PRICING PRICE INFORMATION IN E-COMMERCE

Panos M. Markopoulos

Computer and Information Science Dept.
University of Pennsylvania
markopou@unagi.cis.upenn.edu

Lyle H. Ungar

Computer and Information Science Dept.
University of Pennsylvania
ungar@cis.upenn.edu

ABSTRACT

Shopbots and Internet sites that help users locate the best price for a product are changing the way people shop by providing valuable information on goods and services. This paper presents a first attempt to measure the value of one piece of information: the price charged for goods and services. We first establish a theoretical limit to the value of price information for the first seller in a market that decides to sell price information through a shopbot and quantify the revenues that the seller can expect to receive. We also discuss seller competition in selling price information and analyze the equilibria that our model predicts. We then demonstrate how our analysis can be used to argue about the information value and pricing of other product attributes, such as quality. Finally, we drop our model's assumptions to discuss whether and how much of the theoretical value can actually be realized in equilibrium settings, and the practical problems and implications of our ideas. Our results give counterintuitive predictions about the future forms of some electronic markets, including the possible collapse of the free shopbot model.

1. INTRODUCTION

Shopbots and Internet websites that perform comparisons of prices and product characteristics change the way people shop for goods and services.[1] They reduce search costs and offer the potential for new forms of search and market structures for buying and selling information about products.

Even though the reduction of consumers' search costs has caused increased competition among sellers, when compared with conventional markets of bricks and mortar retailers, Internet markets are far from frictionless and considerable and persistent price dispersion exists [4]. This price dispersion is big enough to make price information valuable to consumers who are price sensitive and want to compare prices before buying a good or service.

[1]An extended abstract of this paper has appeared in the proceedings of the 2001 ACM Conference on Electronic Commerce(EC'01). [14]

Traditionally, price information has been provided for free to potential customers. Even in cases where price information was provided to customers through third parties, who would benefit from it, sellers did not have the technological means to either charge these intermediaries for the price information or prevent them from obtaining it. Furthermore, today, it would be infeasible to charge the buyers for price information, given the lack of a widespread micropayment scheme.

However, in the future, we expect that buyers will employ software agents that will purchase physical goods on behalf of their human owners. It would then be feasible for the shopbots or the sellers themselves to charge buyer agents for this information. The sellers of price information must therefore know how much they can charge for it. In order for the buyer agent to know which price quotes to purchase, it must also know the value of the price information to its human owner. It is thus important to explore how one can price price information.

We argue that price information has economic value and hence can be bought and sold. If shopbots that provide free services and profit only by selling advertisement space do not want to pass the cost of price information on to consumers, they will have to share their advertising revenues with demanding sellers. This will be an issue particularly in markets with high price dispersion and might lead to the collapse of the free shopbot model (or, more accurately, the constant cost model) for these markets. The incentives that sellers have to start charging for price information cannot be overlooked. We calculate below the extra revenues an online bookstore would have, had it charged the optimal amount for this information, to be on average 15.6 cents per shopper, in a market of seven sellers. This can be interpreted as if the price information for a book can generate an extra $\$0.156 \times 7 = \$1.09$ before the actual book is sold through a shopbot in a market with an average book price of \$13.69 [18]. Thus, sellers may start selling their price information that is delivered through shopbots and start competing on this price dimension too.

Yannis Bakos [1] has shown that sellers have strong incentives to manipulate the way buyers access their price information. After separating search costs for price and product information, he shows that sellers prefer buyers to have high search costs when searching for price information, even when product information is relatively easy to acquire.

Two recent papers have considered price information on the Internet as a good that can be traded: Kephart and Greenwald [10] consider the possibility of a shopbot manipulating the buyer's search cost structure. They first determine the optimal buyer behavior, given a certain search cost structure, and then consider what happens if a shopbot can manipulate the buyers' cost structure by charging money for the service it provides. The buyer agent can specify how many price quotes it wants to acquire, but it cannot specify which. The shopbot then randomly samples and delivers the number of price quotes the buyer agent demanded. The authors consider both linear and non-linear cost structures on the number of price quotes.

Baye and Morgan [2] propose a model in which a "gatekeeper" charges a fee to both sellers who advertise their prices via the gatekeeper and buyers who access the price quotes. The gatekeeper is an intermediary in an information market that supports the product market.

In a sense the focus of these two papers is the ability of the price comparing intermediary to realize some of the added value its service generates to other participants in the market. Both papers strengthen our belief that the idea of selling price information has become plausible with the advent of economically motivated software agents.

One characteristic of these approaches is that they assume that the shopbot/intermediary is a monopolist, as it faces no competition. In these models, competition would drive price quote prices down to the marginal cost of acquiring them. The intuition behind our approach is that a seller can manipulate his own product price, and subsequently his information value, but a shopbot cannot. A seller does have monopoly power over his information as he is the only one that can guarantee that the information is accurate.

Our approach, thus, differs from previous work as we consider sellers selling price information directly to price sensitive buyers. This paper is the first, as far as we know, that connects product prices with the information value they generate to the seller of the products, and its central contribution is that it proves that the issue is important in e-commerce and might influence the structure of future Internet markets.

The rest of the paper is structured as follows: In section 2 we give the motivation for our work and our model choices. We present in section 3 our model of a commodity market where people shop through shopbot sites and derive the optimal pricing scheme for price information when only one seller sells price information. We prove that there are economic incentives for a seller to start selling price information, even when every other seller in the market provides it for free. In section 4 we extend our model and consider sellers competing in selling price information and point out that the lowest product price seller might end up lowering total revenues (product sales plus price information sales revenues), unless he prices taking the second lowest price in account. We also analyze two possible equilibria. In section 5 we show how our model can be extended to price quality information. In section 6 we drop our model assumptions to discuss price information value in the real world. More specifically, we examine the obstacles that a seller would have in realizing the maximum value of his price information, as computed in sections 3 and 4. We conclude in section 7.

2. MOTIVATION

This section explains our interest in the pricing of price information and explains our choice to study a model where sellers sell their price information directly to buyers who wish to perform comparison shopping.

2.1. Towards more efficient markets

The most important motivation for our interest in the pricing of price information is related to the socially optimal level of information available to market participants in different markets. It is shown in [1] that sellers will not only fail to invest in better search technology for the buyers, but they will actually *resist* any scheme that reduces the buyers' search costs, unless they (the sellers) can capture the bulk of the buyers' efficiency gains.

For example, different airline reservation systems such as PARS, SABRE and Apollo, were initiated by airlines (TWA, American and United respectively) to help travelers find suitable flights. However, by passing a system utilization fee to the passenger's ticket they could effectively control the amount of the traveler's efficiency gain that they could capture. Furthermore, the reservation systems were used as a tool to promote the controlling airline's offers at the expense of information transparency. Bakos [1] explained in detail the incentives of these sellers to manipulate the amount of information they provide, making search more expensive for the buyer.

Sellers have been observed to effectively oppose search cost reduction efforts initiated by buyers. For example, in the famous case of the "Bargainfinder" electronic agent, music CD stores forbid the agent from accessing their price databases. Even auctioneers have been known to oppose intermediaries who seek to facilitate buyers locating items across many auctions, as in the widely publicized legal dispute between Ebay and different auction aggregators that provided information from a variety of auction websites.

Currently, intermediaries such as Travelocity and Expedia understand the seller side demand for the existence of information "friction" and the sellers' reservation to reducing consumers' search costs. Both web sites do not completely facilitate price comparison, but rather require a lot of "clicking" and patience from the traveler that wishes to find the cheapest available suitable flight.

Furthermore, most shopbot websites are in reality intermediaries, whose incentives do not always match the consumer's. Shopbot-seller deals are quite common as shopbots seek ways to increase their revenues, often at the expense of information transparency [5].

Thus, a deadlock appears in the problem of designing a market mechanism that provides the socially optimal amount of product information to buyers: if that mechanism is not controlled by sellers, they will effectively oppose it and if it is, they will use it to promote their own interests on the expene of information transparency.

Selling price information presents a way out of this deadlock: by selling their product information the sellers can control the amount of buyer efficiency gains they can capture, so they would not oppose such a mechanism. But, additionally, they will have the incentive to improve the richness, timeliness and overall quality of the information they provide in order to make their information "bundle" (that includes product characteristics and price) attractive and thus increase their information revenues.

