage in just a few seconds, far more than a patient, determined human shopper could achieve after hours of manual search.

Since the launch of the BargainFinder [17], a CD shopbot, in 1995, scores of shopbots have emerged. One popular shopbot, `mySimon.com`, provides price and product information for groceries, apparel, office supplies, consumer electronics, toys, and other goods, allowing shoppers to sort by price or by merchant reputation as provided by `Gomez.com`. Another shopbot, `DealTime.com`, reports prices and expected delivery times on a variety of products offered for sale on line, including books, CDs, and movies. It computes the partition of a set of goods across multiple vendors that minimizes the total price, taking shipping charges into account.

Shopbots deliver on one of the great promises of the Internet and electronic commerce: a radical reduction in the economic friction that is associated with the cost of obtaining and distributing information. Transportation costs, menu costs (the costs to sellers of evaluating, updating, and advertising prices), and search costs (the costs to buyers of seeking out optimal price and quality) are all expected to decrease, as a consequence of the digital nature of information as well as the presence of autonomous agents that collect, process, and disseminate that information at little cost. The fluent manipulation of freely flowing information by agents is predicted to have a profound effect on the efficiency and dynamics of markets [1, 9, 18, 6].

In this paper, we explore how shopbots might influence the behavior of electronic markets of the future. More specifically, we study how shopbot-induced changes in the cost of obtaining price information might affect both the individual and the collective behavior of agent buyers and sellers that will participate in and define these electronic markets. Our work builds on economic models that were originally introduced to explain how search costs can engender price dispersion in markets[22, 27, 3]. In such work, economists typically approximate the relevant aspects of human behavior by mathematical algorithms, and compare the collective behavior that arises in their models with the observed behavior of real markets. In contrast, we regard mathematical algorithms, not as approximate descriptions of human behavior, but as precise prescriptions of software agent behavior. (Typically, these prescriptions are based on an effort to maximize each individual agent's economic welfare.) Our goal is not to explain observed market behavior, but to anticipate and influence future market behavior.

The paper is organized as follows. Section 2 presents our model of a simple market populated with electronic sellers, electronic buyers, and a shopbot that readily provides price information. The three main sections analyze the market from the perspective of each type of player in turn: first the sellers, then the buyers, and finally the shopbot. In Section 3, we present a game-theoretic analysis of the sellers' pricing

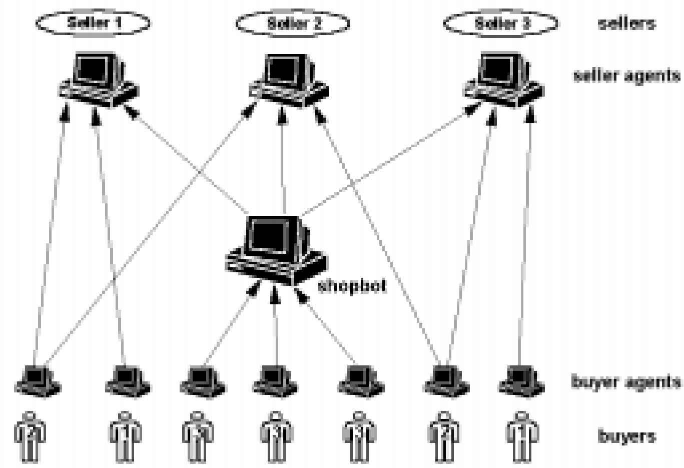

Figure 1. Shopbot Model. Buyers' numerical labels indicate the number of price quotes they compare before selecting a seller.

strategies. In Section 4, we explore the dynamic aspects of buyers' search behavior and the co-evolutionary relationship between pricing and search. Specifically, we investigate the effects of linear search costs, nonlinear search costs, and buyers who do not avail themselves of search mechanisms. In Section 5, we extend our model by providing economic incentives to the shopbot itself, investigating how it might strategically price the information service it provides. In Section 6, we discuss related work, and finally in Section 7, we summarize our findings and discuss interesting directions for future work.

2. Model

We consider a market in which there is a single commodity that is offered for sale by S sellers and of interest to B buyers (see Figure 1). Some or all of the sellers and buyers may be (or be represented by) software agents. Periodically, at the rate ρ_b, buyer b attempts to purchase a unit of the commodity. Each attempted purchase proceeds as follows. Buyer b conducts a search of fixed sample size i, which entails requesting

$0 \leq i \leq S$ price quotes.[1] We refer to such a buyer as a *Search-i* buyer, or a buyer of type i. We allow for a distribution of buyer types across the population; this distribution may be fixed exogenously, or it may arise endogenously through the efforts of buyers to adopt optimal search strategies. 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 does not enter into the buyers' purchasing decisions, however, because buyers must commit to conducting their search before any results become available. In other words, search payments are sunk costs. But when search strategies are determined endogenously, search costs do affect buyers' strategic choices. In this case, buyers are assumed to behave rationally, although they need not be fully informed. In particular, buyer b is assumed to re-evaluate its strategy at random intervals at an average rate $\sigma_b \leq \rho_b$, where typically $\sigma_b \ll \rho_b$. Upon re-evaluation, buyer b estimates a price $\hat{p_i}$ that it would expect to pay for the commodity if it were to abide by strategy i, and then selects the strategy j that minimizes $\hat{p_j} + c_j$, provided that $\hat{p_j} + c_j \leq v_b$. If this condition cannot be satisfied, then the buyer opts out of the market at that time: *i.e.*, $j = 0$. Whenever a rational buyer is fully informed, it makes an optimal decision regarding which search strategy to employ, given the current state of the market.

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

[1] This is referred to as a *nonsequential* [3] or *fixed sample size* [15] search strategy in the economics literature. 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.

systemwide. Now, seller s's profits are given by:

$$\pi_s(\vec{p}, \vec{w}) = (p_s - r_s) h_s(\vec{p}, \vec{w}) g(p_s). \tag{1}$$

The functions $h_s(\vec{p}, \vec{w})$ and $g(p)$ are analyzed in the following section.

3. Sellers

In this section, we analyze the shopbot model from the sellers' perspective. We assume that the buyers' strategies $\vec{w}$ are fixed and determined exogenously, and that the sellers seek to maximize profit. In addition, we assume that the sellers have equal production costs, *i.e.*, $r_s = r$ for all s. Under these assumptions, we derive a Nash equilibrium—a vector of prices at which sellers maximize their individual profits, and from which no seller has any incentive to deviate [20].

If $w_1 = 1$, then all buyers choose sellers at random; thus, the unique Nash equilibrium is such that all sellers charge the monopoly price. At the opposite extreme, when $w_1 = 0$, there are multiple Nash equilibria with the property that at least two sellers charge the production cost, and all sellers receive zero profits. It remains to consider $0 < w_1 < 1$. If the price quantization is sufficiently fine, then there are no pure strategy Nash equilibria. (A proof of this assertion is given in the appendix.) There does exist a set of Nash equilibria in mixed strategies, the symmetric variant of which can be derived as follows.[2]

Let $f(p)$ denote the probability density function according to which sellers set their equilibrium prices, and let $F(p)$ denote the corresponding cumulative distribution function. Following Varian [27], we note that $F(p)$ has no mass points in the range for which it is defined, 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 quantity $h_s(\vec{p}, \vec{w})$.

Recall that $h_s(\vec{p}, \vec{w})$ represents the probability that buyers select seller s as their potential seller. This function is expressed in terms of the probabilistic demand for seller s by buyers of type i, namely $h_{s,i}(\vec{p})$, for $0 \leq i \leq S$. The first component $h_{s,0}(\vec{p}) = 0$. Consider the next component, $h_{s,1}(\vec{p})$. Buyers of type 1 select sellers at random; thus, the

[2] There also exist *asymmetric* Nash equilibria that can arise dynamically in this shopbot model [13].

probability that seller s is selected by such buyers is simply $h_{s,1}(\vec{p}) = 1/S$. 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 $h_{s,2}(\vec{p}) = (2/S)\,[1 - F(p)]$. 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,

$$h_{s,i}(\vec{p}) = \frac{i}{S}[1 - F(p)]^{i-1} \tag{2}$$

and[3]

$$h_s(p, \vec{w}) = \frac{1}{S}\sum_{i=1}^{S} i w_i [1 - F(p)]^{i-1}. \tag{3}$$

Substituting Eq. 3 into Eq. 1, and dropping the subscript s due to symmetry, the expected profit for any given seller is:

$$\pi(p, \vec{w}) = \frac{1}{S}\sum_{i=1}^{S} i w_i [1 - F(p)]^{i-1}(p - r)g(p). \tag{4}$$

A Nash equilibrium in mixed strategies requires that all prices that receive positive probability yield equal profits; otherwise, it would not be optimal to randomize. Thus, the equilibrium profits given by Eq. 4 equal some constant π for all prices p that are included in the support of $f(p)$. The constant π can be derived if we can find a special value of p for which the right-hand side of Eq. 4 can be evaluated. Consider the highest price in the support of $f(p)$: *i.e.*, the smallest value of p for which $F(p) = 1$. Denote this price p_m. Evaluating Eq. 4 at p_m, we obtain $\pi = (w_1/S)(p_m - r)g(p_m)$. This equation has a simple interpretation. Since $F(p_m) = 1$, a seller that charges p_m is catering solely to buyers of type 1, and $1/S$ of such buyers will consider purchasing from it. For these buyers, the seller is in essence a monopolist. The remaining factor, $(p - r)g(p)$, is the profit function for a monopolist. Since a monopolist sets p so as to maximize this profit function, p_m can be identified as the monopolist price, which yields the monopolist profit $\pi_m \equiv (p_m - r)g(p_m)$. Thus the constant π is simply expressed as

$$\pi = \frac{w_1 \pi_m}{S}. \tag{5}$$

[3] In Eq. 3, $h_s(p, \vec{w})$ is expressed as a function of seller s's scalar price p, given that probability distribution $F(p)$ describes the other sellers' expected prices.

Substituting Eq. 5 into the left hand side of Eq. 4, we find that for any p in the support of $f(p)$,

$$(p-r)g(p) = \frac{w_1 \pi_m}{\sum_{i=0}^{S} i w_i [1-F(p)]^{i-1}}. \quad (6)$$

Eq. 6 implicitly defines p and $F(p)$ in terms of one another, and in terms of $g(p)$, for all p such that $0 \leq F(p) \leq 1$.

The function $g(p)$ can be expressed as $g(p) = \int_p^\infty \gamma(x)dx$, where $\gamma(x)$ is the probability density function describing the likelihood that a given buyer has valuation x. For example, if the buyers' valuations are uniformly distributed between 0 and v, with $v > 0$, then the integral yields $g(p) = 1 - p/v$, for $p \leq v$. This case was studied in Greenwald, *et al.* [14].

In this paper, we assume $v_b = v$ for all buyers b, in which case $\gamma(x)$ is the Dirac delta function $\delta(v-x)$, and the integral yields a step function $g(p) = \Theta(v-p)$, defined as follows:

$$\Theta(v-p) = \begin{cases} 1 & \text{if } p \leq v \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

For this distribution of buyer valuations $g(p)$, the monopolist's profit function is simply $(p-r)$ for $p \leq v$, and 0 for $p > v$. This function attains a maximum of $\pi_m = v - r$ at the price $p_m = v$. For the remainder of this paper, we shall assume (without loss of generality) that $v = 1$ and $r = 0$, and hence $\pi_m = 1$.[4] Inserting these values into Eq. 6 and solving for p in terms of F yields:

$$p(F) = \frac{w_1}{\sum_{i=1}^{S} i w_i [1-F]^{i-1}}. \quad (8)$$

Eq. 8 has several important implications. In a population in which there are no buyers of type 1 (*i.e.*, $w_1 = 0$) the sellers charge the production cost r and earn zero profits; this is the traditional Bertrand equilibrium [25]. Also of interest is the case in which the the population consists of just two buyer types, 1 and some $i \neq 1$; in this case, it is possible to invert $p(F)$ to obtain:

$$F(p) = 1 - \left[\left(\frac{w_1}{i w_i} \right) \left(\frac{1-p}{p} \right) \right]^{\frac{1}{i-1}}. \quad (9)$$

The case in which $i = S$ was studied previously by Varian [27]; in his model, buyers either choose a single seller at random (type 1) or search all sellers and choose the lowest-priced among them (type S).

