

Search More, Disclose Less

Chen Hajaj and Noam Hazon and David Sarne and Avshalom Elmalech

Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel

Abstract

The blooming of comparison shopping agents (CSAs) in recent years enables buyers in today's markets to query more than a single CSA while shopping, thus substantially expanding the list of sellers whose prices they obtain. From the individual CSA point of view, however, the multi-CSAs querying is definitely non-favorable as most of today's CSAs benefit depends on payments they receive from sellers upon transferring buyers to their websites (and making a purchase). The most straightforward way for the CSA to improve its competence is through spending more resources on getting more sellers' prices, potentially resulting in a more attractive "best price". In this paper we suggest a complementary approach that improves the attractiveness of the best price returned to the buyer without having to extend the CSAs' price database. This approach, which we term "selective price disclosure" relies on removing some of the prices known to the CSA from the list of results returned to the buyer. The advantage of this approach is in the ability to affect the buyer's beliefs regarding the probability of obtaining more attractive prices if querying additional CSAs. The paper presents two methods for choosing the subset of prices to be presented to a fully-rational buyer, attempting to overcome the computational complexity associated with evaluating all possible subsets. The effectiveness and efficiency of the methods are demonstrated using real data, collected from five CSAs for four products. Furthermore, since people are known to have an inherently bounded rationality, the two methods are also evaluated with human buyers, demonstrating that selective price-disclosing can be highly effective with people, however the subset of prices that needs to be used should be extracted in a different (and more simplistic) manner.

Introduction

Experienced shoppers know that the best way to make sure you are getting the best value for your money is to comparison shop before making a purchase. In today's online world, comparison shopping can be substantially facilitated through the use of commercial comparison shopping agents (CSAs) such as *PriceGrabber.com*, *bizrate.com* and *Shopper.com*. These web-based intelligent software applications allow comparing many online stores prices, saving buyers time and money (Pathak 2010).

The 17th annual release for ShoppingBots and Online Shopping Resources (*shoppingbots.info*) lists more than 350 different CSAs that are available online nowadays. According to the European Commission's market study from 2011¹, 81% of consumers in the EU use CSAs at least once a year (where 48% use them at least once a month).

The plethora of CSAs offering price comparison over the internet and the fact that each CSA covers only a small portion of the sellers offering a given product, suggest that prospective buyers may query more than a single CSA, aiming to find the best (i.e., minimal) price prior to making a purchase. This poses a great challenge to CSAs, as most of them do not charge consumers for accessing their sites and therefore the bulk of their profits is obtained, potentially alongside sponsored links or sponsored ads, via commercial relationships with the sellers they list (most commonly in the form of a fixed cost paid every time a consumer is referred to the seller's website from the CSA) (Moraga-Gonzalez and Wildenbeest 2012). Therefore, if a CSA could influence the probability a buyer will not continue querying additional CSAs, it would certainly improve its expected revenue. In the CSA-buyer setting, the buyer's decision of whether or not to resume exploration is based primarily on the best price obtained so far, her expectations regarding the prices that are likely to be obtained through further CSAs-querying, and the intrinsic cost of querying additional CSAs (e.g., cost of time). Influencing the best price presented to the the buyer can be achieved by increasing the size of the set of sellers whose prices are checked in response to the buyer's query. Yet, this requires consuming more resources and the expected marginal improvement in the best price decreases as a function of the size of the set.

In this paper we show that choosing not to disclose all the prices collected by the CSA can also be beneficial for the CSA, as this enables influencing the buyer's expectations regarding the prices she is likely to run into if querying additional CSAs. The underlying assumption is that the buyer is not a priori familiar with the market price distribution of the specific item she wants to buy, and her expectations are updated each time she obtains an additional set of prices from a queried CSA (Bikhchandani and Sharma 1996). An intelligent price-disclosure strategy can thus decrease the buyer's

¹http://ec.europa.eu/consumers/consumer_research/

confidence in finding a better price at the next CSAs queried and as a result discourage her from any additional querying.

