To be published in proceedings of IJCAI'99.

Shopbots and Pricebots

Amy R. Greenwald and Jeffrey O. Kephart
IBM Institute for Advanced Commerce
IBM Thomas J. Watson Research Center
Yorktown Heights, NY 10598
amygreen@cs.nyu.edu,kephart@watson.ibm.com

Abstract

Shopbots are agents that automatically search the Internet to obtain information about prices and other attributes of goods and services. They herald a future in which autonomous agents profoundly influence electronic markets. In this study, a simple economic model is proposed and analyzed, which is intended to quantify some of the likely impacts of a proliferation of shopbots and other economically-motivated software agents. In addition, this paper reports on simulations of *pricebots* — adaptive, price-setting agents which firms may well implement to combat, or even take advantage of, the growing community of shopbots. This study forms part of a larger research program that aims to provide insights into the impact of agent technology on the nascent information economy.

1 Introduction

Shopbots, agents that automatically search the Internet for goods and/or services on behalf of consumers, herald a future in which autonomous agents become an essential component of nearly every facet of electronic commerce [Chavez and Maes, 1996; Kephart *et al.*, 1998; Tsvetovatyy *et al.*, 1997]. In response to a consumer's expressed interest in a specified good or service, 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.junglee.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 [Lewis, 1997; DeLong and Froomkin, 1998]. Transportation costs, menu costs — the costs to firms of evaluating, updating, and advertising prices — and shopping 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 and related types of autonomous agents in electronic marketplaces, and how might species of computational agents evolve?

DeLong and Froomkin [1998] qualitatively investigate the ongoing emergence of shopbots; in particular, they note that short of violating anti-trust laws, firms will be hard pressed to prevent their competitors from sponsoring shopbots, in which case those who do not do so will experience decreased sales. In this paper, we utilize quantitative techniques to address the aforementioned questions. We propose, analyze, and simulate a simple economic model designed to capture the present role of shopbots as agents of economic change, particularly with regard to consumer preferences, as they decrease the cost of obtaining information in markets known to exhibit price dispersion. Looking ahead several years into the future, we project that shopbots will evolve into economic entities (*i.e.*, utility maximizers) in their own right, interacting with billions of other self-interested software agents. Moreover, we predict the emergence of *pricebots* — economically-motivated agents that set prices so as to maximize the profits of firms, just as shopbots seek prices that minimize costs for consumers. Accordingly, we study adaptive price-setting algorithms which pricebots might utilize to combat the growing community of shopbots, in a full-fledged agent-based economy.

This paper is organized as follows. The next section, Section 2, presents our model, which is analyzed in Section 3 from a game-theoretic point of view. Section 4 describes various adaptive price-setting algorithms and the results of their simulation under the prescribed model. A possible evolution of shopbots and pricebots is discussed Section 5. Concluding remarks and ideas for future work appear in Section 6.

2 Model

We consider an economy in which there is a commodity that is offered for sale by S sellers and of interest to B buyers, with $B \gg S$. Each buyer b generates purchase orders at random times, with rate ρ_b, while each seller s resets its price p_s at random times, with rate ρ_s. The value of the good to buyer b is v_b; the cost of production for seller s is c_s.

A buyer b's utility for a good is a function of price:

$$u_b(p) = \begin{cases} v_b - p & \text{if } p \leq v_b \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

This states that a buyer purchases a good from a given seller if and only if the seller's price is less than the buyer's valuation of the good; if price equals valuation, we make the behavioral assumption that a transaction occurs. We do not assume that buyers are utility maximizers; instead we assume that they consider the prices offered by sellers using one of the following strategies:

1. *Any Seller:* buyer selects seller at random, and purchases the good if the price charged by that seller is less than the buyer's valuation.
2. *Bargain Hunter:* buyer checks the offer price of all sellers, determines the seller with the lowest price, and purchases the good if that lowest price is less than the buyer's valuation. (This type of buyer corresponds to those who take advantage of shopbots.)

The buyer population consists of a mixture of buyers employing one of these strategies, with a fraction w_A using the *Any Seller* strategy and a fraction w_B using the *Bargain Hunter* strategy; $w_A + w_B = 1$. Buyers employing these respective strategies are referred to as type A and type B buyers.