[4] Results for general v and r are easily derived by renormalizing any derived prices appropriately. This price normalization, however, is specific to this particular form of $g(p)$, and is not necessarily applicable to more general distributions of valuations.

Since $F(p)$ is a cumulative probability distribution, it is only valid in the domain for which its value lies between 0 and 1. The upper boundary is $p = 1$, since prices above this threshold lead to decreases in market share that exceed the benefits of increased profits per unit. The lower boundary p^* can be computed by setting $F(p^*) = 0$ in Eq. 8, which yields:

$$p^* = \frac{w_1}{\sum_{i=1}^{S} i w_i}. \tag{10}$$

In general, Eq. 8 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 can use Eq. 8 to evaluate p at equally spaced intervals in $F \in [0,1]$; this produces unequally spaced values of p ranging from p^* to 1.

We now consider the probability density function $f(p)$. For the given choice of $g(p)$, the profits for seller s equal $h_s(p,\vec{w})p = w_1/S$, for $p \leq 1$; thus, the probabilistic demand for seller s is given by $h_s(p,\vec{w}) = w_1/pS$. Differentiating both sides of this expression with respect to p and substituting Eq. 3, we obtain an expression for $f(p)$ in terms of $F(p)$ and p that is conducive to numerical evaluation:

$$f(p) = \frac{w_1}{p^2 \sum_{i=2}^{S} i(i-1) w_i [1 - F(p)]^{i-2}}. \tag{11}$$

The values of $f(p)$ at the boundaries p^* and 1 are as follows:

$$f(p^*) = \frac{\left[\sum_{i=1}^{S} i w_i\right]^2}{w_1 \left[\sum_{i=2}^{S} i(i-1) w_i\right]} \quad \text{and} \quad f(1) = \frac{w_1}{2 w_2}. \tag{12}$$

Fig. 2 depicts the probability density function $f(p)$ for three different buyer strategies representing different mixtures of strategies 1, 2, and S, with $S = 5$ in Fig. 2(a) and $S = 20$ in Fig. 2. When $w_2 = 0$, $f(p)$ is bimodal, with most of the probability density 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. As S increases, the peaks become increasingly accentuated, and $p^* \to 0$. Using Eq. 11, it can be shown that the peak near $p = 1$ is an integrable singularity when $w_2 = 0$, and that the singularity disappears for any finite value of w_2. We have observed that this peak tends to disappear even for fairly small values of w_2. When both w_2 and w_S are of moderate size, one can observe a peak in $f(p)$ at moderate values of p. This peak can be attributed to the viability of a middle-of-the-road strategy in which sellers recognize that they need not undercut all other sellers: if

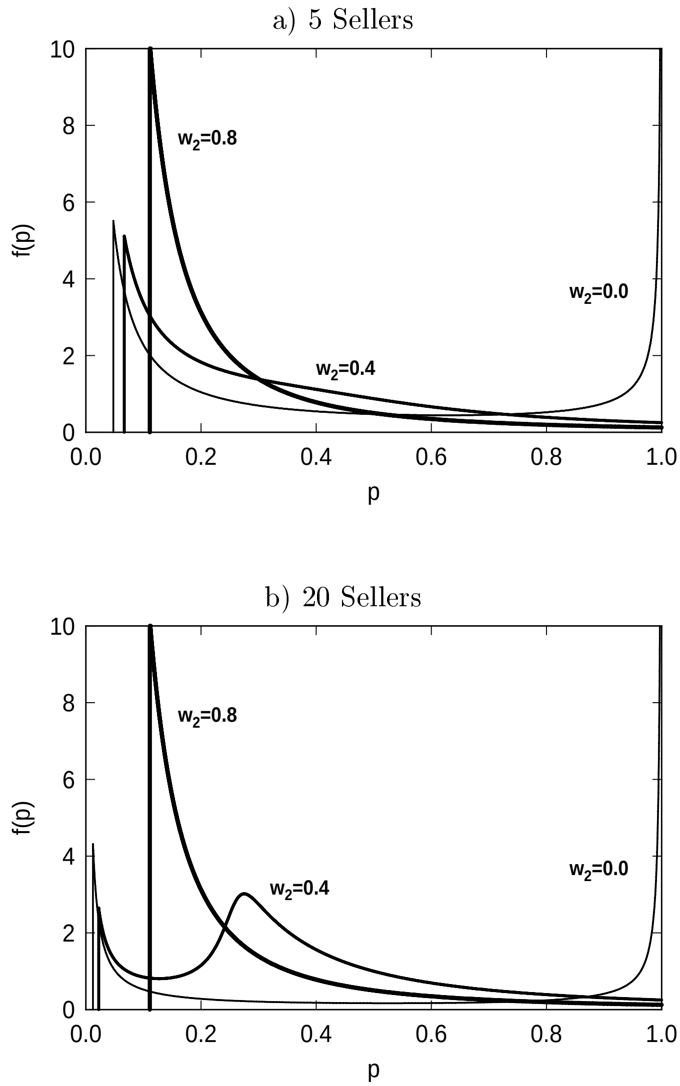

Figure 2. Probability density functions $f(p)$ for three buyer strategy vectors: $(w_1, w_2, w_S) = (0.2, 0.0, 0.8)$, $(0.2, 0.4, 0.4)$, and $(0.2, 0.8, 0.4)$. a) $S = 5$. b) $S = 20$.

a sizeable proportion of buyers only compare two sellers, then a seller need only undercut the seller with which it is randomly paired.

Since we are assuming that $\pi_m = v - r = 1$, by Eq. 5, the profit earned by each seller is $\pi = w_1/S$. Note that this quantity 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 of type 1 are charged less than $v = 1$ on average, and buyers of type $i > 1$ are charged more than $r = 0$ 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.

4. Buyers

Heretofore in our analysis, we have assumed rational decision-making on the part of the sellers, but an exogenous distribution of buyer types. In this section, we allow each buyer to use a rational criterion for choosing its search strategy, and therefore the search strategy vector $\vec{w}$ is determined endogenously.

The rational criterion that buyers employ to select a search strategy is as follows. Recall from Section 2 that search costs are specified by a vector $\vec{c}$, where c_i denotes the cost of comparing the prices of i sellers. A rational buyer estimates the expected price $\hat{p}_i$ that it would pay to the lowest-priced seller among a randomly chosen set of i sellers, given the price distribution $f(p)$, and selects the strategy i^* that minimizes $\hat{p}_i + c_i$. If no strategy i^* satisfies the restriction that $\hat{p}_i + c_i \leq v_b$, then the expected combined cost of discovering prices and purchasing the good exceeds the buyer's valuation, and the buyer opts out of the market, *i.e.*, it selects $i^* = 0$.

How do buyers estimate $\hat{p}_i$? One reasonable method would be to compute estimates using historical price data. In this study, however, we assume that buyers know the price distribution $f(p)$, and therefore compute the expected price $\bar{p}_i$ as the expected value of the lowest of i draws from $f(p)$. Using $\bar{p}_i$ as their estimate of $\hat{p}_i$, buyers can then select an optimal buyer strategy i^* that satisfies

$$i^* \in \arg\min_{0 \leq i \leq S} \bar{p}_i + c_i. \tag{13}$$

The remainder of this section is organized as follows. In Section 4.1, we derive a general expression for the expected price $\bar{p}_i$ in terms of $f(p)$, which allows us to compute the marginal benefit of the i^{th} search, namely $\bar{p}_i - \bar{p}_{i-1}$. Then, in Section 4.2, we assume search costs are linear in the amount of search undertaken. In a market of rational buyers and sellers, $\vec{w}$ and $f(p)$ co-evolve to an equilibrium in which only buyer types 1 and 2 prevail. In the following two sections, we simulate the co-evolutionary dynamics of $\vec{w}$ and $f(p)$ in two more complex scenarios: in Section 4.3 we explore the effect of nonlinear search costs; then, in

Section 4.4, we assume that a portion of the buyer population does not have access to search mechanisms.

4.1. Expected Buyer Prices

Recall that buyers who obtain i price quotes pay the lowest of the i prices observed. (At equilibrium, the sellers' prices never exceed 1 since $F(1) = 1$, so buyers are *always* willing to pay the lowest price.) Thus, the computation of $\bar{p}_i$ depends on the cumulative distribution of the minimal values of i independent samples drawn from the distribution $f(p)$, which is given by $Y_i(p) = 1 - [1 - F(p)]^i$. Differentiating both sides of this equation 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:

$$\bar{p}_i = \int_{p^*}^{v} dp\, p\, y_i(p) = v - \int_{p^*}^{v} dp\, Y_i(p) = p^* + \int_0^1 dF\, \frac{(1-F)^i}{f} \qquad (14)$$

where the first equality is obtained via integration by parts, and the second depends on the observation that $dp/dF = [dF/dp]^{-1} = 1/f$. Combining Eqs. 8, 11, and 14 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. 3(a) depicts sample price distributions for buyers of various types: in particular, $y_1(p)$, $y_2(p)$, and $y_{20}(p)$ are plotted for $S = 20$ 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. 14. The line labeled *Search–1*, which depicts the distribution $y_1(p)$, is identical to the line labeled $w_2 = 0.4$ in Fig. 2(b), since $y_1(p) = f(p)$. The distributions shift toward lower values of p for those buyers who base their buying decisions on information obtained from more sellers.

Fig. 3(b) depicts the expected prices obtained by buyers of various types, assuming 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 extensive but non-maximal searches, *e.g.*, $i = 10$ or $i = 15$. Such buyers initially reap the benefits of increasing the number of buyers of type 20,

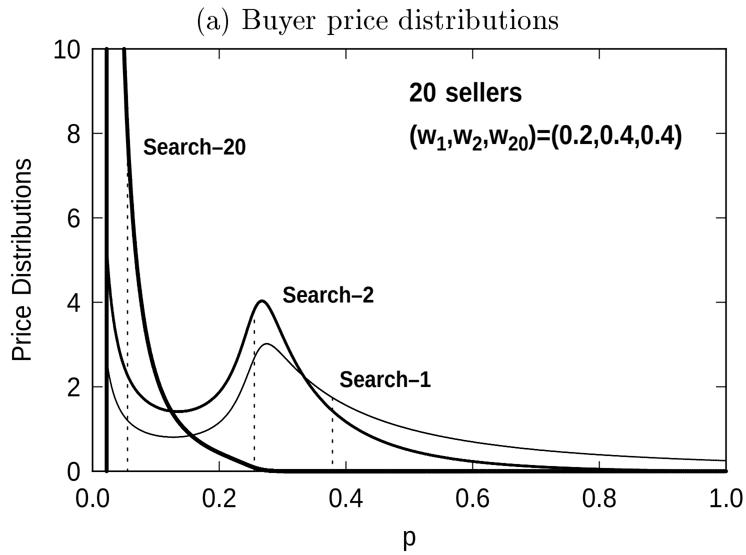

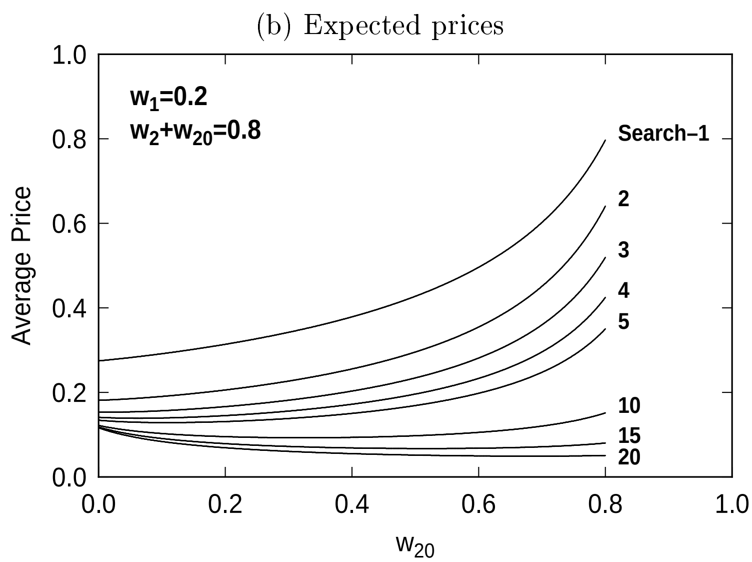

Figure 3. (a) Probability distributions of prices paid by buyers performing 1, 2, and 20 searches in a market with 20 sellers. The buyer strategy vector is $w_1 = 0.2, w_2 = 0.4, w_{20} = 0.4$, corresponding to one of the curves in Fig. 2(b). (b) Expected prices $\bar{p}_i$ paid by buyers performing the indicated amount of search, as a function of w_{20}. There are 20 sellers. The buyer strategy vector is $w_1 = 0.2$, $w_2 = 0.8 - w_{20}$.

but when w_{20} becomes sufficiently large their expected prices increase slightly as well. Given a fixed portion of the population designated as buyers of type 1, Fig. 3(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 some buyers shop randomly.