The contributions of the paper to the study of selective price disclosing are threefold: First, we formally analyze the incentive of buyers to query additional CSAs and CSAs' benefit in selectively disclosing the prices they are acquainted with whenever queried. Evaluating the benefit in all subsets of the original set of prices is computationally intractable, therefore a second contribution is in presenting two price-disclosing methods CSAs can use, aimed to improve the probability that a buyer will terminate her price-search process once applied. Both methods disclose the minimum price known to the CSA, thus the benefit from the partial disclosure does not conflict with increasing the number of prices the CSA initially obtains for potentially finding a more appealing (lower) price. The effectiveness of the methods when the buyer is fully rational is evaluated using real data collected from five comparison shopping agents for four products. The evaluation demonstrates the effectiveness of the resulting subsets of prices achieved with these methods and the tradeoff between their performance and the time they are allowed to execute. Since both methods disclose the minimum price found, the increase in the probability the buyer will not query further CSAs directly maps to an increase in the probability she buys it from the CSA applying the selective price-disclosing. Finally, we evaluate the methods with people, finding out that what might be the best solution for fully-rational buyers is less effective with human buyers. This is partially explained by our experimental finding, according to which people's tendency to terminate their search increases as a function of the number of prices they obtain from the CSA, even if the minimum price remains the same. For the latter population we suggest a simple price disclosing heuristic that is demonstrated to be highly effective in keeping people from querying additional CSAs.

Related Work

The agent-based comparison-shopping domain has attracted the attention of researchers and market designers ever since the introduction of the first CSA (BargainFinder, (Krulwich 1996) (Decker, Sycara, and Williamson 1997; He, Jennings, and Leung 2003; Tan, Goh, and Teo 2010). CSAs were expected to reduce the search cost associated with obtaining price information, as they allow the buyer to query more sellers in the same time (and cost) of querying a seller directly (Bakos 1997; Wan, Menon, and Ramaprasad 2009; Pathak 2010). As such, the majority of CSA research is mostly concerned with analyzing the influence of CSAs on retailers' and consumers' behavior (Clay et al. 2002; Johnson et al. 2004; Karat, Blom, and Karat 2004; Xiao and Benbasat 2007) and the cost of obtaining information (Markopoulos and Ungar 2001; Markopoulos and Kephart 2002; Waldeck 2008).

Much emphasis has been placed on pricing behavior in the presence of CSAs (Pereira 2005; Tan, Goh, and Teo 2010), and in particular on the resulting price dispersion (Baye and Morgan 2006; Tang, Smith, and Montgomery 2010) in markets where buyers apply comparison-shopping. Substantial empirical research, mostly based on data from online books,

CDs and travel markets, has given evidence to the persistence of price dispersion (Clay et al. 2002; Brynjolfsson, Hu, and Smith 2003; Baye, Morgan, and Scholten 2004; Baye and Morgan 2006) in such markets. Other works focused on optimizing CSAs' performance, e.g., by better managing the resources they allocate for the different queries they receive (Sarne, Kraus, and Ito 2007).

Despite the many advances in applying optimal search theories for investigating search dynamics in markets where comparison-shopping principles are applied (Janssen and Moraga-Gonzalez 2004; Waldeck 2008), the absolute majority of the works assume that the CSA and user interests are identical and that the shopbot's sole purpose is to serve the buyer's needs (Markopoulos and Ungar 2002). Other works take the buyer to be the CSA entity (Varian 1980; Stahl 1989; Janssen, Moraga-Gonzalez, and Wildenbeest 2005), i.e., uses the most cost-effective search strategy for minimizing the buyer's overall expense. Naturally, in such cases, the existence of CSAs improves the buyers' performance, resulting in a lower benefit to sellers (Gorman, Salisbury, and Brannon 2009; Nermuth et al. 2009). Those few works that do assume that the CSAs are self-interested autonomous entities (Kephart, Hanson, and Greenwald 2000; Kephart and Greenwald 2002) focus on CSAs that charge buyers (rather than sellers as in today's markets (Wan and Peng 2010)) for their services.

Using smart disclosure in order to lead to a preferred behavior was previously investigated by (Sarne et al. 2011; Azaria et al. 2012). Sarne et al. use manipulation techniques for improving peer designed agents' (PDAs) exploration in order to guide searchers to a strategy that is closer to optimal. Our work's aim is to optimize the opportunity's (i.e., the CSA's) benefit. Moreover, while they deal with PDAs our work deals with fully-rational agents and people. Azaria et al. focus on how automated agents can persuade people to behave in certain ways. The authors assumed that people were aware of the preferences of a central mechanism where in our work the CSA sets his disclosed set of sellers based only on the sampled information without any a-priori knowledge. To the best of our knowledge, the advantage of displaying only subset of the prices known to the CSA has not been researched to date.