It is thus interesting to work towards models that connect the quality of the information that the sellers provide, to some monetary payment that can induce better information. Our model is a first such attempt and lays the foundations for product information markets, not to explicitly demonstrate how monetary payments for product information can be used to achieve information transparency and market efficiency, but rather to explore the feasibility of the idea, proving that monetary payments can actually occur, and that price quote prices are not driven down to marginal cost even in competition settings, because sellers retain a degree of monopoly power over their information.

2.2. Applications to multiagent systems

There has been much interest in the design and development of multiagent systems comprised of economically motivated software agents [22]. Primitive multiagent economies [8] shed light into future forms of online shopping and multi-agent, market-based, solutions are tested in fields ranging from train scheduling to temperature control with the goal of achieving better resource allocation than traditional solutions [15] [7].

A common problem arises in many asynchronous systems of this type: agents that wish to know the current market prices and take action faster than their competitors contribute to a "tragedy of the commons" situation, where the system is strained (bandwidth, computation and databases) as agents demand faster and faster updates on the current price schedule. It is hardly surprising that in those high-speed, ultra-rational settings, any item, such as price information, that has value and is provided for free will tend to be overused.

In [11], Kephart, Hanson and Greenwald study a multi-agent economy where goods are traded between economically motivated software agents that extensively use techniques such as comparison shopping and dynamic pricing. They describe a catastrophic outcome, where the sellers, by trying to increase the rate by which they can update their prices and dominate their competitors, increase the rate by which they request price quotes and engage in a speed competition that ends up consuming all available system bandwidth, before the market breaks down. A solution they propose is making price information expensive to acquire.

3. ONLY ONE SELLER CHARGING FOR PRICE INFORMATION - MAXIMUM THEORETICAL VALUE

Far from following the Bertrand view of markets, where prices are driven down to marginal cost, products sold on the Internet demonstrate a significant and persistent price dispersion [18][6][3]. Smith, Bailey and Brynjolfsson [18] identify product heterogeneity, convenience, lock-in, brand, awareness and price discrimination as potential sources of this price dispersion. It seems plausible that price dispersion exists due to a complex interaction of all these potential sources. This paper assumes for simplicity that the price dispersion on the Internet exists due to the presence of multiple classes of consumers that value products in different ways and demonstrate distinct shopping behaviors. The sellers have optimized by selecting a price for their products that maximizes their revenues across all these classes of consumers. We further assume that one such distinct class is composed of shopbot users who buy products only based on price.

In order to present a closed model, some additional discussion and assumptions are needed. A question that arises is why sellers give price information for products for which they know they do not offer the lowest price. Even though, under our assumptions, sellers will never make a sale through a shopbot when they do not offer the lowest price, they are indifferent in providing this information, since if they deny it, they will not make a sale either. Also, by displaying their information on a shopbot site, they increase their brand recognition among consumers. Furthermore, sellers who choose to sell their price information might make additional revenues from the shopbot, even though they will not make a sale.

We consider a market with N sellers that sell many different undifferentiated goods like books, CDs, electronics, etc. That is, sellers could be bookstores that sell many different books, but each book is a commodity product across all sellers, uniquely identified by its ISBN.

We assume that sellers draw their product prices from a random distribution $f(x)$ with a corresponding cumulative distribution denoted by $F(x)$, constant across time, that is exogenous to our model and common knowledge. This assumption is rather strong since we cannot expect all books, for example, to follow the same price distribution, but it does not effect our results qualitatively. Our results will be an overestimate for items that display higher than the average price dispersion and an underestimate for the items that display a lower than the average price dispersion.

Our assumption that sellers draw their prices from a random distribution, is justifiable if we consider the buyers' view of the market. Another way to describe our model is as if buyers *believe* that sellers draw their prices from a random distribution $f(x)$. Sellers that offer many items in different prices have a price distribution across all the items that they sell. A buyer is likely to have some idea of the overall price levels of a seller based on previous experience[2]. However the buyer is still uncertain about the price of the particular item she is currently interested in. The buyer's best guess would be that the price is randomly distributed according to $f(x)$.

In [3] Baye, Morgan and Scholten argue that the price distribution f is a persistent phenomenon varying across markets, rather than a temporary disequilibrium in Internet markets.

Prospective buyers know the distribution of the prices in the market, but do not know which seller is the cheapest for the particular product they are interested in. They have no preferences for particular sellers and they are willing to buy from the one that offers the lowest price. For this reason they shop through a shopbot site that displays the prices of all sellers in the market for the product that the buyers specify.

Seller j decides to sell his price information and contracts the shopbot to deliver all product information for free, except the price, which will be available at an additional fee. The buyer can click on a button or a link and get j's price immediately by paying a price p. In figure 1 we show such a hypothetical agreement between $Bookfinder.com$ and $Fatbrain.com$. We queried $Bookfinder.com$ for a particular edition of Umberto Eco's "Name of the Rose" and found 7 sellers that offer the book. We changed the appearance of the page to remove used books and other editions of the book and also removed the $Fatbrain.com$ price from the data. We also summarize the same information in table 1.

[2]The StreetPrices.com shopbot displays graphically the price range for a particular product, even before the user has access to actual prices. A potential buyer can thus obtain a good idea of the price dispersion for the product. Other websites, such as BizRate.com will provide information on the overall price levels of a particular seller. It is thus feasible for a buyer to form reasonable expectations for the price levels of a particular seller on a particular product even without past experience.

Retailers at a Shopbot		
Bookstore	**ISBN**	**Price**
AllDirect.com	0156001314	\$8.96
A1Books	0156001314	\$10.75
Kingbooks.com	0156001314	\$11.20
BN.com	0156001314	\$12.00
Amazon.com	0156001314	\$12.00
Powell's Books	0156001314	\$14.00
Fatbrain.com	0156001314	Price Info: 10 cents

Table 1: A hypothetical example of a seller charging for price information through a shopbot

There are three distinct entity classes in our market: the buyers, the sellers and the shopbots. Each class faces a different problem. We address each of these problems in turn.

3.1. The buyer's problem

The buyers have three choices: they either pay for the additional price, do not pay but simply accept the lowest price among the remaining $N-1$ sellers without bothering to learn j's price, or incur a fixed cost of inconvenience c, assumed to be the same for all buyers, to visit the seller website directly. This inconvenience cost can range from a few cents, in the case of a simple web query, to a few dollars in the case of a difficult to search web interface, to infinity for the case that the seller does not provide any price information even at his website[3].

Assuming that the buyers are rational, they will want to learn seller j's price if they expect that the cost reduction would be more than the cost of acquiring the price quote. Given that the minimum price for the good among the remaining $N-1$ sellers is q, the buyer knows that the expected decrease in the minimum price from another search is equal to:

$$g(q) = \int_0^q (q-x)f(x)dx = \ldots = \int_0^q F(x)dx \qquad (1)$$

She is willing to pay $j's$ price information price, p, to learn j's price if $p < g(q)$ and $p < c$. If $p > g(q)$ and $c > g(q)$ the buyer is better off by purchasing the item priced at q without requesting any additional information. Finally, if $c < g(q)$ and $c < p$, the buyer will visit j's website directly to learn j's price.

We expect that in the future, buyers will delegate these "micro-optimization" decisions to software agents that will

[3]An example of a very high alternative search cost would be the case that the buyer has to have knowledge of a promotion code, usually available through direct marketing or printed press advertisements. If the buyer does not happen to have the code it is very difficult to find one and learn the item's price.