4.2. Linear Search Costs

Initially, we model buyer search costs following Burdett and Judd [3], 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 marginal costs of obtaining price quotes. As shown by Burdett and Judd, the buyer cost function $\bar{p}_i + c_i$ is a convex function of i, which is minimized at either a single integer value i^*, or two consecutive integer values i^* and $i^* + 1$. If c_1 is sufficiently large, then $w_0 = 1$ at equilibrium, *i.e.*, all buyers opt out of the market. Otherwise, w_1 is bounded away from 0, since if $w_1 = 0$ then competition drives all sellers' prices to zero (see Eq. 8), making search irrational. But if search is irrational, then $w_1 = 1$, contradicting the original assumption that $w_1 = 0$. Therefore, for sufficiently small values of c_1, equilibrium search strategies are either such that $w_1 = 1$, in which case all sellers charge the monopolistic price, or such that $w_1 + w_2 = 1$, with $w_1 > 0$ and the sellers' prices given by the distribution $f(p)$.

Given that $w_1 + w_2 = 1$, we obtain analytic expressions for the prices expected by buyers of types 1 and 2 by substituting Eq. 9 into Eq. 14:

$$\bar{p}_1(w_2) = p^* + \left(\frac{1-w_2}{2\,w_2}\right)\left[\log\left(\frac{1+w_2}{1-w_2}\right) - \frac{2\,w_2}{1+w_2}\right] \tag{15}$$

$$\bar{p}_2(w_2) = p^* + \left(\frac{(1-w_2)}{2\,{w_2}^2\,(1+w_2)}\right)\left[2\,w_2 + \left(1-w_2^2\right)\,\log\left(\frac{1-w_2}{1+w_2}\right)\right] \tag{16}$$

Fig. 4(a) plots $\bar{p}_1$ (*i.e.*, *Search–1*) and $\bar{p}_2$ (*i.e.*, *Search–2*) as a function of w_2. Both curves decrease monotonically with w_2, since sellers' prices are forced downward when more buyers compare prices. Fig. 4(b) plots the marginal benefit β of *Search–2* over *Search–1*, which is computed by subtracting Eq. 16 from 15:

$$\beta(w_2) = \left(\frac{1-w_2}{w_2}\right)\left[\frac{1}{2w_2}\log\left(\frac{1+w_2}{1-w_2}\right) - 1\right]. \tag{17}$$

Suppose that all buyers are fully informed and rational, and therefore estimate the expected marginal benefit accurately. Then, if buyers are free to choose between the strategies *Search–1* and *Search–2*, the

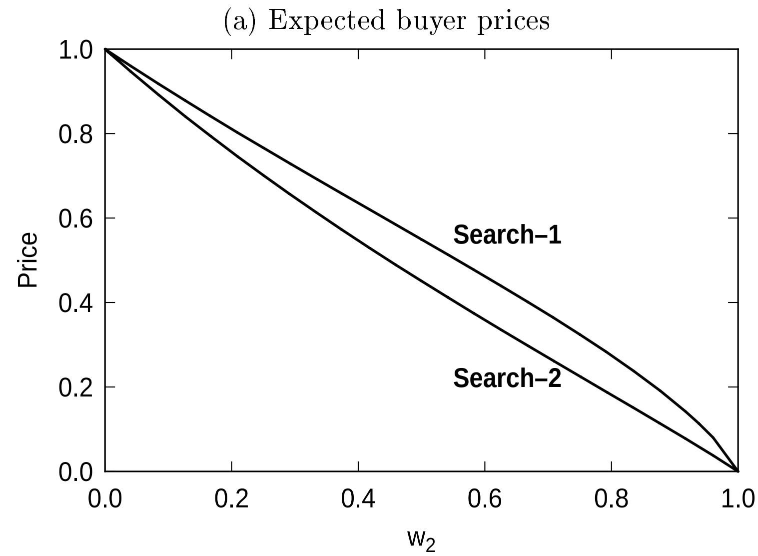

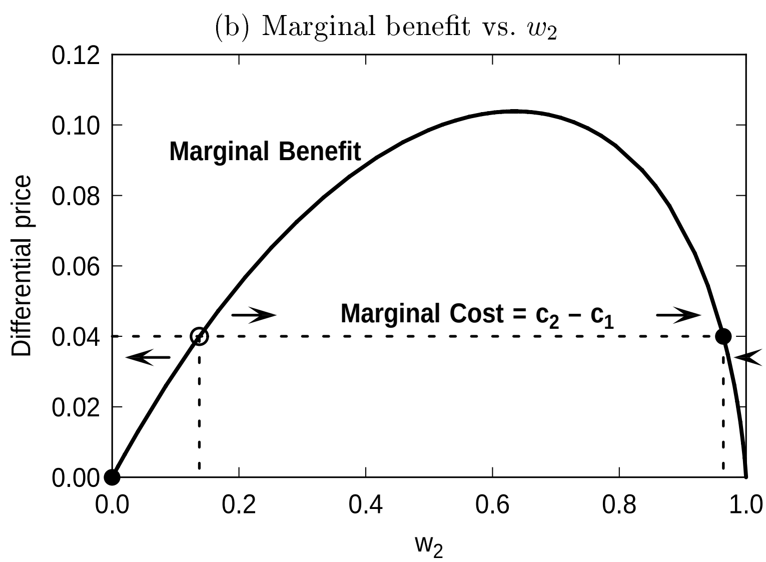

Figure 4. Benefits and costs of search when search costs are linear, and hence $w_1 + w_2 = 1$. a) Expected prices for *Search-1* and *Search-2* buyers as a function of w_2. b) Marginal benefit of *Search-2* strategy over *Search-1* as a function of w_2. Equilibria occur at values of w_2 for which marginal benefit equals marginal cost. Stable and unstable equilibria are represented by solid and open circles, respectively. Arrows indicate the direction in which a population of rational buyers would move if they were permitted to change their search strategies.

buyer population reaches equilibrium when the marginal benefit of a second price quote exactly balances its marginal cost. As an example, Fig. 4(b) graphically illustrates the situation in which the marginal cost $\delta = c_2 - c_1 = 0.04$. The marginal benefit curve and marginal cost (dotted) line intersect twice in this diagram, representing two of the three equilibria that arise in this setting. Above the dotted line, benefit outweighs cost; thus, search is advantageous and there is momentum in the rightward direction. Below the dotted line, cost outweighs benefit; it is therefore not desirable to search, and there is momentum in the leftward direction. Following the direction of the arrows, the filled-in circle that falls on the curve is a stable equilibrium, while the open circle represents an unstable equilibrium. The unstable equilibrium represents a boundary between two basins of attraction: a buyer population in which the initial value of w_2 is greater than this threshold will migrate towards the equilibrium near $w_2 = 1$, while one in which w_2 is initially smaller than this threshold will migrate towards $w_1 = 1$. Finally, there is a second stable equilibrium in the lower left-hand corner of the graph (indicated by a second filled-in circle) where $w_1 = 1$, the equilibrium price is the monopolistic price 1, and $\lim_{w_2 \to 0} \beta(w_2) = \delta = 0$.

4.3. Nonlinear search costs

A typical shopbot such as the one available at `DealTime.com` permits users to choose the number of sellers among whom to search. Since the service is free for buyers at present, and moreover, since the search is very fast—`DealTime` searches prices at a few dozen book retailers within about 20 seconds—there is only a mild disincentive not to request a large number of price quotes. 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 of the form

$$c_j = c_1 + \delta(j-1)^{\alpha} \tag{18}$$

where the exponent α is in the range $0 \leq \alpha \leq 1$. Note that $\alpha = 1$ yields a linear search cost model, while $\alpha = 0$ yields a bundled search cost model, in which the search cost is fixed at $c_1 + \delta$ regardless of the number of searches.

In Section 4.2, we performed an equilibrium analysis for $\alpha = 1$. Here, we examine non-equilibrium behavior for arbitrary α. To do so, we simulate the dynamics of a system in which the strategy vector $\vec{w}$ and the price distribution $f(p)$ co-evolve through the continual re-evaluation of search strategies by buyers and the continual re-setting of the price distribution by sellers. The buyers search strategies evolve via a discrete-time process described below. As the buyers modify their strategies, we assume that the sellers monitor the strategy vector $\vec{w}$, and instantaneously re-compute the symmetric price distribution $f(p)$ according to which they randomly set their prices.

We approximate the co-evolution of $\vec{w}$ and $f(p)$ by a discrete time process in which, at each time step, a fraction η of the buyer population switches to the optimal strategy. The vector $\vec{w}$ evolves as follows:

$$w_i(t+1) = w_i(t) + \eta(\delta_{ij} - w_i(t)), \tag{19}$$

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. Buyers are assumed to determine j by monitoring $f(p)$ (or by using Eq. 8 to compute $F(p)$ from the current state of the strategy vector $\vec{w}$) and then numerically integrating Eq. 14 to compute $\hat{p_j} = \bar{p_j}$.

Fig. 5(a) illustrates the evolution of the components of $\vec{w}$ in a 5-seller system. As in Fig. 4(b), the search costs are linear ($\alpha = 1$), with $c_1 = 0.1$ and $\delta = 0.04$. The value of η is chosen to be 0.002. According to Burdett and Judd [3], $\vec{w}$ must evolve toward an equilibrium consisting of a finite number of type 1 and type 2 buyers. Indeed, we observe such an equilibrium, but the *trajectory* of the strategy vector $\vec{w}$ *en route* to equilibrium is unexpectedly interesting.