Model and Individual Strategies

We consider an on-line shopping environment with numerous buyers (denoted "searchers" onwards), sellers and several comparison shopping agents (CSAs). Sellers' prices are assumed to be associated with a probability distribution function $f(y)$ where y is a possible price that a seller can offer. This assumption is commonly used in e-commerce research (Janssen, Moraga-Gonzalez, and Wildenbeest 2005; Waldeck 2008; Tang, Smith, and Montgomery 2010) and as discussed in the former section is also supported by empirical research in well-established online markets. CSAs are assumed to be self-interested fully-rational agents, aiming to maximize their own net benefit. Once queried by a searcher, a CSA will supply a set of prices in which the requested item can be purchased at different online stores. The number of prices returned by the CSA is a priori unknown to

the searcher and there is no CSA that generally returns more prices than another (Serenko and Hayes 2010). The searcher is assumed, however, to be acquainted with the average number of sellers listed in CSAs responses for a given product, denoted N . The model assumes CSAs do not charge buyers for their services, but rather receive a payment from sellers every time a searcher that was referred to their web-site by the CSA executes a transaction, as the common practice in today's CSAs (Wan and Peng 2010).

A searcher interested in buying a product can either query sellers directly or use CSAs for that purpose. We assume that querying either a seller or a CSA incurs a cost c_{query} (e.g., in the case of people the cost of the time it takes to get to the appropriate website, specify the product of interest, as well as any other required complementary information, and waiting for the results, and in the case of agents the computational and communication costs). Since the cost incurred is equal in both cases, however CSAs return more than a single price quote, it is always beneficial to query CSAs. Based on the price quotes received along her exploration, the searcher needs to decide on each step of her search process whether to terminate her exploration and buy at the best (minimum) price found so far, or query another CSA. The model assumes that searchers execute their price-search on an ad-hoc basis and therefore they are unfamiliar with the distribution function $f(y)$. Instead, they learn the distribution of prices as they move along, based on their queries (Bikhchandani and Sharma 1996). Searchers are assumed to be interested in minimizing their overall expected expense, i.e., the sum of the minimum price they obtain eventually and the costs incurred along the search.

Since the searcher is interested in minimizing her expected overall expense, her state can be represented by the tuple (q, w) where q is the best price revealed so far throughout her search and w is the number of CSAs queried so far. The importance of the latter parameter is because it affects the expected number of redundant results returned by the next queried CSA (i.e., the same price from a seller formerly listed in the results of one of the other queried CSAs). We use $N(w)$ to denote the expected number of "new" results the searcher will obtain from the next CSA she queries, given that she has already queried w CSAs.

Given a state (q, w) , we define the *critical cost*, denoted $c_{critical}$, as the querying cost for which the searcher is indifferent between querying an additional CSA (hence benefiting from the potential improvement to q) and terminating exploration thus saving the additional exploration cost. The critical cost can therefore be calculated as:

$$c_{critical} = \int_{y=0}^q (q - y) f_{N(w)}(y) dy. \quad (1)$$

where $f_{N(w)}(y)$ is the probability distribution of the minimum price among the $N(w)$ new listings in the next CSA's output. The function $f_{N(w)}(y)$ is calculated as the derivative of the probability that the minimum price is equal to or

lesser than y , i.e.:

$$\begin{aligned} f_{N(w)}(y) &= \frac{\partial [1 - (1 - F(y))^{N(w)}]}{\partial y} \\ &= N(w) f(y) (1 - F(y))^{N(w)-1} \end{aligned} \quad (2)$$

If the next CSA presents $N(w)$ prices, then the probability that the minimum price is less than y is $1 - (1 - F(y))^{N(w)}$ (where $(1 - F(y))^{N(w)}$ is the probability all $N(w)$ prices are above y , we are interested in the complementing probability). The derivative of this probability is the p.d.f. of the minimum price. Therefore, the searcher will always prefer querying the next CSA whenever $c_{query} < c_{critical}$.

From the queried CSA's point of view, the condition for generating revenue (i.e., having the searcher buy through that CSA) is that the price it returns is the minimum among the prices obtained throughout the searcher's search. Since the searcher cannot distinguish a priori between the number of prices returned by the different CSAs, the CSA cannot affect the probability that the searcher will get to it earlier or later throughout its search nor the distribution of the minimum price obtained until reaching it. Consider a CSA with the set $Q = \{q_1, \dots, q_n\}$ of available prices. Querying more sellers can potentially result in finding a better best price, hence increasing the probability the searcher will terminate her exploration and/or ending-up buying from that CSA. Nevertheless, this option requires allocating further resources, and can potentially deny service from other prospective searchers. Alternatively, the CSA can choose to disclose only a subset $Q' \subset Q$, attempting to influence the searcher's beliefs concerning the distribution of prices hence discouraging further exploration. Obviously, the CSA will always prefer disclosing the minimum value found, as this is the only decision parameter for the searcher for choosing the CSA to buy through, once deciding to terminate her exploration.²

Naturally, when a CSA decides to disclose only a subset of the prices that were queried, it needs to preserve a minimal number of prices (ρ), otherwise the CSA would seem unreliable and its reputation will suffer. Moreover, in the case where the CSA is the first to be queried by the user, supplying a small subset of prices will preclude an actual estimation of the distribution of prices and will not allow a decision based on the principle given in (1).