A seller s's expected profit per unit time π_s is a function of the price vector $\vec{p}$:

$$\pi_s(\vec{p}) = (p_s - c_s) D_s(\vec{p}) \tag{2}$$

where $D_s(\vec{p})$ is the rate of demand for the good produced by seller s. This rate of demand is the product of the overall buyer rate of demand $\rho = \sum_b \rho_b$, the likelihood of a given buyer selecting seller s as their potential seller, $h_s(\vec{p})$, and the fraction of buyers whose valuations satisfy $v_b \geq p_s$, denoted $g(p_s)$:

$$D_s(\vec{p}) = \rho B h_s(\vec{p}) g(p_s). \tag{3}$$

Note that $g(p_s) = \int_{p_s}^{\infty} \gamma(x) dx$, where $\gamma(x)$ is the probability density function describing the likelihood that a given buyer has valuation x. If $v_b = v$ for all buyers b, then $\gamma(x)$ is the Dirac delta function $\delta(v - x)$, and the integral yields a step function $g(p_s) = \Theta(v - p_s)$:

$$\Theta(v - p_s) = \begin{cases} 1 & \text{if } p_s \leq v \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

Without loss of generality, we define the time scale such that $\rho B = 1$. It follows that $D_s(\vec{p}) = h_s(\vec{p}) g(p_s)$, and π_s is seller s's expected profit per unit sold systemwide.

The probability $h_s(\vec{p})$ that buyers select seller s as their potential seller depends on the distribution of the buyer population, namely (w_A, w_B). In particular,

$$h_s(\vec{p}) = w_A f_{s,A}(\vec{p}) + w_B f_{s,B}(\vec{p}) \tag{5}$$

where $f_{s,A}(\vec{p})$ and $f_{s,B}(\vec{p})$ are the probabilities that seller s is selected by buyers of type A and B, respectively. The probability that a buyer of type A select a seller s is independent of the ordering of sellers' prices; in particular, $f_{s,A}(\vec{p}) = 1/S$. Buyers of type B, however, select a seller s if and only if s is one of the lowest price sellers. Given that the buyers' strategies depend on the relative ordering of the sellers' prices, it is convenient to define the following functions:

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

Now buyers of type b select seller s iff s is *s.t.* $\lambda_s(\vec{p}) = 0$, in which case a buyer selects a particular such seller s with probability $1/(\tau_s(\vec{p}) + 1)$. Therefore,

$$f_{s,B}(\vec{p}) = \frac{1}{\tau_s(\vec{p}) + 1} \delta_{\lambda_s(\vec{p}),0} \tag{6}$$

where $\delta_{i,j}$ is the Kronecker delta function, equal to 1, whenever $i = j$, and 0, otherwise.

The preceding results can be assembled to express the profit function π_s for seller s in terms of the distribution of strategies and valuations within the buyer population. In particular, assuming (as we do from here forward) that all buyers share the same valuation v, and all sellers share the same cost c, then

$$\pi_s(\vec{p}) = \begin{cases} (p_s - c) h_s(\vec{p}) & \text{if } p_s \leq v \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

where

$$h_s(\vec{p}) = w_A \frac{1}{S} + w_B \frac{1}{\tau_s(\vec{p}) + 1} \delta_{\lambda_s(\vec{p}),0} \tag{8}$$

3 Analysis

In this section, we perform a game-theoretic analysis assuming sellers are profit maximizers. In particular, we first show that there is no pure strategy Nash equilibrium, and we then compute and describe the symmetric mixed strategy Nash equilibrium. Recall that $B \gg S$; in particular, the number of buyers is assumed to be very large, while the number of sellers is a great deal smaller. In accordance with this assumption, it is reasonable to consider the strategic decision-making of the sellers alone, since their relatively small number suggests that the behavior of individual sellers indeed influences market dynamics, while the large number of buyers renders the effects of individual buyers' actions negligible. A Nash equilibrium is a vector of prices $\vec{p}^*$ at which sellers maximize their individual profits and from which they have no incentive to deviate [Nash, 1951]. Throughout this exposition, we adopt the notation $\vec{p} = (p_s, p_{-s})$, which distinguishes the price offered by seller s from the prices offered by the remaining sellers.