Initially, $\vec{w}(t = 0) = (0.2, 0.3, 0.0, 0.0, 0.5)$. In this situation, the favored strategy is type 3, which begins increasing at the expense of w_1, w_2 and w_5. As w_5 diminishes, however, the total amount of search in the system diminishes, causing $f(p)$ to flatten and shift in such a way that eventually the favored strategy changes from 3 to 2. Thereafter, w_2 grows at the expense of w_3 and the other components until the system nearly stabilizes. In this simulation, near but imperfect equilibrium is achieved: due to the finite size of η, 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 $\delta = 0.04$ in Fig. 4(b). Although there is a second value of w_2 satisfying $\delta = 0.04$ in this figure, namely $w_2 = 0.1375564$, it is the unstable equilibrium. 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 co-exist is based on the assumption that search costs are linear in the amount of search. In order to investigate the effect of nonlinear search costs, we re-ran the previous experiment with identical parameters, except for the exponent α, which was reduced from 1.0 to 0.25. Fig. 5(b) depicts the ensuing dynamics. Interestingly, in this case

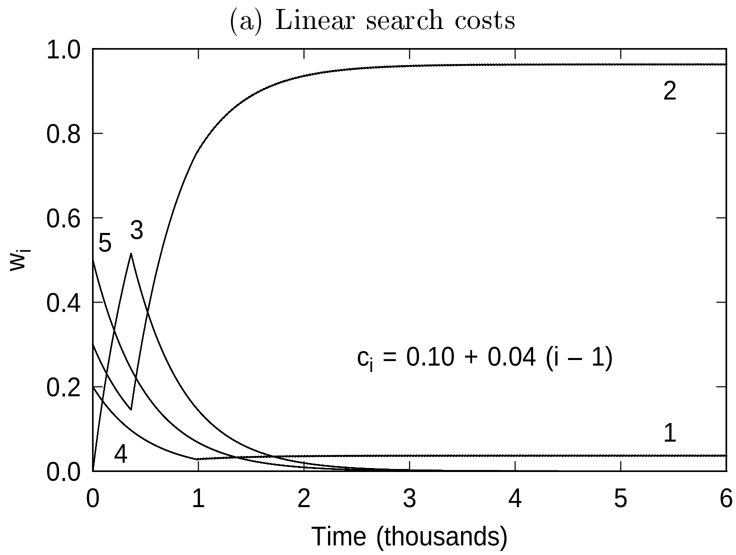

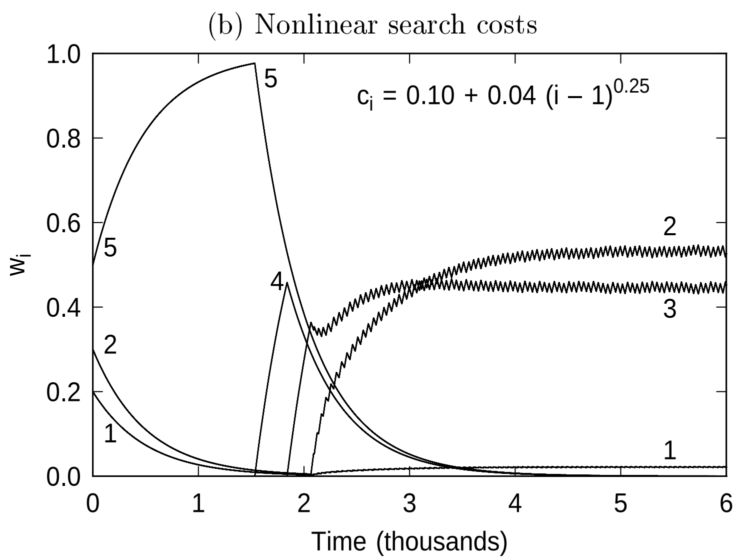

Figure 5. (a) Evolution of indicated components of buyer strategy vector $\vec{w}$ for 5 sellers, with linear search costs $c_i = 0.10 + 0.04(i-1)$. Final equilibrium oscillates with small amplitude around the theoretical solution involving a mixture of strategy types 1 and 2. (b) Evolution of indicated components of buyer strategy vector $\vec{w}$ for 5 sellers, with nonlinear search costs $c_i = 0.10 + 0.04(i-1)^{0.25}$. Final equilibrium oscillates in a complex small-amplitude limit cycle around a mixture of strategy types 1, 2, and 3.

the system evolves to an equilibrium in which *types 1, 2 and 3 co-exist*: $\vec{w}$ oscillates around the value $(0.0217, 0.5357, 0.4426, 0.0000, 0.0000)$ in a complex, small-amplitude limit cycle with period 104. As η is increased, the amplitude of the oscillations grows, but their period diminishes. Although these oscillations are an artifact of the finite size of η, and disappear in the limit $\eta \to 0$, they indicate that the system would undergo large-scale oscillations if the buyers were to revise their strategies synchronously (*i.e.*, $\eta = 1$) rather than asynchronously.

4.4. Lower limit on w_1

Today's shopbots are used by only a small fraction of shoppers. This behavior is at least in part due 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 buyers are of type 1 and incur a search cost of 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.

Fig. 6(a) depicts a simulation in which $\lfloor w_1 \rfloor = 0.04$, with linear search costs $c_i = 0.10 + 0.01(i-1)$. Starting from an initial strategy vector $\vec{w}(t=0) = (0.04, 0.20, 0.00, 0.00, 0.76)$, the system evolves to an equilibrium in which only types 1 and 4 co-exist, with $w_1 = 0.04$ and $w_4 = 0.96$. In numerous experiments with w_1 bounded below and linear search costs, we have observed that the system stabilizes with a mixture of buyer types 1 and i, where i is not necessarily 2, as it is in the traditional economic setting that was analyzed in Section 4.2. The particular strategy i that arises depends on the lower bound $\lfloor w_1 \rfloor$ and the incremental search cost δ; this dependence is illustrated in Table I. Higher values of $\lfloor w_1 \rfloor$ lead to higher values of i at equilibrium (more extensive search) while higher values of δ lead to lower values of i at equilibrium (less extensive search).

Table I. Search strategies that co-exist with *Search-1*, as a function of the lower limit $\lfloor w_1 \rfloor$ and the incremental cost δ.

$\lfloor w_1 \rfloor$	$\delta = 0.002$	$\delta = 0.010$	$\delta = 0.040$
0.01	5	2	2
0.04	5	2–5	2
0.20	5	5	2–3

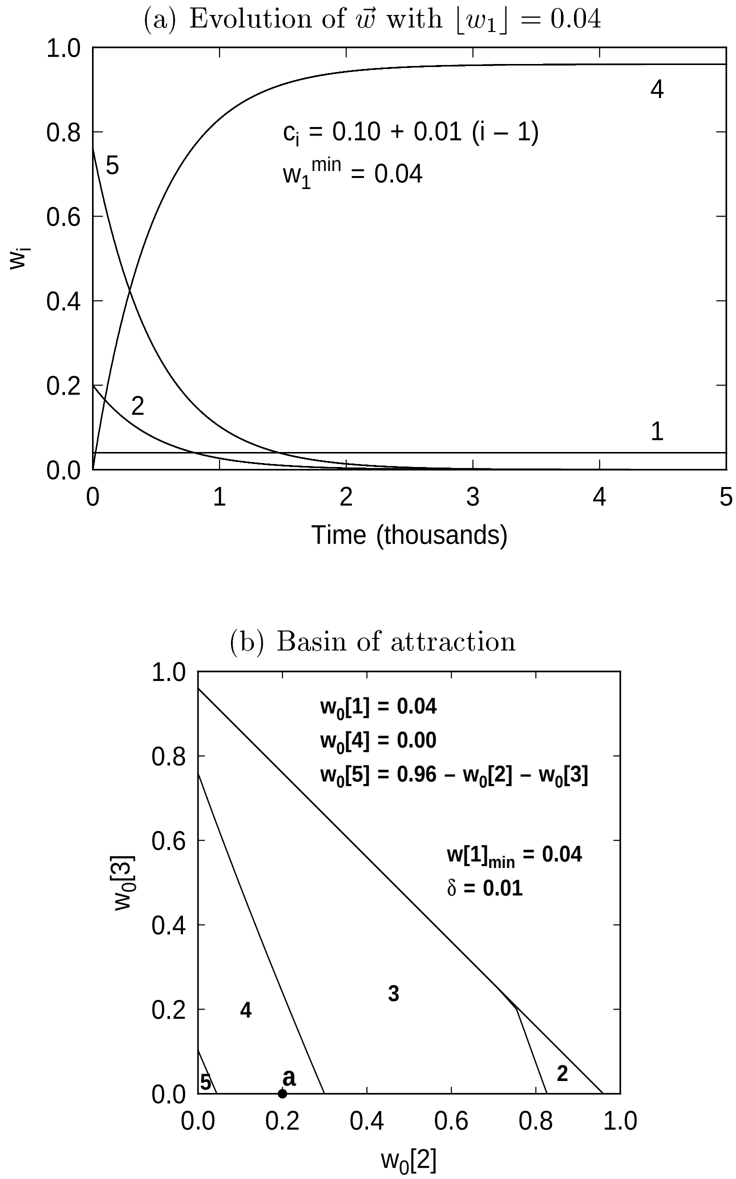

Figure 6. (a) Evolution of indicated components of buyer strategy vector $\vec{w}$ for 5 sellers, with linear search costs $c_i = 0.10 + 0.01(i-1)$ and $\lfloor w_1 \rfloor = 0.04$. Starting from an initial state with $w_1 = 0.04$, $w_2 = 0.20$, and $w_5 = 0.76$, $\vec{w}(t)$ evolves towards an equilibrium in which only w_1 and w_4 are nonzero. (b) Two-dimensional cross-section of basin of attraction for $(\lfloor w_1 \rfloor, \delta) = (0.04, 0.01)$, associated with the middle cell of Table I. Numerical labels indicate the strategy i that mixes with strategy 1 in equilibrium, given that $\vec{w}$'s initial state is in the indicated region. The simulation depicted in (a) is associated with the point marked a.

Multiple equilibria arise for the table entries $(\lfloor w_1 \rfloor, \delta) = (0.04, 0.01)$ and $(\lfloor w_1 \rfloor, \delta) = (0.20, 0.04)$, and precisely which equilibrium is reached is determined by the initial setting of the strategy vector. The effect of initial conditions on equilibrium selection in the case $(\lfloor w_1 \rfloor, \delta) = (0.04, 0.01)$ is illustrated in Fig. 6(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*[23]—forms a contiguous, smoothly bounded region, a two-dimensional cross-section of which is depicted in Fig. 6(b).

5. Shopbot

Thus far, we have followed the traditional price dispersion literature in regarding the search costs $\vec{c}$ as fixed. We now explore what might happen if a shopbot were to behave as an economically-motivated agent itself—strategically setting the search costs in an effort to maximize its own profit. Although shopbots do not currently use such a business model, it would be legitimate for them to charge buyers some fraction of the savings that they provide.[5] Since we are now viewing the market from the shopbot's perspective, we consider the search cost vector $\vec{c}$ as a *price schedule* that the shopbot sets so as to maximize its profits.

How might a shopbot compute a profit-maximizing price schedule? At first glance, the following method might seem plausible:

1. For each of a large number of candidate price schedules $\vec{c}$, simulate the market dynamics to determine the asymptotic equilibrium $\vec{w}(\vec{c})$.

2. Compute the shopbot's profit, namely[6]

$$\pi_{\mathrm{sh}}(\vec{c}, \vec{w}(\vec{c})) = \sum_{i=1}^{S} c_i w_i(\vec{c}). \tag{20}$$

3. Select the price schedule $\vec{c}$ that maximizes $\pi_{\mathrm{sh}}(\vec{c}, \vec{w}(\vec{c}))$.

However, several complications prevent such an approach from being useful in practice. First, the computational burden associated with this algorithm could be severe, as it entails optimization of a function in

[5] Baye and Morgan [2] propose an interesting alternative in which a "gatekeeper" intermediary charges fees to sellers who advertise their prices via the gatekeeper, and charges buyers to access the advertised prices. Markopoulos and Ungar [19] consider yet another scenario in which the sellers themselves charge for their own price information.

[6] Eq. 20 depends on the assumption that the shopbot has no production costs.

$S-1$ dimensions, and each function evaluation requires simulation. Second, recall from Section 4.3 that the evolutionary dynamics of buyer strategies may not lead to an equilibrium. Indeed, as illustrated in Fig. 5(b), $\vec{w}(t)$ can take the form of a complex limit cycle when search costs are nonlinear. Third, even when c_i is linear in i, and $\vec{w}$ reaches an equilibrium, there can be multiple equilibria, as illustrated in Figs. 4(b) and 6(b). In these cases, the chosen equilibrium depends on the initial value of $\vec{w}$ at time $t=0$, which is beyond the shopbot's control. Finally, the shopbot is unlikely to be the only means by which price information may be obtained. It is more natural to suppose that buyers have alternative mechanisms by which to discover prices, such as manual search, and that buyers weigh the costs and benefits of using the shopbot against those of using any alternative search mechanisms in order to decide which mechanism to use and how much search to undertake.

Taking these factors into account, we allow for a possibly time-dependent expected shopbot profit $\pi_{\mathrm{sh}}(\vec{c}, \vec{w}(t))$, given by:

$$\pi_{\mathrm{sh}}(\vec{c}, \vec{w}(t)) = \sum_{i=1}^{S} c_i \, w_i(t) \, \Theta(c_{\mathrm{alt},i} - c_i) \Theta(1 - (\bar{p}_i + c_i)) \tag{21}$$

where $\Theta(x)$ represents the step function, equal to 1 for $x \geq 0$ and 0 otherwise[7], and $\vec{c}_{\mathrm{alt}}$ represents the alternative search cost schedule. The step functions in Eq. 21 reflect the fact that, for each search strategy i, the shopbot is motivated to set c_i such that (i) it undercuts $c_{\mathrm{alt},i}$, and (ii) it is low enough to ensure that the sum of the search cost and the expected price of the good does not exceed the buyer valuation.