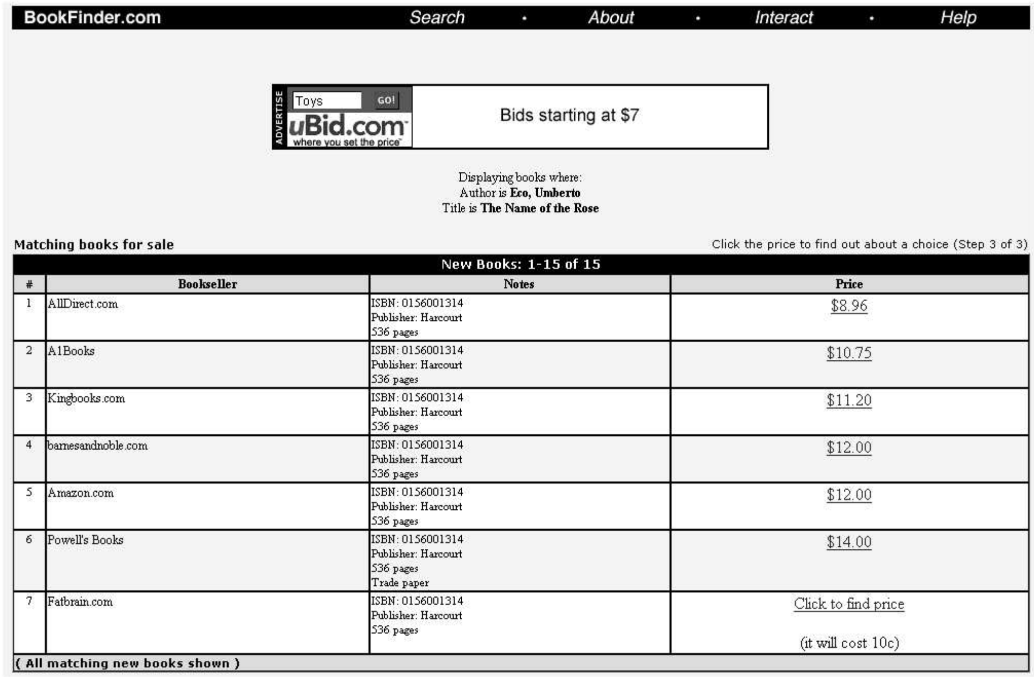

BookFinder.com Search · About · Interact · Help

Toys GO! ADVERTISE uBid.com where you set the price — Bids starting at $7

Displaying books where:
Author is **Eco, Umberto**
Title is **The Name of the Rose**

Matching books for sale — Click the price to find out about a choice (Step 3 of 3)

New Books: 1-15 of 15

#	Bookseller	Notes	Price
1	AllDirect.com	ISBN: 0156001314 Publisher: Harcourt 536 pages	$8.96
2	A1Books	ISBN: 0156001314 Publisher: Harcourt 536 pages	$10.75
3	Kingbooks.com	ISBN: 0156001314 Publisher: Harcourt 536 pages	$11.20
4	barnesandnoble.com	ISBN: 0156001314 Publisher: Harcourt 536 pages	$12.00
5	Amazon.com	ISBN: 0156001314 Publisher: Harcourt 536 pages	$12.00
6	Powell's Books	ISBN: 0156001314 Publisher: Harcourt 536 pages Trade paper	$14.00
7	Fatbrain.com	ISBN: 0156001314 Publisher: Harcourt 536 pages	Click to find price (it will cost 10c)

(All matching new books shown)

Figure 1: A hypothetical example of a seller charging for price information through a shopbot

be authorized to purchase product information under budget constraints. These agents will try to minimize the buyer's total cost (product plus search costs) while providing total process transparency to their human owner.

3.2. The seller's problem

Seller j's problem is how to price his price information to maximize revenue. Buyers have queried the shopbot site and found that the current lowest price for the product they require is q, where q follows some distribution[4] $f_1^{(N-1)}(x)$ that depends on $f(x)$ and the number of other sellers, $N-1$. Seller j would know that, since he too can query the shopbot and obtain the free price quotes.

Again, we repeat our assumption that the seller does not change his price. The price has been set by the seller to maximize revenues across all different categories of customers, including those that shop directly from the sellers web site without intermediaries. For example the seller might not be allowed to price discriminate (charging different customers different prices) or might be afraid that a price reduction might ignite price wars with other sellers (See section 6.3).

The seller can safely set the price of his price information to ϵ below $g(q)$, given by equation 1, knowing that rational buyers will always want to know his price. However, j cannot charge more than c, the inconvenience cost to the buyer of visiting j's website directly. So, the seller would set his price to

$$p(q) = \min(g(q), c), \tag{2}$$

where, again, q is the minimum price observed so far.

We have assumed that should the buyer decide to pay for seller j's price, the price will be displayed immediately when requested (the shopbot pre-fetches the price but withholds it), so there are no wait costs.

The expected revenue per customer for seller j from selling price information as ϵ goes to zero is thus:

$$\Pi = \int_0^\infty f_1^{(N-1)}(q)p(q)dq \tag{3}$$

It is informative to estimate how much an online bookstore would be able to charge for its price information to shopbots or shopbot users. We use internet book price dispersion data, including all costs (shipping etc.), collected in [4] and assume that all books follow the same price distribution. That means that our results would be an underestimate in the cases of books that exhibit higher price dispersion than the average, and an overestimate for the books that have lower than average price dispersion. We have fitted the authors' de-meaned experimental data with the normal price distribution with mean zero and standard deviation 2.0.

[4]This is simply the 1st order statistic, the distribution of the lowest value in $N-1$ draws from distribution f, $f_1^{(i)}(y) = i(1-F(y))^{i-1}f(y)$. [9]

The shopper session described in figure 1, is a representative case of price dispersion[5], described well by the normal distribution with standard deviation 2.0, the shopper has discovered a minimum price of $8.96. The expected gain that a shopper would have from knowing one additional price is $p(\$8.96)$, close to 9.8 cents. The seller could charge just below 9.8 cents for this book's price, to make sure that rational shoppers would pay.

The average revenue per buyer, the seller would expect from selling price information in this market with seven sellers is, from equation 3, approximately 15.6 cents per shopper, assuming an alternative search cost $c = 50$ cents[6] to obtain the same information by other means. For 20 sellers the expected revenue per buyer drops to 4.3 cents.

If seller j prices his price information with the method suggested above, a rational buyer will always purchase it. So, seller j makes the revenues described by equation 3 every time a consumer uses the shopbot. Furthermore, the seller makes sure that when he happens to sell the cheapest product, this information will eventually be conveyed to the buyer and thus he will not risk losing a product sale.

We use the term *"multiplicative effect"* to refer to the fact that a buyer requires multiple price quotes for a single purchase. Conversely, a seller would potentially be selling price information more often than products. For example in a market of seven bookstores, assuming each bookstore is equally likely to carry a cheapest book, each bookstore is expected to make 1/7 of total sales. If a seller is the only one that sells price information, pricing according to equation 1, rational buyers would always purchase this information. So for an average price information price of 15.6 cents, the bookstore's price information would on average generate $\$0.156 \times 7 = \$1.092$ per book that the bookstore sells. For an average internet book price of $13.69 the information revenues would bring a revenue increase of approximately 8%. Or, even more accurately, in a market of seven sellers and an average book price of $13.69, each seller expects to sell books with an average price of the expectation of the first order statistic of the price for seven sellers, which is about $11, which raises the revenue increase to approximately 10%. Furthermore, since it is arguably far less costly to sell information rather than products, the profit increase would be considerably higher.

We have plotted seller j's revenues from selling price information, per book that seller j sells, in figure 2, as a function of the number of sellers in the market. We assume that all other sellers give their price information for free while seller j prices it according to equation 2, and that the buyers have an alternative search cost $c = 50$ cents. We used

[5]Indeed, the average price for this book is $11.49 and the expectation for f^6_{min} is -2.53, meaning that we expect the minimum price to be $11.49-$2.53=$8.96, a pleasant coincidence when we tested the data.

[6]Based on a back of the envelope calculation for what 2 minutes of time worth to the average shopper, given today's US salaries.

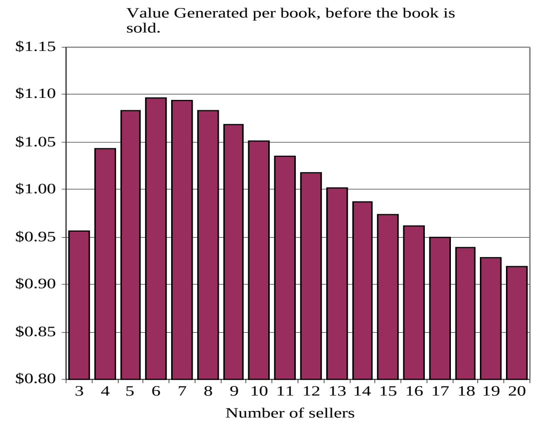

Figure 2: Maximum theoretical value that a book's price information can generate, before the actual book is sold. Only one of the sellers charges for price information.

the formula $\Pi \cdot N$ for the seller's information revenues per product sold, where N is the number of sellers in the market and Π is defined in equation 3.

We see that the price information value for a seller in a market of 20 sellers is not significantly lower than the same value in a market of 3 sellers, or the peak value for 6 sellers. Intuitively, this is because in a market of three sellers, each seller would sell books more often (once every three sales, on expectation), so each book would not generate much value before the sale. More sellers in the market would mean that a price quote would be requested more times, before the book is actually sold. So, even though the price information value per buyer is considerably lower, for twenty sellers, its importance in terms of the additional seller revenue is still significant.