Traditional economic models consider the case in which all buyers are bargain hunters: *i.e.*, $w_B = 1$. In this case, prices are driven down to marginal cost; in particular, $p_s^* = c$, for all sellers s (see, for example, Tirole [1988]). In contrast, consider the case in which all buyers are of type A, meaning that they randomly select a potential seller: *i.e.*, $w_A = 1$. In this situation, tacit collusion arises, in which all sellers charge the monopolistic price, in the absence of explicit coordination; in particular, $p_s^* = v$, for all sellers s. Of particular interest in this study, however, is the dynamics of interaction among buyers of various types: *i.e.*, $0 < w_A, w_B < 1$.

We begin our analysis with the following observation: *at equilibrium, at most one seller s charges $p_s^* < v$.* Suppose that two distinct sellers $s' \neq s$ set their equilibrium prices to be $p_{s'}^* = p_s^* < v$, while all other sellers set their equilibrium prices at the buyers' valuation v. In this case, $\pi_s(p_s^* - \epsilon, p_{-s}^*) = [(1/S)w_A + w_B](p_s^* - \epsilon - c) > [(1/S)w_A + (1/2)w_B](p_s^* - c) = \pi_s(p_s^*, p_{-s}^*)$, for small values of ϵ, whenever $w_B > 0$, which implies that p_s^* is not an equilibrium price for seller s. Now suppose that two distinct sellers $s' \neq s$ set their equilibrium prices to be $p_{s'}^* < p_s^* < v$, while all other sellers set their equilibrium prices at v. In this case, seller s prefers price v to p_s^*, since $\pi_s(v, p_{-s}^*) = [(1/S)w_A](v - c) > [(1/S)w_A](p_s^* - c) = \pi_s(p_s^*, p_{-s}^*)$, which implies that p_s^* is not an equilibrium price for seller s. Therefore, at most one seller charges $p_s^* < v$.

On the other hand, *at equilibrium, at least one seller s charges $p_s^* < v$.* Given that all sellers other than s set their equilibrium prices at v, seller s maximizes its profits by charging price $v - \epsilon$, since $\pi_s(v-\epsilon, p_{-s}^*) = [(1/S)w_A + w_B](v - \epsilon - c) > [(1/S)(w_A + w_B)](v - c) = \pi_s(v, p_{-s}^*)$, for small values of ϵ, whenever $w_B > 0$. Thus v is not an equilibrium price for seller s. It follows from these two observations that at equilibrium, exactly one seller s sets its price below the buyers' valuation v, while all other sellers $s' \neq s$ set their equilibrium prices $p_{s'}^* \geq v$. Note, however, that $\pi_{s'}(v, p_{-s'}^*) = [(1/S)w_A](v - c) > 0 = \pi_{s'}(v', p_{-s'}^*)$, for all $v' > v$, if $w_A > 0$, implying that all other sellers s' maximize their profits by charging price v. Thus, the unique form of pure strategy equilibrium which arises in this setting requires that a single seller s set its price $p_s^* < v$ while all other sellers $s' \neq s$ set their prices $p_{s'}^* = v$. The price vector (p_s^*, p_{-s}^*), with $p_{-s}^* = (v, \ldots, v)$, however, is not a Nash equilibrium. While v is in fact an optimal response to p_s^*, since the profits of seller $s' \neq s$ are maximized at v given that there exists low-priced seller s, p_s^* is not an optimal response to v. On the contrary, $\pi_s(p_s^*, v, \ldots, v) < \pi_s(p_s^* + \epsilon, v, \ldots, v)$. In particular, the low-priced seller s has incentive to deviate. It follows that there is no pure strategy Nash equilibrium in the proposed model of shopbots.

There does, however, exist a symmetric mixed strategy Nash equilibrium. Let $f(p)$ denote the density function according to which sellers set their prices, and let $F(p)$ be the corresponding cumulative distribution function. [1] The event that seller s is the low-priced seller occurs with probability $[1 - F(p)]^{S-1}$. Substituting this into Eq. 5, we obtain the demand expected by seller s:

$$h_s(p) = w_A \frac{1}{S} + w_B[1 - F(p)]^{S-1} \tag{9}$$

The precise value of $F(p)$ is determined by noting that at equilibrium expected profits are equal for all sellers, and moreover the expected profit level is given by the guaranteed minimum achieved at price v, namely $(1/S)w_A(v - c)$. Now, by setting $\pi_s(p) = h_s(p)(p - c)$ equal to this value and solving for $F(p)$, we obtain:

$$F(p) = 1 - \left[\left(\frac{w_A}{w_B S}\right)\left(\frac{v - p}{p - c}\right)\right]^{\frac{1}{S-1}} \tag{10}$$