We explore both linear and nonlinear shopbot price schedules in the remainder of this section. First, in Section 5.1, we assume that the shopbot restricts itself to a linear price schedule of the form $c_i = c_1 + \delta(i-1)$. This assumption permits an exact analysis that yields the optimal values of c_1 and δ. An interesting complication arises, however, due to the existence of multiple equilibria. It is insufficient for the shopbot to merely fix its prices based on the optimal values of c_1 and δ. It must must also somehow ensure that the market evolves to its preferred equilibrium. We present one technique by which the shopbot can guide buyers and sellers towards this equilibrium in Section 5.2. Finally, in Section 5.3, we lift the restriction of price schedules to a

[7] By letting $\Theta(x)$ be 1 for $x \geq 0$, rather than for $x > 0$, we are in effect assuming that buyers will prefer the shopbot to the alternative mechanism if the prices are equal. One could offer plausible rationales supporting preference for either the shopbot or the alternative search mechanism when prices are equal; this choice is more cosmetic than anything else as it allows us to avoid carrying infinitesimals along through the analysis.

linear form. Rather than undertaking an analysis of time-dependent asymptotic behavior in this case, we introduce an adaptive heuristic for setting the shopbot's price schedule. We observe that the shopbot can earn significantly higher profits by setting its price schedule such that c_i is a highly nonlinear function of i. Under both scenarios, we demonstrate that the shopbot can improve the welfare of buyers and sellers.

5.1. Optimal Linear Price Schedule

In this section, we derive the optimal linear price schedule, which has the form $c_i = c_1 + \delta(i-1)$. We assume there exists an alternative search mechanism with a cost that is strictly proportional to the number of price quotes: $c_{\mathrm{alt}}(i) = ic'$. This could correspond to the cost of manual search, for example, which is at present the most natural alternative to the shopbot.

The shopbot's goal is to set the values c_1 and δ so as to maximize its expected profits. As discussed in Section 4.2, when costs are linear in the number of comparisons, a population of rational, fully-informed buyers evolves towards an equilibrium in which $w_1 + w_2 = 1$. Suppose that the shopbot takes a long-term view in that it does not factor transients into its calculations, but seeks to set a price schedule $\vec{c}$ that maximizes $\pi_{\mathrm{sh}}(\vec{c}, \vec{w})$, with $\vec{w}$ assumed to be its asymptotic equilibrium value.

Given that $\vec{w}$ stabilizes at an equilibrium in which the marginal benefit β is equal to the marginal cost δ, the shopbot's profit is:

$$\begin{aligned} \pi_{\mathrm{sh}}(\vec{c}) &= w_1 c_1 + w_2 c_2 \\ &= c_1 + w_2\,\delta \\ &= c_1 + w_2 \beta \end{aligned} \tag{22}$$

subject to the following conditions, derived from the step functions in Eq. 21:

$$c_1 \leq c' \tag{23}$$

$$c_2 \leq 2c' \tag{24}$$

$$c_1 \leq 1 - \bar{p}_1 \tag{25}$$

$$c_2 \leq 1 - \bar{p}_2 \tag{26}$$

Note that Cond. 26 can be restated as $c_1 + \delta \leq 1 - (\bar{p_1} - \beta)$, which reduces to Cond. 25, given that $\delta = \beta$ at equilibrium. Thus, only Conds. 23, 24, and 25 need be considered.

First, let us temporarily ignore Conds. 24 and 25. From Cond. 23 and Eq. 22, it is evident that c_1 should be set to c', and that δ should be set so as to maximize $w_2\beta$. We can use Eq. 17, which makes explicit the dependence of β on w_2, to solve numerically for the optimal value of w_2. The result is $w_2^* = 0.781796$, at which point $\beta^* = \beta(w_2^*) = 0.095741$.

It would appear that the shopbot could encourage the buyer population to evolve towards w_2^* by setting its price schedule such that $\delta^* = \beta^* = 0.095741$. However, fixing δ at this value is not guaranteed to have the desired effect. Recall from Fig. 4(b) that there are *two* values of w_2 that correspond to $\delta^* = 0.095741$. In addition to the desired value $w_2^* = 0.781796$, there is an unstable solution at $w_2 = 0.465602$. If the initial value of w_2 exceeds 0.465602, the buyer population will evolve towards the desired equilibrium $w_2^* = 0.781796$. However, if the initial value of w_2 is less than 0.465602, the buyer population will evolve towards the undesirable equilibrium solution $w_2 = 0$. Interestingly, as shall be illustrated in Section 5.2, the shopbot can cope with low initial values of w_2 by using a time-dependent pricing strategy to shift the system towards the desired equilibrium. Therefore, we shall proceed with an analysis that assumes that the desired equilibrium can be reached regardless of initial conditions.

According to our analysis thus far, $c_1 = c'$ and $c_2 = c' + \delta^*$. However, this violates Cond. 24 whenever $c' < \delta^*$. In this regime, the shopbot cannot charge as much as δ^* for the second quote because buyers would opt for the cheaper alternative search. Instead, to optimize its profits, the shopbot must match the cost of the alternative search method by setting $\delta = c'$ whenever $c' < \delta^*$. At equilibrium, this implies that $\beta(w_2) = c'$, or inversely, $w_2 = \beta^{-1}(c')$. Again, there are two possible solutions that maximize $w_2\beta(w_2)$. The favored solution is the one for which the value of w_2 is larger. Again, we assume that the shopbot can use time-dependent pricing to manipulate the buyer population toward the favored solution.

Next, consider the effect of Cond. 25. Substitution of Eqs. 10 and 15 into Cond. 25 yields the constraint $c_1 \leq \phi(w_2)$, where

$$\phi(w_2) \equiv 1 - \left[\frac{1-w_2}{2w_2}\right] \log\left(\frac{1+w_2}{1-w_2}\right). \tag{27}$$

This constraint comes into play when $c' > \phi(w_2^*) = 0.706946$.[8] In this regime, c' is so large that, with $c_1 = c'$ and $\delta = \delta^*$, the expected price plus the search cost exceeds every buyer's valuation. Consequently, all buyers opt out of the market, and the sellers and the shopbot earn no profits. To optimize its profit, the shopbot must therefore reduce δ below δ^*, which has the effect of increasing w_2 above w_2^*, thus lowering

[8] Thus, Conds. 24 and 25 are never violated simultaneously, and can be treated separately, as we have done.

the expected prices for all buyers. The shopbot's optimal strategy is to set δ to the value at which buyers are marginally interested in shopping, which occurs when Cond. 25 is satisfied with equality, *i.e.*, $c_1 = 1 - \bar{p}_1 = \phi(w_2)$. Thus, the optimal value of w_2 in this regime is determined by inverting $\phi(w_2) = c'$, which yields $w_2 = \phi^{-1}(c')$. To ensure that this value of w_2 is attained, the shopbot must set $\delta = \beta(w_2) = \beta(\phi^{-1}(c'))$.

Finally, note that since $w_2 \leq 1$, the results of the previous paragraphs hold only when $c' \leq 1$. If $c' > 1$, then the alternative search mechanism is not viable because a single price quote costs more than the buyer's valuation of the good. In this case, the shopbot is effectively a monopolist, and its optimal strategy is to set $c_1 = 1$ and $\delta = 0$. This price schedule encourages all buyers to be comparison shoppers ($w_2 = 1$), forcing sellers' prices to zero. Thus, sellers earn no profits, and buyers gain no surplus because the entire amount of their valuation is paid to the shopbot for the first price quote. The shopbot extracts all of the market surplus.

The analytic results for c_1, δ, and w_2 in the four relevant ranges of c' are summarized in Table II and illustrated in Fig. 7. Table III displays, over the four relevant ranges of c', three quantities of interest that can be computed from the values in Table II. The shopbot's profit π_{sh} is computed using Eq. 22. The total seller profits, $S\pi$, are computed using Eq. 5 with $\pi_m = 1$, which yields $w_1 = 1 - w_2$. Finally, the average buyer surplus σ is computed as follows:

$$\begin{aligned}
\sigma &= w_1(1 - (\bar{p}_1 + c_1)) + w_2(1 - (\bar{p}_2 + c_2)) \\
&= w_1(1 - \bar{p}_1) + w_2(1 - \bar{p}_2) - \pi_{\mathrm{sh}} \\
&= \phi(w_2) + w_2\beta(w_2) - \pi_{\mathrm{sh}} \\
&= w_2 - \pi_{\mathrm{sh}}
\end{aligned}$$

Table II. Optimal linear shopbot price schedule in terms of c_1 and δ, and distribution of the buyer population w_2, as a function of alternative search cost c'. Numerical values of the special constants appearing in the table are: $\delta^* = 0.095741$, $w_2^* = 0.781796$, and $\phi(w_2^*) = 0.706946$.

c'	c_1	δ	w_2
$0 \leq c' \leq \delta^*$	c'	c'	$\beta^{-1}(c')$
$\delta^* \leq c' \leq \phi(w_2^*)$	c'	δ^*	w_2^*
$\phi(w_2^*) \leq c' \leq 1$	c'	$\beta(\phi^{-1}(c'))$	$\phi^{-1}(c')$
$c' > 1$	1	0	1

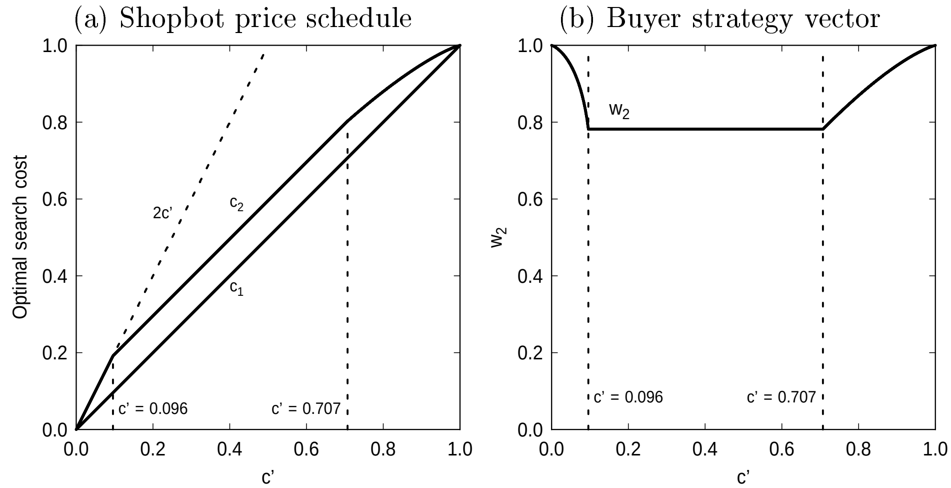

Figure 7. (a) Optimal linear shopbot price schedule as a function of alternative search cost c'. (b) Equilibrium values of buyer strategy vector $\vec{w} = (w_1, w_2)$.

The total social welfare is $\pi_{\rm sh} + S\pi + \sigma = \pi_{\rm sh} + 1 - w_2 + w_2 - \pi_{\rm sh} = 1$.[9] In other words, the market is efficient regardless of the value of c'.