Surprisingly, the value does not depend monotonically on the number of sellers and it peaks when six sellers compete. This is an effect of the buyer's reservation price of 50 cents. In markets with very few competitors, seller j's price information if very valuable, but the seller cannot price it higher than the buyers' alternative search cost. As a result the information revenues per product sold are relatively low.

3.3. The shopbot's problem

We assume that the shopbot's utility is an increasing function of the traffic it generates. We further assume that the shopbot does not subsidize the buyers' price quote payments. We deal with the case where a shopbot can subsidize the buyer's search costs in section 4.4, where we show that in certain markets shopbot services will not be free for shoppers.

The problem that the shopbots face is whether or not to

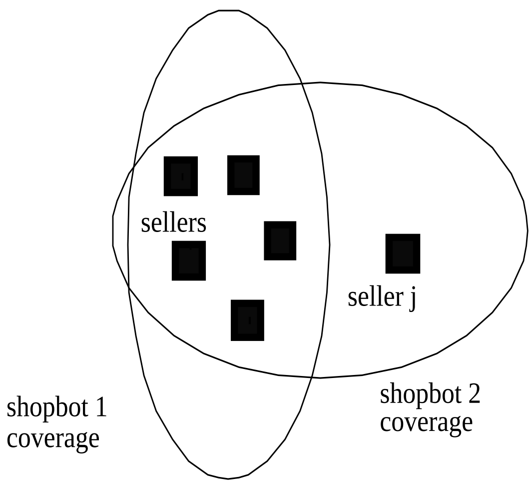

Figure 3: Shopbot 1 covers all sellers except j, while shopbot 2 covers all sellers, but charges to reveal j's price.

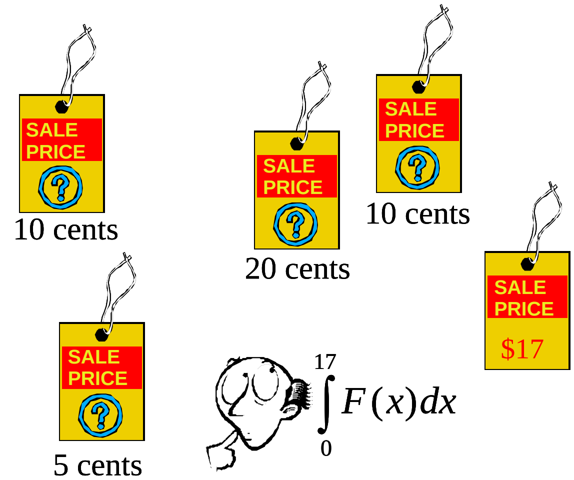

Figure 4: Sellers demand a payment for their product price information to be displayed to the buyer as part of a shopbot's search results. The buyer has already discovered a seller that charges $17 for his product.

accept an arrangement with the jth seller. We consider the case where a shopbot $S1$ does not wish to help sellers charge for price information and that the seller contracts another shopbot $S2$ to do so. (See figure 3.)

A prospective buyer would have to choose between the two shopbots. Assume that she can only choose one, since if she wants to visit shopbot 1, and then after getting a minimum price go to shopbot 2 to see if there are other sellers in the market, the model becomes equivalent to the model studied in section 3.1.

A rational buyer knows that the expected cost, should she visit $S1$ is given by $E(f_{min}^{N-1})$, while should she visit $S2$, is $E(f_{min}^{N})$. If it is known to the buyer that the cost to get an additional price quote is less than $E(f_{min}^{N-1}) - E(f_{min}^{N})$, then the buyer will always visit $S2$. If the seller charges for price information using the method described in section 3.2, then the shopper would visit $S2$. Thus, the first shopbot that makes an exclusive deal with one or more sellers will attract all traffic, and other shopbots will also want to make similar deals.

We have shown in this section that if all sellers reveal their price information for free, it pays for a seller to start charging for his price information. We now proceed to study an equilibrium where all sellers charge for their price information.

4. SELLER COMPETITION

In this section we present a model where all sellers charge money to reveal their price.

We consider a shopbot that displays price information of N different sellers that sell many different undifferentiated goods. Again, product prices follow a random distribution $f(p)$ with the equivalent cumulative distribution denoted by $F(p)$, constant across time, that is exogenous to our model. Shopbot users know the distribution of the prices in the market, but do not know which seller is the cheapest for the particular product they are interested. They have no preferences for particular sellers and they are willing to buy from the one that offers the lowest price.

4.1. The buyer's problem

In the extreme case where all sellers charge for price information, the prospective buyer cannot find any shopbot that provides price information for free[7].

In figure 4 we display the buyer's problem, when more than one sellers sell their price information. The buyer has discovered a seller that charges $17 for the product, either by purchasing the price quote or because the seller provides his information for free. How will the buyer continue with his search? should he stop and accept the $17 product?

If the price quote prices carry no information about the hidden product prices we can refer to the optimal consumer behavior in the classic economic literature [20][17][12][21]. Optimal search behavior has been explored in connection to search costs that are exogenously given, as in the case of travel time or communication costs. In the classic economic

[7]Unless the shopbot wishes to absorb the cost for the buyer. In this case the buyer's problem is trivial, since all price information is displayed for free. This scenario is explored in section 4.4

literature, search costs do not convey any information about what the sellers might charge for their products and the sellers cannot use search costs to signal the buyer about their product prices.

In such settings, with the sellers drawing their product prices from the same distribution, the buyer would keep searching as long as the expected decrease in the minimum price for another search is less than what the buyer has to pay for it. If the buyer believes that all the product prices follow the same distribution, the search would start from the prices that are cheaper to acquire. We assume that if two price information prices are the same, the buyer chooses one at random. If the current minimum price is q, then the expected decrease in the minimum price from another search is a function of q and f, as described by equation 1.

The buyer would start by requesting the price quotes that are cheaper to acquire and stop when the lowest price quote price p is greater than the expected marginal product price decrease: $p > g(q)$, where q is the minimum product price observed so far. The optimal sequential decision rule is for the shopper to continue searching if the lowest price observed up to that point is greater than b, where b is the solution to $g(b) = p$.

However, in our model, the problem is far more complex. Consider the seller in figure 4 that demands 20 cents for his price quote. In the classic buyer search literature, the buyer would visit this seller last or she might not visit him at all. In our model there is no simple answer to the question "What should the buyer do?". It depends on what the buyer believes and can conclude about the product prices, given the price quote prices. In other words it is a matter of how the buyer translates the sellers' signals.

We have solved the problem in [13] for the case of one buyer and two sellers with arbitrary product price distributions. The buyers, that know the sellers' product price distributions, observe the sellers' price quote prices and use Bayes' rule to update their beliefs about the underlying product prices.

In the next section we present one of the Nash equilibria to a simplified version of the game for N sellers, where they all draw their prices from the same distribution. In this simple Nash equilibrium the buyers do not need to update their beliefs about the underlying product prices until they actually discover one of the two cheapest sellers. This is due to the fact that all sellers appear identical to the buyer, because they charge the same price quote price.

4.2. The seller's problem

The first problem a seller would face, in a market for price information, is knowing the product prices of competitors, which would not be available for free to shopbots or price-comparison web sites. Instead of incurring this cost, a seller might choose to buy only price information for some of the competitors products and make price information pricing decisions based on estimates of product price dispersion. Or, sellers may provide price information to a shopbot that would then sell it to potential buyers for a commission. That shopbot would function as a trusted seller proxy. It would have all available information to make optimal pricing decisions to maximize its own commission and seller revenue from selling product price information.

The price p^* that jointly maximizes price information revenue for N sellers is the price that maximizes $\Pi(p)$,

$$\Pi(p) = \frac{p}{N}(1 + \sum_{i=2}^{N} Q(p,i)) \tag{4}$$

where $Q(p,i)$ is the probability that a seller will be the ith to be queried given price p for his price information. The formula is derived in the Appendix.

The model we presented in this paper is actually a somewhat simplified version of reality. In fact, selling product price information according to equation 4 might actually make the seller worse off if the pricing is not done carefully: As we saw in section 4.1, it is not guaranteed that the buyer will actually discover the cheapest seller's price. Thus, the cheapest seller for a particular product, who would have captured all buyers, had he provided price information for free, will now make only a fraction of sales and may end up lowering total revenue.

A solution to this problem would be for the cheapest seller to price price information cheaply enough so that a rational shopper would always request it, even though the second lowest price in the market has been discovered. If the cheapest seller knows that the second lowest price in the market is q_2, then by charging

$$I(q_2) = \int_0^{q_2} (q_2 - x) f(x) dx - \epsilon \tag{5}$$

for price information, he makes sure that the buyer will always eventually discover the lowest price.