Notice that $F(p) = 0$ for $p = p^*$ defined as follows:

$$p^* = c + \frac{w_A(v - c)}{w_A + w_B S} \tag{11}$$

and $F(p) = 1$ for $p = v$. Thus, Eq. 10 is valid only in the range $p^* \leq p \leq v$. [2]

The functions $F(p)$ and $f(p)$ are plotted in Figure 1. When w_B exceeds a critical threshold $w_B^{\text{crit}} = \frac{S-2}{S^2+S-2}$ (equal to 0.1071 for $S = 5$), $f(p)$ is bimodal. In this regime, as either w_B or S increases, the probability density concentrates either just below v, where sellers expect high margins but low volume, or just above p^*, where they expect low margins but high volume; moreover, the latter solution becomes increasingly probable. Since p^* itself decreases under these conditions (see Eq. 11), it follows that both the average price paid by buyers and the average profit earned by sellers decrease. These relationships have a simple interpretation: buyers' use of shopbots catalyzes competition among sellers, and moreover, smaller fractions of shopbot users induce competition among larger numbers of sellers.

4 Simulations

When sufficiently widespread adoption of shopbots by buyers forces sellers to become more competitive, it seems likely that sellers will respond by creating *pricebots* that automatically set prices so as to maximize profitability. It is unrealistic, however, to expect that pricebots will simply compute the mixed strategy Nash equilibrium and distribute their prices accordingly. The real business world is fraught with uncertainties that undermine the validity of traditional game-theoretic analyses: sellers lack perfect knowledge of buyer demands, and have an incomplete understanding of competitors' strategies. In order to be profitable, pricebots will need to continually adapt to changing market conditions.

[1] As the equilibrium is symmetric, we abbreviate $p \equiv p_s$ and we suppress dependence on $\vec{p}$.

[2] A similar derivation of the mixed strategy equilibrium appears in Varian [1980].

In this section, we discuss simulations of two adaptive pricing strategies, and we compare the resulting price and profit dynamics with the game-theoretic equilibrium. Recently, empirical studies of sophisticated learning algorithms have revealed that learning tends to converge to pure strategy Nash equilibria in games for which such equilibria exist [Greenwald *et al.*, 1998]. As there does not exist a pure strategy Nash equilibrium in the shopbot model, it is of particular interest to study the outcome of adaptive pricing schemes.

4.1 Pricing Strategies

We consider three pricing strategies, each of which makes very different demands on the required level of informational and computational power of agents:

GT The *game-theoretic* strategy is designed to reproduce the mixed strategy game-theoretic equilibrium computed in the previous section, provided that it is employed by all seller agents. It makes use of full information about the buyer population, and assumes that its competitors also use the GT strategy. It therefore generates a price chosen randomly from the probability density function derived in the previous section.

MY The myopically optimal, or *myoptimal*, [3] pricing strategy [Kephart *et al.*, 1998] uses information about all the buyer characteristics that factor into the buyer demand function, as well as the competitors' prices, but makes no attempt to account for competitors' pricing strategies. Instead, it is based on the assumption of static expectations: even if one seller is contemplating a price change under myoptimal pricing, this seller does not assume that this will elicit a response from its competitors; instead it assumes that competitors' prices will remain fixed.

The myoptimal seller uses all of the available information and the assumption of static expectations to perform an exhaustive search for the price p^* that maximizes its expected profit π. In our simulations, we compute π according to Eqs. 7 and 8. The optimal price p^* is guaranteed to be either the valuation v or ϵ below some competitor's price, where ϵ is the *price quantum*, or the smallest amount by which one seller may undercut another, set to 0.002 in these simulations. This limits the search for p^* to S possible values.

DF The *derivative-following* strategy is far less informationally intensive than either the myoptimal pricing strategy or the game-theoretic strategy. In particular, this strategy can be used in the absence of any knowledge or assumptions about one's competitors or the buyer demand function. A derivative follower simply experiments with incremental increases (or decreases) in its price, continuing to move its price in the same direction until the observed profitability level falls, at which point the direction of movement is reversed. The price increment δ is chosen randomly from a specified probability distribution; in the simulations described here the distribution was uniform between 0.01 and 0.02.