Table III. Shopbot's profit $\pi_{\rm sh}$, total seller profits $S\pi$, and average buyer surplus σ.

c'	$\pi_{\rm sh}$	$S\pi$	σ
$0 \leq c' \leq \delta^*$	$c'(1+\beta^{-1}(c'))$	$1-\beta^{-1}(c')$	$\beta^{-1}(c')(1-c')-c'$
$\delta^* \leq c' \leq \phi(w_2^*)$	$c'+\delta^* w_2^*$	$1-w_2^*$	$\phi(w_2^*)-c'$
$\phi(w_2^*) \leq c' \leq 1$	$\phi^{-1}(c')$	$1-\phi^{-1}(c')$	0
$c' > 1$	1	0	0

Fig. 8(a) graphically displays the information contained in Table III. For comparison, Fig. 8(b) depicts the same quantities under a scenario in which no shopbot is present, and price information costs c' per quote. The behavior of the market in the absence of a shopbot has two regimes of interest. If c' is less than the maximal value of $\beta(w_2)$, namely 0.103872, then there are three equilibria. The first one is the trivial stable equilibrium at $w_2 = 0$. The two non-trivial equilibria are

[9] For non-normalized values of the buyer valuation v and the production cost r, these quantities can be rescaled to yield total social welfare of $v - r$. This simple rescaling assumes that all buyers have equal valuations v, which motivates the shopbot to manipulate the market to ensure that all buyers buy the good.

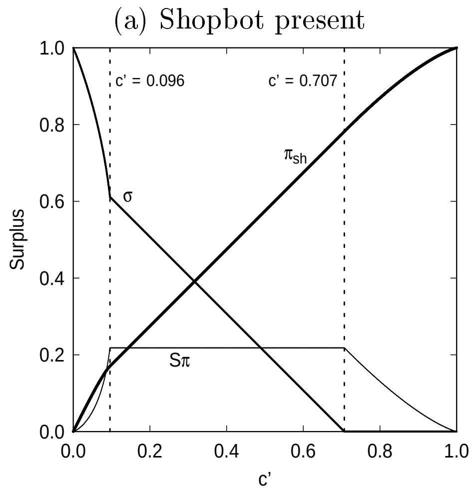

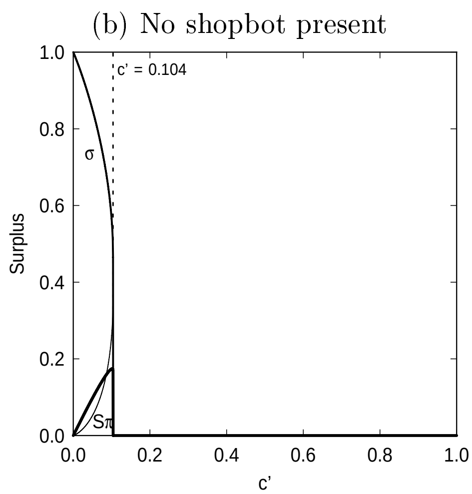

Figure 8. Average buyer surplus σ, total seller profits $S\pi$, and total search cost as a function of alternative search cost c'. (a) Shopbot present. Total search cost is equivalent to shopbot profit π_{sh}. (b) No shopbot present. Unlabeled curve that is roughly straight for small c' represents buyers' total expenditure on search.

the solutions to $\beta(w_2) = c'$. Of these, the one with the larger value of w_2 is stable, while the one with the smaller value of w_2 is unstable. If the initial value of w_2 exceeds that of the unstable non-trivial equilibrium solution, then the buyer population evolves towards the non-trivial stable equilibrium. Otherwise, the system evolves to the trivial stable equilibrium. Assuming that the buyer population stabilizes at the non-trivial equilibrium, the total seller profits and buyer surplus are given by precisely those formulas that appear in the first row of Table III, but the range of validity extends up to $c' = \max \beta(w_2) = 0.103872$ rather than $c' = \delta^* = 0.095741$. This equivalence is apparent for low values of c' in Figs. 8(a) and 8(b).

If $c' > 0.103872$, then there is no solution such that $\beta(w_2) = c'$. In this case, the system evolves to the trivial stable equilibrium at $w_2 = 0$. Since none of the buyers compare prices, sellers are free to behave as monopolists. If the sellers are unaware of (or do not take into account) the buyers' search cost c'—a tacit assumption in the foregoing analysis[10]—then they attempt to maximize their profits by charging the buyers exactly their valuations. But since $c' > 0$, the sum of the product and search costs would exceed the buyers' valuations. No goods would be bought or sold, resulting in zero surplus for the buyers, the sellers, and the shopbot.

In the presence of a shopbot (Fig. 8(a)), the market behavior is quite different. The shopbot deliberately sets the differential search cost to $\delta < c'$ in order to encourage buyers to search. This lowers prices, making it affordable for buyers to stay in the market. This helps both buyers and sellers. For $c' < 0.706946$, buyers receive a positive surplus, and sellers receive a positive profit. For $c' \geq 0.706946$, buyers receive zero surplus, but the sellers continue to eke out a profit that diminishes monotonically towards zero as $c' \rightarrow 1$.

5.2. Strategic Manipulation of Equilibria

In this section we address the issue that was deferred in Section 5.1: when there are multiple equilibria, how can the shopbot control which equilibrium is reached, regardless of initial conditions? The trick is to use a time-dependent pricing strategy to strategically manipulate the equilibria and their basins of attraction so as to guide the market towards the desired equilibrium.

Figure 9 illustrates such a scenario, in which it is assumed that $c' = 0.25$. Initially, the shopbot sets the "optimal" linear price schedule as computed from Table II: $c_1 = 0.25$ and $\delta = \delta^* = 0.095741$. The initial strategy vector is taken to be $\vec{w} = (0.95, 0.05, 0, 0, 0)$, reflecting a situation in which there is very little search on the part of the buyers. This initial state lies in the basin of attraction of the trivial (and unfavorable) equilibrium $w_2 = 0$. Throughout the simulation, the

[10] If the sellers were to take into account the effect of search costs upon buyers' search behavior, they could not legitimately regard $\vec{w}$ as fixed and independent of their own pricing decisions. This additional consideration would require an extensive overhaul of the game-theoretic analysis of Section 3.

market evolves according to Eq. 19, with $\eta = 0.004$. As the co-evolution of buyers' strategies and sellers' price distributions proceeds, sellers' prices rise towards 1. Any buyer who remains in the market pays nearly its valuation *plus* a search cost, and thus receives a negative surplus. Therefore, the buyers opt out of the market at their first opportunity, as is reflected in the growth of the w_0 component in Fig. 9b.

Recognizing that the buyer population is destined for an undesirable equilibrium in which the shopbot earns no profit, the shopbot drops its prices to zero at time $t = 1000$. Now in fact the shopbot earns no profit, but since buyers can obtain an unlimited number of quotes for free, the optimal strategy now shifts to $i = 5$. The buyers (who reevaluate their strategy once every 250 time steps on average) gradually shift to this optimal strategy, and hence w_5 approaches 1. This extensive search pins sellers' prices at low values, resulting in essentially no profits for sellers. On the other hand, the buyers experience both low prices and zero search costs, which causes their surplus to rise towards 1.

Finally, at time $t = 2000$, the shopbot reverts back to its original "optimal" linear price schedule. Now that most of the buyers are engaging in extensive search, the buyer population is in a state that lies within the favorable equilibrium's basin of attraction. But since the shopbot is once again charging for search, the buyers search less, ultimately stabilizing at the optimal state in which $w_2 = w_2^* = 0.781796$ and $w_1 = 1 - w_2^*$. In this state, the shopbot realizes its optimal profit of 0.324850, and the total seller profit and the buyer surplus assume the calculated values in Table II and Fig. 8a for $c' = 0.25$: 0.218204 and 0.456946, respectively. In this scenario, the self-interested manipulation of the market by the shopbot benefits the buyers and sellers, who otherwise would not have bought or sold any goods.

5.3. Heuristic Nonlinear Price Schedule

In the foregoing analysis, we have assumed that the shopbot adheres to a linear price schedule. In this section, we investigate how a shopbot might attempt to maximize its profits if freed from this constraint. Rather than attempting to compute the optimal nonlinear price schedule (which may be dependent upon initial conditions and time), we introduce an adaptive heuristic by which the shopbot dynamically sets a nonlinear price schedule.

One heuristic for dynamically setting the price schedule is to set $c_1 = c'$, and to set all remaining prices such that $\bar{p}_i + c_i$ is constant for all i. In other words, at time t, set the marginal price of quote i, namely $c_i - c_{i-1}$, equal to the marginal benefit of quote i, given the current vector $\vec{w}(t)$. This heuristic ensures that all search strategies

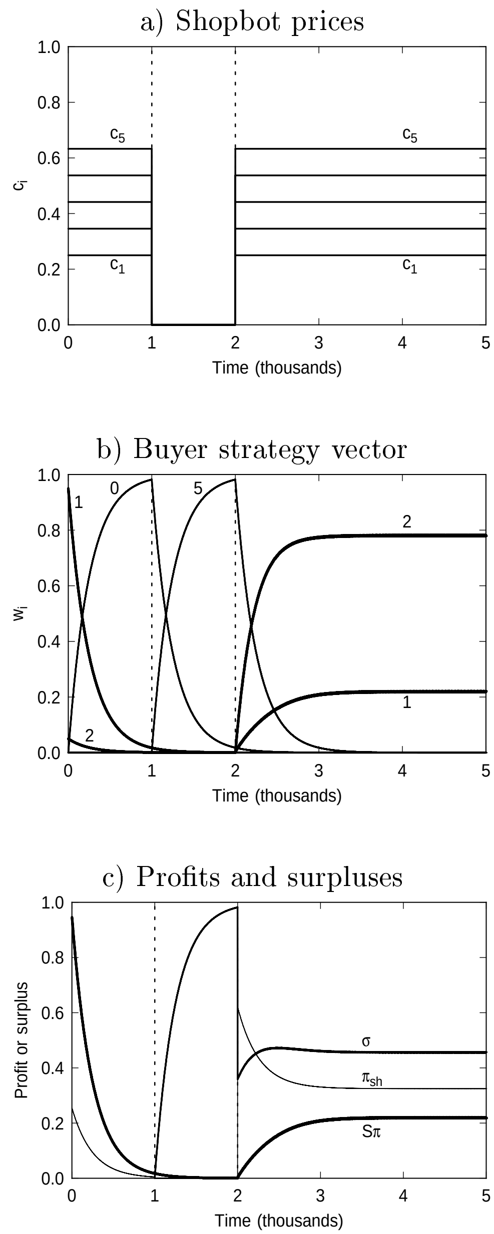

Figure 9. Example of how a shopbot can adjust its prices dynamically so as to shift the market from an unfavorable equilibrium to one that maximizes the shopbot's profit. a) Shopbot prices vs. time. b) Evolution of buyer search strategies vs. time, with components $w_i(t)$ labeled according to i. Asymptotic strategy vector is 0.218204 for search strategy 1, 0.781796 for strategy 2, and 0 for all other strategies. c) Shopbot profit, total seller profits, and average buyer surplus vs. time. Asymptotic values are $\pi_{\text{sh}} = 0.324850$, $S\pi = 0.218204$, and $\sigma = 0.456946$, respectively.

are equally desirable to buyers, and encourages $\vec{c}$ and $\vec{w}$ to reach an equilibrium. It is possible for this heuristic to be overly aggressive in forcing equilibrium, causing the shopbot to settle on a non-optimal price schedule, but in practice it typically yields reasonable solutions.

Figure 10 depicts this heuristic in action. Up until time $t = 3000$, the simulation is identical to Fig. 9. By this time, the buyer strategy $\vec{w}$ and the shopbot profit have nearly converged to equilibrium. In particular, the shopbot profit has nearly converged to the optimal value that can be attained by a linear price schedule, 0.324850. Then at time $t = 3000$, the heuristic described above is put into effect. On a time scale of a few thousand time steps, the shopbot prices $\vec{c}$ and the buyer strategies $\vec{w}$ co-evolve to an equilibrium in which the c_i are nonlinear in i. The nonlinearity is roughly of the form assumed in Eq. 18, with exponent $\alpha \approx 0.45 < 1$. Thus, the incremental cost of quote i decreases roughly as the square root of i.

The nonlinear price schedule that arises via this heuristic yields the shopbot an equilibrium profit of 0.417165—more than a 28% increase over what is attainable with a linear price schedule. The increase in profits comes about because now about half of the buyers are requesting 4 or 5 price quotes, and paying substantially more for them than they would have paid for 2 price quotes in the linear price schedule. As a side-effect of the shopbot's switch to a nonlinear price schedule, the total seller profit increases by over 14%, from 0.218204 to 0.249344. These increased profits come at the expense of the buyers, whose surplus drops by 27%, from 0.456946 to 0.333491.