One Nash equilibrium of the one-shot game we described dictates that the cheapest seller prices his price information according to equation 5 and all other sellers follow by setting their price information price to the same value as the cheapest product seller.

The buyer would keep purchasing price quotes until she discovers the cheapest product price. Since all price quote prices are set according to equation 5, when the cheapest product price is discovered, all price quote prices become too expensive for the buyer to pay. The buyer purchases from the cheapest seller that has just been discovered. On average we expect that the buyer will need $N/2$ price quotes, before the cheapest seller is discovered.

For a seller, the expected information revenue per buyer, given that N sellers sell their price information is:

$$R_{buyer} = \int_0^{\infty} f_2^{(N)}(q)p(q)(\frac{1}{N} + \frac{1}{2}(1 - \frac{1}{N}))dq \qquad (6)$$

given that the cheapest product seller sets his price quote price according to the second cheapest product in the market (equation5), and that other sellers pool by charging the same amount. $f_2^{(N)}$ is the second order statistic for N sellers[8]. That is because, for every second lowest product price q the average revenue per customer is $p(q)$, with probability $1/N$ (corresponding to the case where the seller is the cheapest product price seller and will definitely sell his price information) plus $p(q)$, with probability $\frac{1}{2}(1 - \frac{1}{N})$, which corresponds to the case where the seller is not the cheapest and only half the times his price information will be requested, before the actual cheapest seller is found.

The expected value that the price information of a product generates, before the product is sold, in a market of N sellers that sell their price information is:

$$R_{product} = N \cdot R_{buyer} \qquad (7)$$

We have plotted R_{buyer} and $R_{product}$ in figure 5, assuming normal product price distribution, for different number of competitors and different price distribution standard deviations. R_{buyer} is monotonically decreasing because it depends on the expected difference of the two lowest prices. This difference is decreasing as the number of competitors increases, due to our assumption that sellers draw their prices from an exogenous distribution f. However, this seems to be verified by Baye, Morgan and Scholten [3] who found that indeed the difference in the prices of the two cheapest products decreases as the number of sellers increases. Our model's predictive power is thus strengthened by real world data.

A seller is always better off, though, by selling his price information than giving out price information for free. This suggests that in our model, no seller will want to reveal price information for free to a shopbot.

4.3. Equilibrium analysis

To show one Nash equilibrium, we developed a simplified version of the game as follows: The sellers draw their product prices from the same distribution, observe all outcomes, and the cheapest seller moves first by choosing a price for his price information. All other sellers play simultaneously by setting their price quote prices. Finally the risk neutral buyer starts purchasing price quotes sequentially. Furthermore, we assume that in any circumstances the sellers would rather sell their product than information. In other words, the cheapest seller would never set a price for his price information that would allow another seller to "steal" the product sale from him, for example if that other seller provides his price information for free.

One Nash equilibrium of the one-shot game we described dictates that the cheapest seller prices his price information according to equation 5 and all other sellers follow by setting their price information price to the same value as the cheapest product seller.

It is not hard to see why this behavior is a Nash equilibrium for the one-shot game. A simple proof consists of checking to see if any player has incentives to alter their strategy, given that everybody else plays the way we described. First, a seller, other than the cheapest one cannot alter their strategy, by choosing a different price quote price and increase his expected profits. The reason is that by choosing a different price than everybody else, he stands out as the non-cheapest seller. The buyer will just ignore him and never purchase his price quote, for any non-zero price quote price. This is because the buyer knows that it will always be sequentially optimal to keep purchasing price quotes until the cheapest seller is discovered. So, it does not make sense to purchase the price quote of a seller for which she knows is not the cheapest. The cheapest seller will not choose a price quote price lower than the value given by equation 5, because he would be reducing expected information revenues: the buyer is always willing to pay as high a price as given by equation 5, so there is no reason to choose a lower price quote price. The cheapest seller cannot choose a price quote price, higher than the value given by equation 5, because the following may occur: the second cheapest seller would charge the same price quote price as the cheapest seller and the buyer might happen to request the second cheapest seller's price quote first because she cannot distinguish between a-priori similar sellers that have chosen the same price quote price. At the moment that the buyer discovers the second cheapest seller's product price, all price quote prices become too expensive to continue obtaining, because they would be set to a value, higher than equation 5. The buyer would stop purchasing more price quotes, and buy the second cheapest seller's product. The cheapest seller would have lost the product sale, which would, by assumption, reduce his profits. Finally, the buyer has no incentive to alter his behavior of sequentially purchasing price quotes, until he discovers the cheapest seller. Equation 5 guarantees that if the lowest currently discovered product price is not lower than the second cheapest product price in the market, it is always optimal for the buyer to keep purchasing price quotes, because on expectation she reduces her total cost (search cost plus product cost). The game is thus in equilibrium, since no player has an incentive to deviate.

Our claim that a Nash equilibrium consists of all sellers

[8] In other words, this is how the second lowest price in the market is distributed.

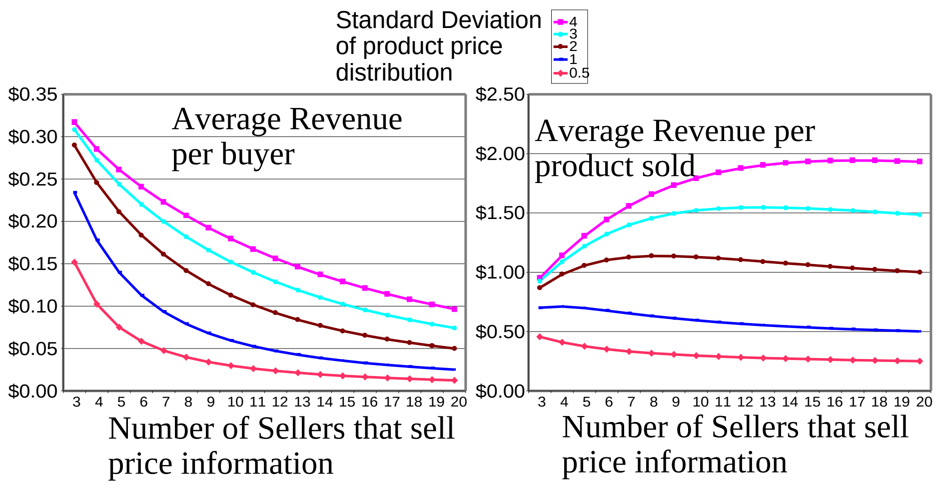

Figure 5: The series represent different standard deviations of the assumed normal product price distribution, in dollars. A reservation price of 50 cents per price quote is assumed for all buyers. The series that have standard deviation 2.0 correspond to real internet book price dispersion data.

charging for their price quotes according to equation 5 includes a "hidden" assumption that might not always hold. We have assumed that when the buyer obtains the second cheapest product seller's price quote, she does not realize that this is indeed the second cheapest product seller. If prices are drawn from continuous distributions, as in our model, a buyer would certainly know that she has just discovered the second cheapest seller. The buyer will actually know what the second cheapest product price in the market is, just by observing the price quote prices and reversing equation 5. When that price is discovered the buyer will use a different rule for determining which price quotes to obtain: in the new rule the buyer will use the new information that all but one of the remaining price quotes are useless.

Thus, in order for the cheapest seller to achieve this equilibrium — since he is the one that sets the pace of the game and other sellers follow — he must randomize his price quote price by subtracting a random quantity from equation 5, that would make the buyer uncertain that he has discovered the second cheapest seller, if this actually occurs. All other sellers would then match the cheapest seller's price quote price. Strictly speaking, according to our model, which uses continuous distributions for the product prices, equation 5 is an upper bound of the equilibrium that the sellers can achieve.

If the cheapest seller is not allowed to randomize then we proceed to describe a fixed equilibrium.

If the buyer discovers the second cheapest seller that charges q_2 and there are K sellers remaining with unknown price quotes, then he is not willing to pay more than:

$$I'(q_2) = min\left(\frac{\int_0^{q_2} F(x)dx}{K \cdot F(q_2)}, c\right) \tag{8}$$

The worst case for the cheapest seller would be for the buyer to discover the second cheapest seller first, where $K = N - 1$.

The expected information revenue per buyer is given by:

$$R'_{buyer} = \int_0^{\infty} f_2^{(N)}(q) I'(q) (\frac{1}{N} + \frac{1}{2}(1 - \frac{1}{N})) dq \tag{9}$$

which is equivalent to equation 6, using equation 8 for the pricing of price information.