[3] In the game-theoretic literature, this adaptive strategy is known as Cournot best-reply dynamics [Cournot, 1838].

4.2 Price and Profit Dynamics

We have simulated an economy with 1000 buyers and 5 sellers employing various mixtures of pricing strategies. In each of the simulations depicted below, each buyer's valuation of the good $v = 1$, and each seller's production cost $c = 0.5$. The mixture of buyer types is set at $w_B = 0.75$, *i.e.*, 75% are bargain hunters.

The simulation is asynchronous: at each time step, a buyer or seller is randomly selected to carry out an action (*e.g.*, buying an item or resetting a price). The chance that a given agent is selected for action is determined by its rate; the rate ρ_b at which a given buyer b attempts to purchase the good is set to 0.001, while the rate ρ_s at which a given seller reconsiders its price is 0.00002. Each simulation was iterated for 100 million time steps.

GT Pricing Strategy

Simulations verify that, if agents are GT strategists, the cumulative distribution of prices closely resembles the derived $F(p)$ (to within statistical error), and moreover, the time-averaged profit for each seller is $\bar{\pi} = 0.0255 \pm 0.0003$, which is nearly the theoretical value of 0.0250.

MY Pricing Strategy

Fig. 2(a) illustrates the cyclical price wars that typically occur when all 5 sellers use the myoptimal pricing strategy. Regardless of the initial value of the price vector, a pattern quickly emerges in which prices are positioned near the monopolistic price $v = 1$, followed by a long episode during which the sellers successively undercut one another by ϵ. During this latter phase, no two prices differ by more than $(S-1)\epsilon$, and the prices fall linearly with time. Eventually, when the lowest-priced seller is within ϵ above the value $p^* = 0.53125$, the next seller finds it unprofitable to undercut, and instead resets its price to $v = 1$. The other sellers follow suit, until all but the lowest-priced seller are charging $v = 1$. At this point, the lowest-priced seller finds that it can maintain its market share but increase its profit dramatically — from $p^* - .5 = 0.03125$ to $0.5 - \epsilon$ — by raising its price to $1 - \epsilon$. No sooner than the lowest-priced seller raises its price does the next seller who resets its price undercut, thereby igniting the next cycle of the price war.

Fig. 3(a) shows the sellers' profits averaged during the intervals between successive resetting of prices. The upper curve represents a linear decrease in the average profit attained by the lowest-priced seller as price decreases, whichever seller that happens to be. The lower curve represents the average profit attained by sellers that are not currently the lowest-priced; near the end of the cycle they suffer from both low market share and low margin. The expected average profit can be computed by averaging the profit given by Eqs. 7 and 8 over one price-war cycle: $\pi_s^{\mathrm{MY}} = (1/S)[(1/2)(v + p^*) - c]$, which yields $\pi_s^{\mathrm{MY}} = 0.053125$ in this instance. The simulation results match this closely: the average profit per time step is 0.0515, which is just over twice the average profit obtained via the game-theoretic pricing strategy.

Since prices fluctuate over time, it is of interest to compute the probability distribution of prices. Fig. 4(a) depicts the cumulative distribution function for myoptimal pricing. This measured cumulative density function has exactly the same endpoints $p^* = 0.53125$ and $v = 1$ as those of the mixed strategy equilibrium, but the linear shape between those endpoints (which reflects the linear price war) is quite different from what is displayed in Fig. 1(a).

DF Pricing Strategy

Fig. 2(b) shows the price dynamics that result when 5 derivative followers are pitted against one another. Recall that derivative followers do not base their pricing decisions on any information that pertains to other agents in the system — neither sellers' price-setting tendencies nor buyers' preferences. Nonetheless, their behavior tends towards what is in effect a collusive state in which *all* sellers charge nearly the monopolistic price. This is tacit collusion as defined, for example, in Tirole [1988], so-called because the agents do not communicate at all and there is consequently nothing illegal about their collusive behavior. Note that DF sellers accumulate greater profits than myoptimal or game-theoretic sellers. According to Fig. 3(b), sellers that are currently lowest-priced can expect an average profit of 0.30 to 0.35, while the others can expect roughly the game-theoretic profit of 0.025. Averaging over the last 90 million time steps (to eliminate transient effects), we find that the average profit per seller is 0.0841. This is near the absolute collusive limit of $(1/S)(v - c) = 0.10$, which would be obtained if all sellers were to fix their prices at 1.