In order to understand why the use of a nonlinear price schedule by the shopbot increases the total seller profits, note that, by Eq. 5 and the assumption that $\pi_m = 1$, the total seller profits $S\pi = w_1$. Since the nonlinear pricing heuristic results in a near-even four-way split among search strategies, the asymptotic value of w_1 is very close to 0.25. This is higher than the value $w_1^* = 1 - w_2^* = 0.218204$ that is attained with the optimal linear pricing structure, and the increase in seller profits follows directly from this increase in w_1. The increase in seller profits is not general, however. If the number of sellers increases appreciably beyond 5, then the even splitting of the non-zero components of $\vec{w}$ leads to a *decrease* in w_1, and therefore a decrease in total seller profits as well. We have observed this effect in simulations with 10 and 20 sellers.

The decrease in average buyer surplus after the introduction of the heuristic follows from the fact that the overall surplus in the market is a constant, since the shopbot adjusts its prices to ensure that all buyers purchase the good. Therefore, for $S = 5$, the increase in both the shopbot's and the sellers' profits must be offset by a decrease in the average buyer surplus. The buyer surplus appears to be relatively

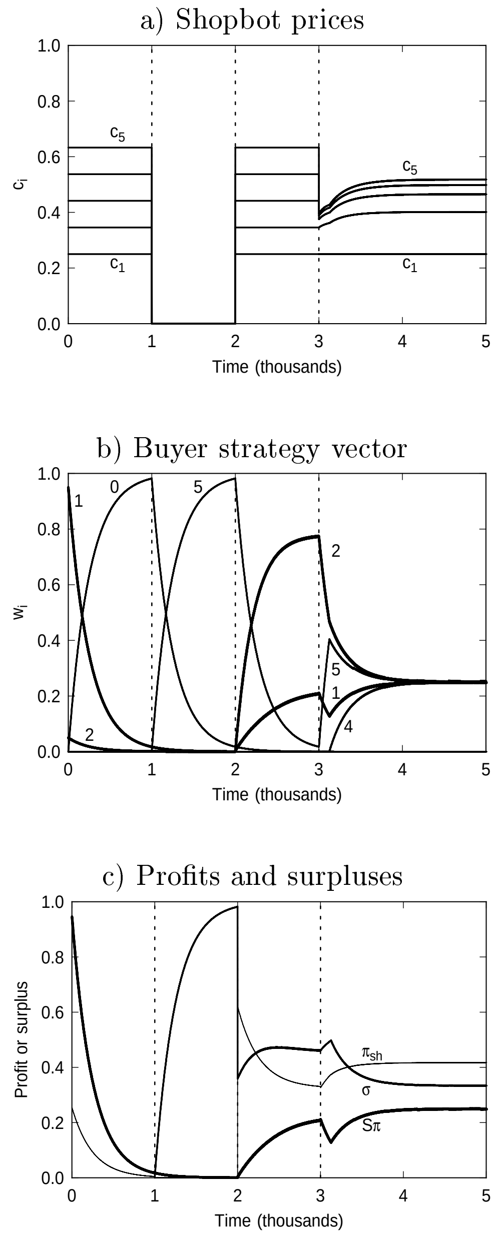

Figure 10. Scenario similar to that of Fig. 9 except that at times $t \geq 3000$ an adaptive pricing algorithm is employed by the shopbot. The resultant nonlinear price schedule yields higher profits. (a) Shopbot prices vs. time. Asymptotic price schedule is $\vec{c} = (0.250000, 0.401388, 0.465005, 0.498231, 0.517981)$. (b) Evolution of buyer search strategies vs. time, with components $w_i(t)$ labeled according to i. The asymptotic strategy vector is 0.249344 for search strategies 1, 2, and 4, 0.251969 for strategy 5, and 0 for strategies 0 and 3. (c) Shopbot profit, total seller profits, and buyer surplus vs. time. Asymptotic values are $\pi_{\text{sh}} = 0.417165$, $S\pi = 0.249344$, and $\sigma = 0.333491$, respectively.

insensitive to the number of sellers, never varying by more than 5% as S ranges from 5 to 20. As the number of sellers increases, the decrease in total seller profits results in an increase in the shopbot's profit.

The co-evolution of $\vec{c}(t)$ and $\vec{w}(t)$ after $t = 3000$ is somewhat complex. Initially, c_5 drops so low that search strategy 5 is consistently favored. After less than 200 time steps, however, the heuristic begins to equalize the strategies, until ultimately the equilibrium strategy vector is nearly evenly divided among search strategies 1, 2, 4, and 5, with no buyers pursuing search strategies 0 and 3. Since the heuristic is designed to create indifference among the various strategies, the buyers' avoidance of strategy 3 would appear to be a numerical artifact. However, we have consistently obtained the same result under several different experimental conditions, including different initial conditions, different durations of the period of free search, and different values of η. In addition, for related scenarios with more sellers, we typically find that some components of the equilibrium strategy vector tend to zero, and that often the remainder is divided evenly (or nearly so) among the non-zero components. Most likely, there is a subtle effect associated with the finite size of η that prevents exact indifference among the strategies.

In summary, a shopbot can increase its profit by using a well-tuned nonlinear price schedule. This price schedule can influence seller profit and buyer surplus in ways that are not intuitively obvious. Relative to the optimal linear price schedule, the nonlinear price schedule may result in increased or diminished seller profits, depending on the number of sellers, while the average buyer surplus is typically diminished. These results should be regarded as illustrative rather than general. A more general investigation would explore how nonlinear shopbot pricing affects the market under different pricing heuristics assuming different distributions of buyer valuations. Furthermore, as shopbots become more prevalent, one can expect the emergence of competition among shopbots. The ability of shopbots to manipulate the market will be greatly diminished when this occurs, and the resultant shopbot pricing strategies and market dynamics are likely to be radically different from what we have explored here for a monopolist shopbot.

6. Related Work

The study of the economics of information was launched in 1961 with the seminal work of Stigler [24]. Stigler cited several empirical examples of price dispersion, which he attributed to the costly search procedures faced by consumers. He stressed the utility of trade journals and organizations that specialize in the collection and dissemination of product information, such as *Consumer Reports*. Stigler also reminded us that, in medieval times, localized marketplaces thrived in spite of the heavy taxes levied on merchants, demonstrating how worthwhile it was for sellers to participate in such markets rather than undertake individual searches for buyers. Shopbots create analogous localized marketplaces; accordingly, we observe some sellers sponsoring shopbots and paying commissions on sales, essentially paying for the right to participate in the shopbot marketplace.

Since Stigler, many economists have developed and analyzed formal models that attempt to explain the phenomenon of price dispersion.[11] We have already mentioned the work of Varian [27] and Burdett and Judd [3]. The former study is a special case of our model of shopbot economics in the case of two types of buyers, namely type 1 and type S, existing in fixed proportions. The latter authors, on the other hand, consider independent buyer decisions among a set of search rules of fixed sample size. In this paper, we integrate these two approaches by specifying a fixed minimal proportion of buyers of type 1 while allowing all other buyers to choose their sample sizes. Moreover, we extend these ideas with the notion that shopbots themselves may well price for the information services they provide, acting as economic agents in their own right; this creates an additional strategic variable, namely the cost of search, as determined by shopbots attempting to maximize their profitability.

In the absence of shopbots, a model of sellers' price adjustment was studied in Diamond [7], in which a somewhat paradoxical outcome arises: for any positive search costs, no consumers search and all sellers charge the monopolistic price v. More recent work on the dynamics of price-setting includes the evolutionary approach of Hopkins and Seymour [15], where it is argued that the symmetric Nash equilibrium in mixed strategies is dynamically unstable. This raises concerns about the validity of assuming that $f(p)$ describes sellers' behavior. Since we also assume that sellers are represented by software agents, however, it is perhaps plausible to consider pricing in this manner, as $f(p)$ is computed according to a well-specified algorithm. Nonetheless, we are currently studying learning algorithms that evolve stably to *asymmetric* mixed strategy Nash equilibria in this model [13]; these may be more likely to arise as equilibrium outcomes in a dynamic economy of software agents.

[11] We mention only a handful of papers that make up this large body of literature, but refer the reader to the bibliography included in Hopkins and Seymour [15] for additional sources.

Bakos [1] has recently considered the potential impact of reduced buyer search costs on the electronic marketplace. His model allows for product differentiation, but it does not allow for a variety of buyer types. It remains to incorporate product differentiation into our model of buyer search. Baye and Morgan [2] consider the consequences of introducing into a commodity market a profit-seeking "gatekeeper", an intermediary that charges both buyers and sellers for price information. Although their assumptions about the market and the nature of the intermediary differ in detail from ours, they observe similar effects. In particular, a gatekeeper can wield enormous power over market dynamics, encouraging buyer search and improving social welfare. But to effectively maximize profits, gatekeepers have to contend with the existence of multiple equilibria.

7. Conclusions and Future Work

In our explorations of the potential economic impact of shopbots on electronic markets, we have proposed a model that is similar to several that have been investigated by economists interested in understanding the phenomenon of price dispersion. However, in contrast to these previous studies, our goals are *prescriptive* rather than *descriptive*. We are interested in designing economically-motivated software agents, and in understanding the market behavior that is likely to emerge from their collective interactions. Our different orientation is reflected in many ways, including our emphasis on the constructive computation of price distributions (as opposed to theorems and lemmas regarding their basic properties), and our extensive study of non-equilibrium behavior (as opposed to proofs of the existence of equilibria).

Arguing that nonlinear search cost schedules are likely to arise naturally, or could 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 take advantage of search mechanisms, can lead to similar outcomes. Finally, the evolutionary dynamics of buyer strategies can be complicated, possibly settling into a complex limit cycle. Even when buyer strategies stabilize, there may be multiple equilibria, with the choice of equilibrium governed by the initial state of the buyer strategies, among other factors.

Placing ourselves in the role of shopbot designers, we explored the strategic pricing of price information. Through modeling, analysis, and simulation, we validated one of our earlier assumptions by showing that nonlinear cost schedules come about as the natural consequence of endowing shopbots with economic incentives. Even in the face of competition from an alternative search mechanism, shopbots can wield a good deal of control over markets; specifically, they can manipulate search costs so as to extract an appreciable fraction of the market surplus. In addition, we showed how buyers can benefit from the shopbot's self-interested price manipulation. The power of an individual shopbot would of course be diminished considerably if there were competition among two or more shopbots. A study of coupled markets in which shopbots compete as providers of price and product information services to buyers or buyer agents would be fascinating.

In closing, we mention two important future directions of our work. First, it is necessary to develop a more faithful representation of buyer behavior that explicitly takes both product and vendor attributes into account. By conducting an empirical study of the behavior of buyers who use the shopbot provided at `DealTime.com` (known as `evenbetter` at the time of their study), Brynjolfsson and Smith [4] have found that brand and consumer loyalty can lead buyers to select higher-priced vendors. Shopbots like `mysimon.com` are already starting to provide additional information about product attributes, and moreover, they have experimented with eliciting utility functions from buyers that can be used to compute personalized rankings of product offerings. It would be of both theoretical and practical interest to analyze and simulate a model that accounts for these various aspects of buyer behavior, perhaps along the lines of Bakos' analysis of horizontal differentiation in electronic markets [1] or the analysis of vertical differentiation of Sairamesh *et al.* [21]. In addition to studying the effect of more realistic buyer behavior, it would be desirable to introduce more realistic models of pricing behavior. Previous studies of asynchronous [14] and adaptive [13] pricing in markets with fixed-strategy buyers have revealed interesting nonequilibrium behavior, which may well become even more complex in nature when the buyers are adaptive.

Acknowledgments

The authors gratefully acknowledge the support of IBM's Institute for Advanced Commerce.