The expected information revenue per product sold is due to the "multiplicative effect" equal to:

$$R'_{product} = N \cdot R'_{buyer} \tag{10}$$

One Nash equilibrium is for the cheapest seller to price according to equation 8 and all other sellers to pool, by charging the same price quote price.

The proof is along the same lines as for equation 5. Sellers other than the cheapest product seller cannot deviate as that would reveal to the buyer that they are not the cheap ones. The cheapest seller will not charge anything less than equation 8, as he would be reducing expected information revenue and he cannot charge more than equation 8 by any fixed amount k, as that would allow the buyer to conclude the price of the second cheapest seller in the market, and once that seller is discovered the price quote of the cheapest

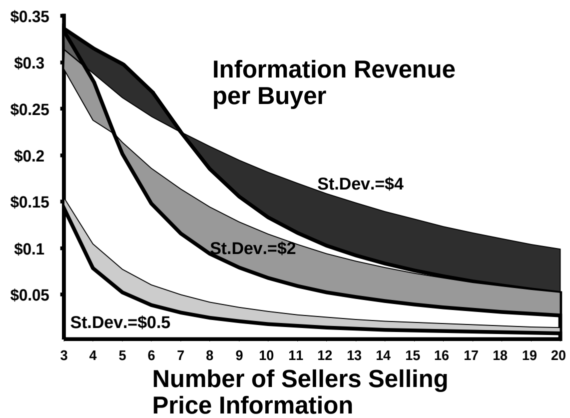

Figure 6: Average seller information revenue per buyer. The series represent different standard deviations of the assumed normal product price distribution, in dollars. A reservation price of 50 cents is assumed for all buyers. Thick lines represent revenue attainable when pricing is done with equation 8. Thin lines represent revenue attainable with equation 5.

seller, priced higher than equation 8 would be too expensive for the buyer to acquire. The buyer would be better off purchasing the second cheapest seller's product. The buyer herself has no other information to work with, except the price of the second cheapest seller. Even if that seller is discovered, it is still marginally optimal for the buyer to keep buying price quotes, as long as they are priced according to equation 8.

We have plotted R'_{buyer} and $R'_{product}$ in figures 6 and 7 respectively. Thick lines represent the equilibrium attainable with equation 8. Thin lines represent the equilibrium attainable with equation 5. Strictly speaking, thick lines represent an equilibrium when sellers are not allowed to randomize their price information pricing. Thin lines can be interpreted either as representing an equilibrium when the buyer cannot deduce that she has discovered the second lowest price, or as an upper bound of the revenue in an equilibrium where the cheapest seller subtracts a random quantity from equation 5. Notice that in some cases, the equilibrium revenue attainable with equation 8 is higher than the upper bound of the randomized equilibrium revenue. In those cases the sellers would always want to use equation 8 to price their price information. In the cases where the thin line is over the thick line, the sellers have the incentive to start using equation 5 with the randomization rule that we described.

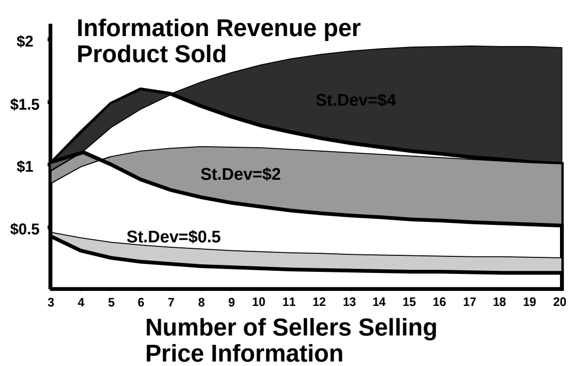

Figure 7: Average seller information revenue per product that the seller sells. Assumptions, same as figure 6.

4.4. The shopbot's problem

In this section we make predictions in accordance with our model about shopbot competition in delivering product information to potential buyers. We find that shopbot revenues reduce to marginal cost, while their ability to completely subsidize buyers' search costs depends on the price dispersion.

We use a simple utility model for the shopbot that assigns a fixed revenue S_r per shopbot user. This revenue could be ad banners revenue, a shopbot utilization fee, or some other form of revenue. For further simplicity, we assume that all competing shopbots have the same fixed revenue S_r per user and incur zero marginal costs. We ignore any service set up costs as sunk costs.

Prospective buyers are indifferent about which shopbot to use and they will visit the one that will minimize their expected costs (product plus search costs). We further assume that sellers sell their price information through the shopbots, pricing according to equation 8.

The total cost the buyer expects to incur is given by:

$$C_{buyer} = R'_{product} + \int_0^\infty x f_1^{(N)}(x)dx \tag{11}$$

where $R'_{product}$ is given by equation 10 and the second term is the expected minimum product price.

In our model the shopbots will engage in Bertrand competition, subsidizing the buyer's search costs down to $C_{buyer} - S_r$. Search will be free in markets with $S_r > R'_{product}$, where the shopbots can afford to subsidize the full cost of the buyer's search. For normal price distributions we can show that the buyers' expected search costs are strictly increasing with the normal distribution's standard deviation, and thus in markets with high product price dispersion shopbots are less likely to be able to completely subsidize the buyers' search costs.

Thus, by allowing shopbots to subsidize the buyers' search, our model predicts that shopbot competition will not affect seller information revenues but will rather drive shopbot revenues down to marginal cost.

Sellers of widgets at a Shopbot		
Seller number	**widget price**	**widget quality**
1	$8.96	Quality Info: 9 cents
2	$10.75	Quality Info: 11 cents
3	$11.20	Quality Info: 15 cents
4	$12.00	Quality Info: 5 cents
5	$12.00	Quality Info: 10 cents
6	$14.00	Quality Info: 12 cents
7	$14.99	Quality Info: 10 cents

Table 2: A hypothetical example of a seller charging for quality information through a shopbot

5. PRICING OTHER PRODUCT ATTRIBUTES

In the previous sections we dealt with the pricing of price information. We chose to do so in order to provide an objective measure of the value of product information. We showed that price information can be traded — with the help of software agents — and that the size of the market for price information can be a significant percentage of the product market itself. However, our analysis can be used to reason about the value and correct price of other product attributes, such as quality, delivery time, guaranties and special offers, to name a few. Or even the value of information for bundles of product attributes.

In this section we briefly describe the case where product quality information is traded between sellers and utility maximizing comparison shoppers. We use a simple utility model for a population of identical buyers.

Let the utility the buyers obtain by purchasing a product from seller i equal $U_b(q_i, p_i) = \alpha q_i - (1-\alpha)p_i$ where q_i is the quality of seller i's product and p_i is seller i's product price. The coefficient α controls the relative importance of price and quality on the buyer's valuation of the product and is taken to be the same for all buyers. All sellers draw their prices and qualities from independent distributions with PDF f_p and f_q respectively.

Assume that the user can visit a shopbot that displays all p_i's but the sellers demand a payment to release information on q_i. We assume, for simplicity, that the buyer has no other way of obtaining the quality information unless she pays for it. We further assume that the buyer will not purchase a product for which she has no quality information.

In table 2 we show a hypothetical example of sellers charging for quality information through a shopbot that displays price information for free.

If the most attractive product that the buyer has discovered has utility u, then the expected increase in utility from acquiring the quality information of seller i, that charges p_i for the product, is given by:

$$G(u, p_i) = \int_{\frac{u+(1-\alpha)p_i}{\alpha}}^{\infty} f_q(x)(\alpha x - (1-\alpha)p_i - u)dx \quad (12)$$

Equation 12 is the equivalent of equation 1 for our model of quality information value. It is simply the expected increase in utility integrated over all values of quality that actually result in utility increase.

We can show one Nash equilibrium to a simplified version of the quality information game that is defined as follows: All sellers draw their prices and qualities from the same distributions f_p and f_q respectively and observe each other's values. The seller with the highest utility moves first by setting a price for his quality information. All other sellers move simultaneously by setting prices for their quality information. The buyer, that can only observe product prices and qualities that are priced at zero, is then free to choose any utility maximizing strategy to search the product space. Again, a seller is assumed to be always better off selling his product than selling his information.

One Nash equilibrium is equivalent to the Nash equilibrium described in section 4.3. Given that the second highest product utility in the market is u_2, the highest utility buyer l sets his quality information price to $G(u_2, p_l)$ and all other sellers $j \neq l$ follow by choosing $G(u_2, p_j)$ for their quality information. Using similar arguments as in section 4.3, we can show that the highest utility seller would not choose a higher or lower price than this, no other seller would want to deviate by not charging $G(u_2, p_j)$ for his quality information, and that the buyer would then be randomly buying quality information until the highest utility is discovered, at which point all remaining quality information prices become too expensive to acquire. No player can thus deviate without lowering his expected utility.