How do derivative followers manage to collude? Like myoptimal sellers, DF sellers are capable of engaging in price wars; such dynamics are visible in Fig. 2(b). However, these price wars tend to involve only two sellers, and the positive feedback that drives them depends critically on both the sequence of price increments and the timing of the asynchronous moves by the sellers. Downward trends are therefore very easily disrupted. For example, if A's price is currently above B's, but A reduces its price by an amount insufficient to undercut B, then A's profits decrease, so that A raises its price in subsequent time steps. Soon after A breaks the downward cycle, B discovers that it can improve profits by increasing its price, and does so. Simulations clearly show that upward trends in price are much faster and more certain than downward trends. The tendency of a society of DF sellers to reach and maintain high prices is reflected in the cumulative distribution function, shown in Fig. 4(b).

It is also of interest to study the interplay among GT, MY, and DF sellers. Typically, we find that, when a myoptimal seller is introduced into a population of DF or GT sellers, it substantially outplays them, and their profits decline significantly.

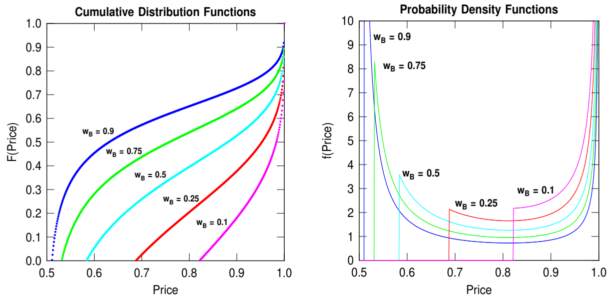

Figure 1: Nash Equilibria for $S = 5, v = 1, c = .5$, and $w_B \in \{.1, .25, .5, .75, .9\}$

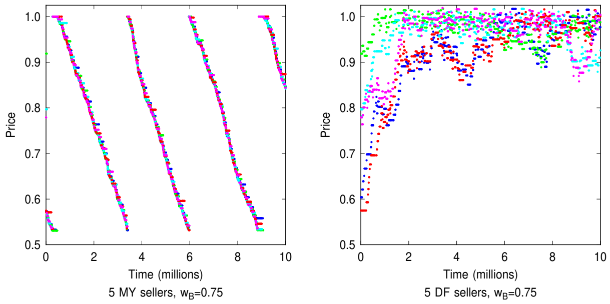

Figure 2: Price dynamics for a) 5 MY sellers and b) 5 DF sellers during the first 10 million time steps.

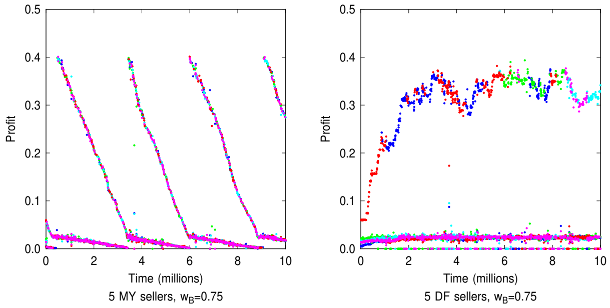

Figure 3: Profit dynamics for a) 5 MY sellers and b) 5 DF sellers during the first 10 million time steps.

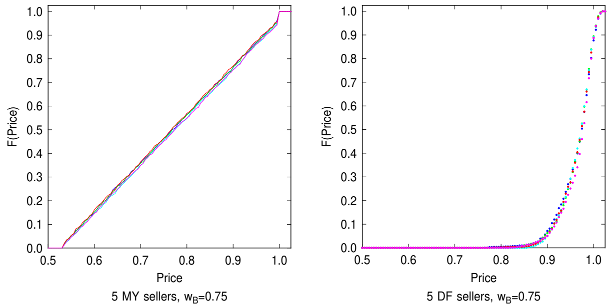

Figure 4: Cumulative distribution functions for a) 5 MY sellers and b) 5 DF sellers, derived from prices observed between times 10 million and 100 million.