Appendix

In this appendix, it is shown that no pure strategy Nash equilibria exist in our model of shopbot economics whenever (i) prices are chosen from a sufficiently large (or continuous) set, and (ii) $0 < w_1 < 1$.[12] If $w_1 = 1$, then the unique Nash equilibrium is such that all sellers charge the monopoly price p_m. If $w_1 = 0$, then at equilibrium at least two sellers charge the competitive price r and all sellers earn zero profits (see [10, 11]). Our proof proceeds by contradiction: we assume the existence of a pure strategy Nash equilibrium, derive the unique form of such an equilibrium, and then argue that a strategy profile of this form is not individually optimizing for all sellers.

Assuming the existence of a pure strategy Nash equilibrium, we rederive Eq. 3 of Section 3, which describes the probability $h_s(\vec{p}, \vec{w})$ that buyers choose seller s. As in Section 3,

$$h_s(\vec{p}, \vec{w}) = \sum_{i=0}^{S} w_i h_{s,i}(\vec{p}) \tag{28}$$

where $h_{s,i}(\vec{p})$ denotes the demand for seller s by buyers of type i. Since we are now assuming pure rather than mixed strategies, it is convenient to define the following functions in deriving $h_{s,i}(\vec{p})$, a quantity whose probabilistic analog was previously expressed in terms of cumulative distribution functions:

- $\mu_s(\vec{p})$ is the number of sellers charging a higher price than s;
- $\tau_s(\vec{p})$ is the number of sellers charging the same price as s, excluding s itself;
- $\lambda_s(\vec{p})$ is the number of sellers charging a lower price than s.

The quantity $h_{s,i}(\vec{p})$ is is the product of the probability x_s that s is among i potential sellers selected at random, and the conditional probability y_s that s is either the lowest-priced seller or the lucky one in a set of lowest-priced sellers, given a set of i potential sellers that includes s. The probability x_s that s is one of a set of i potential sellers selected at random is simply $x_s = \binom{S-1}{i-1} / \binom{S}{i}$.

The conditional probability y_s is a function of $\lambda_s(\vec{p}), \tau_s(\vec{p})$, and $\mu_s(\vec{p})$. For convenience, fix seller s, and abbreviate $\nu \equiv \nu_s(\vec{p})$, for $\nu \in \{\lambda, \mu, \tau\}$. Given τ sellers charging the same price as seller s, the probability y_s

[12] The case of small, discrete price sets is considered for a related shopbot model in Appendix A of Greenwald and Kephart [12]. In some cases, it is possible to obtain pure strategy Nash equilibria.

that s is one of the lowest-priced sellers among i potential sellers, and moreover, that s is randomly selected from among $0 \leq t \leq \tau$ such lowest-priced sellers, is given by:

$$y_s(\lambda, \tau, \mu) = \sum_{t=0}^{\tau} \frac{1}{t+1} x_s(\lambda, \tau, \mu, t) \tag{29}$$

where $x_s(\cdot, t)$ denotes the probability that s is one of the t lowest-priced sellers of i potential sellers. The quantity $x_s(\cdot, t)$ is computed by determining the number of ways in which to arrange $i-1$ of the remaining $S-1$ sellers, exclusive of seller s, such that t sellers charge the same price as s, but no seller charges a price less than s, divided by the total number of ways in which to choose $i-1$ other potential sellers from the remaining $S-1$: *i.e.*,

$$x_s(\lambda, \tau, \mu, t) = \frac{\binom{\lambda}{0}\binom{\tau}{t}\binom{\mu}{i-1-t}}{\binom{S-1}{i-1}}. \tag{30}$$

Combining the expressions for x_s, y_s, and z_s and simplifying, we arrive at the following:

$$h_{s,i}(\vec{p}) = \binom{S}{i}^{-1} \frac{1}{\tau+1} \sum_{t=0}^{\tau} \binom{\tau+1}{t+1} \binom{\mu}{i-(t+1)}. \tag{31}$$

Finally, as in Section 3, let $\rho B = 1$ and $g(p) = \Theta(v-p)$, which yields the sellers' profit function $\pi = (p-r)h_s(\vec{p}, \vec{w})$ with $h_s(\vec{p}, \vec{w})$ defined by Eqs. 28 and 31.

We now proceed to derive the unique form of a possible pure strategy Nash equilibrium. Suppose that the sellers are ordered $s_1, \ldots, s_j, \ldots, s_S$ such that the indices $i < j$ whenever equilibrium prices $p_i^* \leq p_j^*$. First, note that equilibrium prices $p_j^* \in (r, v]$, since $p_j < r$ yields strictly negative profits, but $p_j = r$ and $p_j > v$ yield zero profits, and $p_j \in (r, v]$ yields strictly positive profits. The following observation describes the form of a pure strategy Nash equilibrium whenever $0 < w_1 < 1$: *at equilibrium, no two sellers charge identical prices.*

Suppose that, on the contrary, two distinct sellers offer equivalent prices: *i.e.*, $r < p_1^* < \ldots < p_j^* = p_{j+1}^* < \ldots < p_S^* \leq v$. In this case, seller j stands to gain by undercutting seller $j+1$ by ϵ, implying that p_j^* is not in fact an equilibrium price. In particular,

$$\begin{aligned} \pi_j(p_j^* - \epsilon, p_{-j}^*) &= \textstyle\sum_{i=0}^{S} w_i \binom{S}{i}^{-1} \binom{\mu+1}{i-1} (p_j - \epsilon - r) \\ &> \textstyle\sum_{i=0}^{S} w_i \binom{S}{i}^{-1} \left[\binom{\mu}{i-1} + \frac{1}{2}\binom{\mu}{i-2} \right] (p_j - r) \\ &= \pi_j(p_j^*, p_{-j}^*) \end{aligned}$$

for sufficiently small values of ϵ. The price quantum ϵ is sufficiently small precisely when the set of possible prices is sufficiently large: the set of prices can be continuous or it can be discretized such that $\epsilon = (v-r)/(|P|-1)$, where $|P|$ is the cardinality of the price set.

Further, we also observe that seller S charges price v at equilibrium, since $\pi_S(v, p^*_{-S}) = \frac{1}{S} w_A(v-r) > \frac{1}{S} w_A(p_S - r) = \pi_S(p_S, p^*_{-S})$, for all $p^*_{S-1} < p_S < v$. Therefore, the relevant price vector consists of S distinct prices ordered such that $r < p^*_1 < \ldots < p^*_{S-1} < p^*_S = v$.

The price vector $(p^*_1, \ldots, p^*_j, \ldots, p^*_S)$, however, is not in fact a Nash equilibrium. It is the case that $p^*_S = v$ is an optimal response to p^*_{-S}, since the profits of seller S are maximized at v given that there exists lower priced sellers $1, \ldots, S-1$. However, price $p^*_j < p_S = v$ is *not* an optimal response to p^*_{-j}. On the contrary, for all sellers $1 \leq j \leq S-1$, seller j has an incentive to deviate, since whenever $\epsilon < p_{j+1} - p_j$,

$$\begin{aligned}
\pi_j(p^*_{j+1} - \epsilon, p^*_{-j}) &= \textstyle\sum_{i=0}^{S} w_i \binom{S}{i}^{-1} \binom{\mu}{i-1} (p_{j+1} - \epsilon - r) \\
&> \textstyle\sum_{i=0}^{S} w_i \binom{S}{i}^{-1} \binom{\mu}{i-1} (p_j - r) \\
&= \pi_j(p^*_j, p^*_{-j}).
\end{aligned}$$

It follows that there is no pure strategy Nash equilibrium in our model of shopbot economics, assuming sellers choose prices from a large (or continuous) set and $0 < w_1 < 1$.

References

1. Bakos, J. Y.: 1997, 'Reducing Buyer Search Costs: Implications for Electronic Marketplaces'. *Management Science* **43**(12).
2. Baye, M. R. and J. Morgan: 2001, 'Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets'. *American Economic Review* **91**, 454–474.
3. Burdett, K. and K. L. Judd: 1983, 'Equilibrium price dispersion'. *Econometrica* **51**(4), 955–969.
4. Bynjolfsson, E. and M. D. Smith: 2000, 'The Great Equalizer? Consumer Choice Behavior at Internet Shopbots'. Working paper. Draft available on World Wide Web at http://ecommerce.mit.edu/erik.
5. Chavez, A. and P. Maes: 1996, '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.
6. DeLong, J. B. and A. M. Froomkin: 1998, 'The Next Economy?'. In: D. Hurley, B. Kahin, and H. Varian (eds.): *Internet Publishing and Beyond: The Economics of Digital Information and Intellecutal Property.* Cambridge, Massachusetts: MIT Press.
7. Diamond, P.: 1971, 'A Model of Price Adjustment'. *Economic Theory* **3**, 156–168.

8. Eriksson, J., N. Finne, and S. Janson: 1996, 'Information and interaction in MarketSpace — towards an open agent-based market infrastructure'. In: *Proceedings of the Second USENIX Workshop on Electronic Commerce*.
9. Green, H.: 1998, 'Good-bye to Fixed Pricing?'. *Business Week* pp. 71–84.
10. Greenwald, A.: 1999, 'Learning to Play Network Games'. Ph.D. thesis, Courant Institute of Mathematical Sciences, New York University, New York.
11. Greenwald, A. and J. Kephart: 1999, 'Shopbots and Pricebots'. In: *Proceedings of Sixteenth International Joint Conference on Artificial Intelligence*, Vol. 1. pp. 506–511.
12. Greenwald, A. and J. Kephart: 2000, 'Shopbots and Pricebots'. In: A. Moukas, C. Sierra, and F. Ygge (eds.): *Lecture Notes on Artificial Intelligence: Proceedings of the IJCAI Workshop on Agent-mediated Electronic Commerce*. Springer-Verlag.
13. Greenwald, A. and J. Kephart: 2001, 'Probabilistic Pricebots'. In: *Proceedings of the Fifth International Conference on Autonomous Agents*. pp. 560–567.
14. Greenwald, A., J. Kephart, and G. Tesauro: 1999, 'Strategic Pricebot Dynamics'. In: *Proceedings of First ACM Conference on Electronic Commerce*. pp. 58–67.
15. Hopkins, E. and R. M. Seymour: 1996, 'Price Dispersion: An Evolutionary Approach'. Unpublished Manuscript.
16. Kephart, J. O., J. E. Hanson, D. W. Levine, B. N. Grosof, J. Sairamesh, R. B. Segal, and S. R. White: 1998, 'Dynamics of an Information Filtering Economy'. In: *Proceedings of the Second International Workshop on Cooperative Information Agents*.
17. Krulwich, B.: 1996, 'The BargainFinder agent: Comparison price shopping on the Internet'. In: J. Williams (ed.): *Agents, Bots and Other Internet Beasties*. SAMS.NET publishing (MacMillan), pp. 257–263. URLs: http://bf.cstar.ac.com/bf, http://www.geocities.com/ResearchTriangle/9430.
18. Lewis, T. G.: 1997, *The Friction-Free Economy: Marketing Strategies for a Wired World*. New York: HarperCollins.
19. Markopoulos, P. and L. Ungar: 2001, 'Pricing Price Information in E-Commerce'. In: *Proceedings of Third ACM Conference on Electronic Commerce*.
20. Nash, J.: 1951, 'Non-cooperative Games'. *Annals of Mathematics* **54**, 286–295.
21. Sairamesh, J. and J. O. Kephart: 2000, 'Price Dynamics of Vertically Differentiated Information Markets'. *Decision Support Systems* **28**(1,2).
22. Salop, S. and J. Stiglitz: 1977, 'Bargains and rip-offs: a model of monopolistically competitive price dispersion'. *Review of Economic Studies* **44**, 493–510.
23. Schuster, H. G.: 1988, *Deterministic Chaos: An Introduction*. New York: VCH.
24. Stigler, G. J.: 1961, 'The Economics of Information'. *Journal of Political Economy* **69**(3), 213–225.
25. Tirole, J.: 1988, *The Theory of Industrial Organization*. Cambridge, MA: The MIT Press.
26. Tsvetovatyy, M., M. Gini, B. Mobasher, and Z. Wieckowski: 1997, 'MAGMA: an agent-based virtual market for electronic commerce'. *Applied Artificial Intelligence*.
27. Varian, H.: 1980, 'A model of sales'. *American Economic Review, Papers and Proceedings* **70**(4), 651–659.