To calculate a seller's expected information revenues per buyer we proceed as follows: Let f^P be the joined distribution of product prices, $f^P(p_1, \ldots, p_N) = f_p(p_1) \ldots f_p(p_N)$. Let $f_2^{p_1 \cdots p_N}(x)$ be the distribution of the second highest product utility, given that the N sellers charge $p_1, p_2, \ldots, p_N$ for their products[9]. The sellers' expected information revenues per buyer is:

$$R_{buyer} = \frac{N+1}{2N} \cdot$$
$$\cdot \int_0^{\infty} \ldots \int_0^{\infty} \int_0^{\infty} f^P f_2^{p_1 \cdots p_N}(x) G(x, p_1) dx dp_1 \ldots dp_N \quad (13)$$

This is equivalent to equation 6 presented in section 4.2. The seller's expected information revenue per product that the seller sells is due to the multiplicative effect:

$$R_{product} = N \cdot R_{buyer} \quad (14)$$

Due to the complexity of equation 13 we used Monte-Carlo simulations to plot the sellers' expected information revenue per buyer, and the sellers' expected information revenue, as a percentage of product sales revenue with a 95% confidence interval in figures 8 and 9 respectively.

[9] We refrain from providing a formula for this distribution as we will not be using it for calculations.

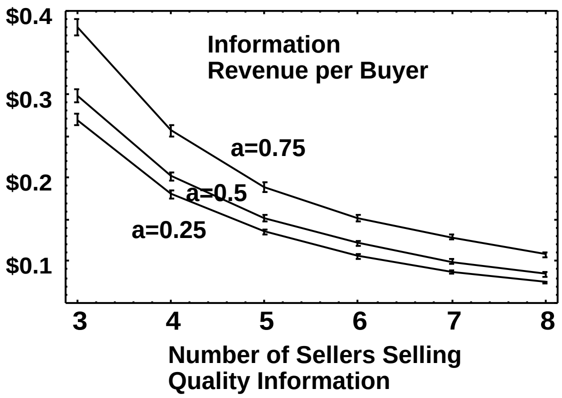

Figure 8: Average seller quality information revenue per buyer as a function of sellers in the market. Different series represent different values of the α constant that controls the relative importance of quality on the buyers' utility. For f_q and f_p we used the normal distribution with mean = 10 and standard deviation = 2.

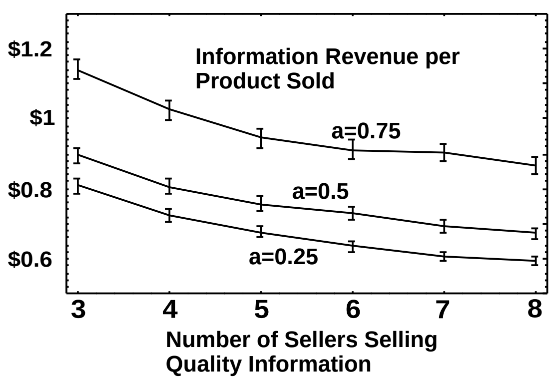

Figure 9: Average seller quality information revenue per product sold as a function of sellers in the market. Different series represent different values of the α constant that controls the relative importance of quality on the buyers' utility. For f_q and f_p we used the normal distribution with mean = 10 and standard deviation = 2.

In both figures, price and quality values are drawn from a normal distribution with mean 10 and standard deviation 2. Notice that figure 8 can also be interpreted as the buyer's expected search cost.

Again, we see the "multiplicative effect" of product information: while the average information revenue per buyer reduces about 66% before it starts to stabilize, as we increase the number of competing sellers from three to eight, the corresponding information revenue per product sold reduces only about 20%. It is also worth noting that the graphs are monotonically decreasing as we increase the number of sellers. This is because we have considered an infinite alternative search cost, as opposed to the 50 cents alternative search cost we have assumed in previous figures.

6. PRICE INFORMATION VALUE IN THE REAL WORLD

In the previous section we have calculated the maximum theoretical value of price information in an electronic market. We have found that, with today's observed price dispersion in Internet markets, this value represents a substantial additional revenue to the seller. The question that arises is how much of this theoretical value can actually be realized. Practical issues, such as the market power structure and the seller's ability to price discriminate and protect product price information from shopbots that do not pay for it, have great impact on the actual value that can be realized. We relax our model assumptions to address those issues.

6.1. Relaxing our model's assumptions

Our model assumes that sellers sell their product information to buyers that wish to easily compare across many different offers. Our intuition is that sellers will always maintain a degree of monopoly power regarding their own product information. After all, they are the ones that can choose their product offerings and guarantee that the information is accurate. In our model we were careful to price product information so that product sales are not affected: at equilibrium, the highest utility seller would eventually sell his product as he would have done if he had provided his information for free.

However, since it is conceivable that intermediaries will still be playing a role of finding and aggregating all this information we might expect that they will have a more active role in the information market. Depending on the degree of their monopoly power they might even be the ones that price and sell the information. Intermediaries would want to maximize their information revenues without considering any potential "distortions" in the market. By the term "distortion" we mean that the buyer might not choose to buy the same product before and after the introduction of an information market. This happens because the information regarding the best product offering might not be discovered, if pricing is not done according to the equilibria we described. The intermediaries would rather price price information using, for example, equation 4, which maximizes information revenue, without considering the effect on product sales revenue. The “multiplicative effect” of the value of product information would still be valid, since it is still true that a buyer would require multiple price quotes for a single purchase.

We have further assumed that the sellers would rather sell their products than information. This seems to be a reasonable assumption in most cases but it is also conceivable that (because of the “multiplicative effect”) product information might turn out to be more important than product revenue. In that case it is possible that the seller that offers the best product might be willing to risk not selling it, in return of increasing his information revenue.

In that case, it is likely that sellers would employ dual strategies: equations 5 or 8 for the cases where a product sale is always more important than information revenues and equation 4 otherwise.

6.2. Power structure - Sellers, Buyers and Intermediaries

In section 3.3 we saw that rational price sensitive buyers will abandon a shopbot that refuses to sell price information on behalf of a buyer. But real buyers may still prefer to use the shopbot with the smaller market coverage. This could be because they only know or trust that particular shopbot or because they value other services that the shopbot offers to them, like better interface or customization. In this case the buyer would lose sales because his product is not covered by the shopbot the buyers prefer. The seller could be forced to abandon product price selling or to share revenues with the shopbot. This would not be a problem if buyers used personalized agents to help them find a cheapest price. These agents would most likely not have “preferences” and their only goal would be to minimize search plus product costs for their owners.

A “popular” shopbot clearly enjoys market power that allows it to dictate the rules. In other markets the power group might be the buyers[10] and they would simply punish any effort by a seller or shopbot to charge for price information.

It is true that e-commerce has allowed the creation in some markets of powerful new intermediaries, such as Ebay, while reducing the power of intermediaries elsewhere. Traditionally, intermediaries that connected geographically dispersed markets enjoyed market power and high profit margins. In the electronic era, these intermediaries see their role diminished (disintermediation). Intermediaries in e-commerce will be probably serving an information need for their industry and geography with much lower margins than before. See [16] for an estimate of the value different players retain in B2B markets. Phillips and Meeker argue that the buyers are likely to be the group with real market power in environments with high seller concentration and similarly the sellers would enjoy power in markets with high buyer concentration. The intermediaries might only be able to dictate the rules in markets with fragmented buyers and sellers.

6.3. Practical and legal considerations

Sellers can theoretically use the information a shopbot provides to learn about their competitors prices and try to undercut them. Effectively, the sellers price discriminate by offering better prices to price sensitive shoppers. Before Books.com was acquired by BN.com it provided a free price-comparison service that allowed buyers to compare prices across some of the well known on-line bookstores. If the buyer discovered a cheaper price, Books.com would automatically undercut the cheapest competitor. The method we used to measure price information value in previous sections does not work in these settings, as we assumed that prices are constant. However it is not clear if sellers will be employing shopbots to undercut competitors in the future. It seems that if a competitor of Books.com were to employ the same methods then Books.com might have been forced to withdraw to fixed pricing.[11]

Before a true market for price information emerges, sellers would have to better understand how buyers use shopbots or shopping intermediaries. It appears that the common assumption that buyers use shopbots to locate a best price is only approximately correct: it is more plausible that buyers use shopbots to easily compare across a range of characteristics. In [19] the authors offer evidence that buyers regularly use shopbots to buy products from branded retailers that do not offer a lowest price. They show that buyers use shopbots to make an optimal price/brand decision, rather than locate a cheapest product. It is thus the case that other product characteristics, besides price and price dispersion, have to be taken into account when a seller decides to sell his price information.