5 Evolution of Shopbots and Pricebots

In additional simulations, we investigated a situation in which all five sellers use identical pricing strategies, but one of the sellers resets its price more quickly than the others. We observed that the faster price-setter earns substantially more profit than the others because, for example, in the case of myoptimal agents, it undercuts far more often than it itself is undercut. In the absence of any throttling mechanism, it is advantageous for sellers to re-price their goods as quickly as possible, but this could potentially lead to an arms race in which sellers do so with ever-increasing frequency. In such a world, a human price setter would undoubtedly be too slow and costly, and would be replaced with a pricebot (likely one based on a more sophisticated algorithm than any explored in Section 4!). Almost certainly, this strategy would make use of information about the buyer population, which could be purchased from other agents. Even more likely, however, the strategy would require knowledge of competitors' prices. How would the pricebot obtain this information? From a shopbot, of course!

With each seller seeking to re-price its products faster than its competitors, shopbots would quickly become overloaded with requests. A pricebot representing `amazon.com` might submit a million or more queries (one per book title) to a shopbot every hour — or maybe even every minute! Since shopbots must query individual sellers for prices, they would in turn pass this load back to `amazon.com`'s competitors: *e.g.*, `barnesandnoble.com`, `kingbooks.com`. The rate of pricing requests made by sellers could easily dwarf the rate at which similar requests would be made by human buyers, eliminating the potential of shopbots to ameliorate market frictions.

A typical solution to an excess demand for shopbot services would be for shopbots to charge pricebots for price information. Today, shopbots tend to make a living by selling advertising space on their Web pages. This appears to be an adequate business model so long as requests are made by humans. Agents, however, are unwelcome customers because they are are not influenced by advertisements; as a result, agents are either barely tolerated or excluded intentionally. By charging for the information services they provide, shopbots would be economically-motivated agents, creating the proper incentives to deter excess demand, and welcoming business from other agents. Once shopbots begin to charge for pricing information, it would seem natural for sellers — the actual owners of the desired information — to themselves charge the shopbots for their information. The sellers could use another form of pricebot to dynamically price this information. This scenario illustrates how the need for agents to dynamically price their services could quickly percolate through an entire economy of software agents. The alternative is "meltdown" due to overload which could occur as agents become more prevalent on the Internet. Rules of etiquette followed voluntarily today by web crawlers and related programs could be trampled in the rush for competitive advantage.

6 Conclusion

Game-theoretic analysis of a model of a simple commodity market established a quantitative relationship between the degree of shopbot usage among buyers and the degree of price competition among sellers. This motivated a comparative study of various *pricebot* algorithms that sellers might employ in an effort to gain an edge in a market in which shopbots have increased the level of competition. Pricebots were shown to be capable of inducing price wars, yet even so they may earn profits that are well above game-theoretic equilibrium levels. Future work will explore the dynamics of markets in which more sophisticated shopbots base their search on product attributes as well as price, and in which pricebots use more sophisticated learning algorithms such as Q-learning.

References

[Chavez and Maes, 1996] Anthony Chavez and Pattie 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.

[Cournot, 1838] A. Cournot. *Recherches sur les Principes Mathematics de la Theorie de la Richesse.* Hachette, 1838.

[DeLong and Froomkin, 1998] J. Bradford DeLong and A. Michael Froomkin. The next economy? In Deborah Hurley, Brian Kahin, and Hal Varian, editors, *Internet Publishing and Beyond: The Economics of Digital Information and Intellecutal Property.* MIT Press, Cambridge, Massachusetts, 1998.

[Greenwald *et al.*, 1998] A. Greenwald, E. Friedman, and S. Shenker. Learning in network contexts: Results from experimental simulations. *Submitted for Publication*, May 1998.

[Kephart *et al.*, 1998] J. O. Kephart, J. E. Hanson, D. W. Levine, B. N. Grosof, 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.

[Lewis, 1997] T. G. Lewis. *The Friction-Free Economy: Marketing Strategies for a Wired World.* Hardcover HarperBusiness, 1997.

[Nash, 1951] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[Tirole, 1988] Jean Tirole. *The Theory of Industrial Organization.* The MIT Press, Cambridge, MA, 1988.

[Tsvetovatyy *et al.*, 1997] M. Tsvetovatyy, M. Gini, B. Mobasher, and Z. Wieckowski. MAGMA: an agent-based virtual market for electronic commerce. *Applied Artificial Intelligence*, 1997.

[Varian, 1980] H. Varian. A model of sales. *American Economic Review, Papers and Proceedings*, 70(4):651–659, September 1980.