Furthermore, it is not clear at present if an open web site has the right to block search engines and shopbots from accessing it. Recently Bidder’s Edge agreed to settle the case with Ebay, who claimed that its site was being “trespassed” by auction aggregators[12], by agreeing to stop displaying Ebay’s auction prices on its website without compensating

[10] An example would be oligopsonistic markets, where there are much fewer buyers than the potential sellers.

[11] In [11] the authors point out that this would also cause an endless pricing loop.

[12] As a side note, this also strengthens our belief that price information can be overused when provided for free.

Ebay. This controversial issue is currently far from settled as it is not legally clear if sellers maintain ownership over their product information, if it is available for free in their website.

However, it seems that at least technologically, the seller will be able to effectively forbid shopbots from accessing its site, by employing mechanisms that stop non-human visitors. For example, randomly changing the user interface from session to session would foil any attempts from shopbots to access web site information.[13]

To our knowledge, none of these practices has been tested in a legal dispute yet.

7. CONCLUSIONS AND FUTURE RESEARCH

We envision a future where buyers will be using software agents to compare product characteristics across a number of sellers. It would then be feasible for the sellers to charge buyer agents for the product information they offer. The complex space of consumer behavior and preferences will result in a range of complementary product offers which will in turn make product information valuable to the utility maximizing consumer. The result will be the emergence of product information markets, where information about products will be traded, rather than the actual products themselves.

In this paper we have focused mainly on the value of one product attribute, price, and showed how our model can be extended for the case of quality information. We showed that, because of the "multiplicative effect" of product information, the sellers' revenues from the product information market will be a significant percentage of the product market revenues. Arguably, in terms of profits, the importance of product information markets would only increase as it is cheaper to sell information rather than products.

The first seller to charge for price information can set the price information price so that all rational shoppers would be willing to pay for it. However, when more than one seller sells price information, it could be the case that a shopper would stop short of requesting a particular seller's price information. Pricing price information so that it would be cheap enough for the buyer to always want it (even if the second lowest price is known) is a solution that guarantees that the cheapest product seller will not jeopardize the product sale, while earning extra revenue from his information. We showed two plausible Nash equilibria in the price information selling game and proved that, because the sellers maintain a degree of monopoly power over their information, product information selling does not reduce to Bertrand competition.

We further showed that price comparison web sites will only be able to absorb the buyer search costs in markets with relatively small price dispersion, while in some markets that exhibit significant price dispersion, search will no longer be free.

Our model is also the first one to connect the "quality" of the information that the sellers provide to a monetary payment that can potentially induce better information. In our simple model better information "quality" simply meant lower price but it would be interesting to extend our ideas to settings where sellers have incentives to manipulate the information they provide or they find it too expensive to provide all the available information. In those cases product information markets can be used to induce better information and can potentially lead to higher information transparency and more efficient markets.

A product information market would be operating in parallel with the product market

8. ACKNOWLEDGEMENTS

The authors would like to thank Professor Ravi Aron for his valuable comments and suggestions on an initial draft of the paper.

This research was funded in part by National Science Foundation Grant SBR 96-02053.

9. APPENDIX

Sellers jointly maximizing their information revenues: The expected price information revenue per customer that uses a shopbot as a function of the price p the seller sells its price information is
$\Pi(p) = p[Prob($p is the lowest PIP[14]$)+Prob($p is the 2nd lowest PIP but the buyer still pays for it$) + \cdots + Prob($p is the Nth lowest PIP but the buyer still pays for it$)]$

When sellers collude, they choose the same price p for their price information. Thus, the probability that a seller is ith lowest but the buyer still pays for the seller's price information is $1/N \cdot Prob($The marginal expected decrease in product price is greater than $p)$

If q is the current minimum product price discovered then the marginal expected decrease in product price is given in section 2.1: $g(q) = \int_{-\infty}^{q}(q-x)f(x)dx$ and the distribution of q is (from section 2.3) $f_{min}^{i-1}(q) = (i-1)(1-F(q))^{i-2}f(q)$

So, for a given current minimum product price q the probability that the marginal expected decrease in product price is greater than p is

[13] A number of methods are used for this purpose currently, that require an "intelligent" action on the part of the user, like a randomly positioned "click here" button, or a "copy the letters from a figure" test.

[14] Price Information Price

$Q(p,i) = \int_p^\infty ((i-1)(1-F(q))^{i-2} f(q) \int_{-\infty}^q (q - x) f(x) dx) dq$, for $i \geq 2$

The expected price information revenue per customer is thus:

$$\Pi(p) = \frac{p}{N}(1 + \sum_{i=2}^{N} Q(p,i)) \qquad (15)$$

assuming that p is always less than c, the cost to the shopper of visiting a seller's website directly.

10. REFERENCES

[1] J. Yannis Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 43(12), 1997.

[2] Michael R. Baye and John Morgan. Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review*, 2001. Forthcoming.

[3] Michael R. Baye, John Morgan, and Patrick Scholten. Price dispersion in the small and in the large: Evidence from an internet price comparison site. working paper, 2002.

[4] Erik Brynjolfsson and Michael D. Smith. Frictionless commerce? a comparison of internet and conventional retailers. *Management Science*, 46(4), 2000.

[5] Lysbeth B. Chuck. Online shopping: Confusion, glut, overload, and misinformation. *Searcher*, 8(4):46–53, 2000.

[6] Karen Clay, Ramayya Krishnan, and Eric Wolf. Pricing strategies on the web: Evidence from the online book industry. In *Proceedings of the ACM Conference on Electronic Commerce, ACM EC'00*, 2000.

[7] S.H. Clearwater, R. Costanza, R. Dixon, and B. Schroeder. *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter Saving Energy Using Market-Based Control. World Scientific Publishing, Singapore, 1996.

[8] Amy R. Greenwald and Jeffrey O. Kephart. Shopbots and pricebots. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, 1999.

[9] Robert V. Hogg and Allen T. Craig. *Introduction to Mathematical Statistics*. The Macmillan Company, New York, 2 edition, 1965.

[10] Jeffrey O. Kephart and Amy R. Greenwald. *Shopbot Economics*, chapter forthcoming. Kluwer Academic Publishers, 2000.

[11] Jeffrey O. Kephart, James E. Hanson, and Amy R. Greenwald. Dynamic pricing by software agents. *Computer Networks: Special Issue on Trends and Research in E-Commerce, Elsevier Publishers*, 32(6):731–752, 2000.

[12] S. Lippman and J. McCall. The economics of uncertainty. selected topics and probabilistic methods. *Handbook of Mathematical Economics*, 2(6):213–225, 1982.

[13] Panos M. Markopoulos, Ravi Aron, and Lyle H. Ungar. Meta-prices for price information. working paper, 2002.

[14] Panos M. Markopoulos and Lyle H. Ungar. Extended abstract: Pricing price information in e-commerce. In *Proceedings of the ACM Conference on Electronic Commerce, ACM EC'01*, 2001.

[15] David Parkes and Lyle H. Ungar. An auction-based method for decentralized train scheduling. In *Proceedings of the 5th International Conference on Autonomous Agents, Agents'01*, 2001.

[16] Chuck Phillips and Mary Meeker. The b2b internet report. Available on line by Morgan Stanley Dean Witter at http://www.msdw.com/mrchuck/B2BReport.html, 2000.

[17] Michael Rothschild. Models of market organization with imperfect information: A survey. *The Journal of Political Economy*, 81(6):1283–1308, 1973.

[18] Michael D. Smith, Joseph Bailey, and Erik Brynjolfsson. *Understanding Digital Markets: Review and Assessment*, chapter forthcoming. MIT Press, 2001.

[19] Michael D. Smith and Erik Brynjolfsson. Consumer decision-making at an internet shopbot: Brand still matters. *The Journal of Industrial Economics*, 49(4):541–558, 2001.

[20] G. Stigler. The economics of information. *The Journal of Political Economy*, 69(6):213–225, 1961.

[21] Martin Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.

[22] Michael P. Wellman and Peter R. Wurman. Market-aware agents for a multiagent world. *Robotics and Autonomous Systems*, 24(3-4):115–125, 1998.