Methods

The number of CSAs searchers are likely to query is not large. For example, a recent consumer intelligence report (Knight 2010) reveals that the average number of CSAs visited by motor insurance switchers in 2009 was 2.14. Therefore, our price-disclosing methods primarily apply to the case where the CSA is the first to be queried by the searcher, and the goal is to minimize the probability that the searcher will decide to query another CSA. This has also many theoretical-based justifications: By querying another

²Also, this one price has a small influence over the distribution perceived by the searcher as a whole.

CSA other than the first, the probability of making a purchase through the first drops from 1 to less than 0.5, as the searcher may continue and query additional CSAs, and since every CSA query the same number of sellers (give or take) the probability of each of them to be the one associated with the minimum price is equal. If the CSA is not the first to be queried, i.e., if it is the k th to be queried then an upper bound for the benefit from partial disclosing of prices is an improvement $1/k - 1/(k+1) = 1/k(k+1)$ in the probability that the agent will be the one through which the product will be purchased (e.g., if the agent is the fourth to be queried then the maximum improvement is 5%). The actual improvement that can be achieved is far less than $1/k(k+1)$ since after querying the k th ($k > 1$) CSA the probability the searcher will query further CSAs substantially decreases anyhow compared to the case of $k = 1$. This is because the probability of having a price that is good (low) enough, such that an additional costly CSA query is not justified, increases as k increases. Furthermore, the chance of running into sellers which prices have already been supplied within the query of the former k CSAs increases as k increase, therefore the benefit in exploring the $(k+1)$ th seller decreases. Finally, since the searcher's beliefs concerning the distribution of prices are based on all prices obtained from the CSAs queried so far, the extent of effect the partial price disclosing will have on the searcher's belief of the actual distribution of prices is limited whenever the CSA is not the first to be queried.

By increasing the termination probability from p to p' the increase in a CSA's revenue is bounded in the interval $[(p' - p)/2, p' - p]$ (since on the extreme case of querying an infinite number of CSAs the increase in revenue is $(p' - p)M$, where M is the amount it charges sellers, and on the other extreme when only one additional CSA will be queried the increase in revenue is $(p' - p)M/2$). Based on the observed distribution $f(y)$, a fully-rational searcher's decision of whether or not to query an additional CSA depends on the relation between c_{query} and $c_{critical}$. As noted above, if $c_{query} \geq c_{critical}$ the searcher will terminate the search. Since the CSA does not know the value of c_{query} for each searcher, the CSA cannot estimate the improvement that achieved in the termination probability when using the different methods. Instead, it can measure the reduction in $c_{critical}$. The lower the value of $c_{critical}$ is, the lesser the number of searchers that will decide to query additional CSAs.

In order to find the set of prices that yields the minimal $c_{critical}$, the CSA can theoretically check all combinations of $\rho \leq k < n$ subset of prices. Since the minimal price must be included, the number of combinations to check is $\sum_{k=\rho-1}^{n-1} \binom{n-1}{k}$. For example, if the CSA sampled $n = 30$ prices and the minimum number of prices (denoted ρ) is $\rho = 10$ then it needs to check 530 million combinations. For each such combination, there is a need to estimate the distribution of prices $f(y)$ based on the subset's prices and calculate the critical cost. Obviously, this method is infeasible. Today's e-commerce is characterized by quick interactions, and a price-disclosing method should return a result within seconds or milliseconds. Even a pre-processing step

will not do much, since sellers change their prices quite often, and any change in price might have a large effect on the critical cost. We therefore propose two heuristic methods for choosing a subset of prices to disclose, Monte-Carlo based disclosing and Interval disclosing.

Monte-Carlo based Disclosing This method randomly samples different subsets of prices. At first, the CSA chooses a random number of prices $\rho \leq k < n$ to disclose to the searchers. Then it randomly chooses a set of $k - 1$ prices out the $n - 1$ known prices (since the minimum price is necessarily part of the subset that will be returned to the searcher), estimates the probability distribution $f(y)$ based on this subset, and calculates the critical cost. This process is repeated as long as the CSA is able to hold its response to the searcher. We thus get an anytime algorithm, since the greater the number of subsets that can be sampled, the lower the critical cost that will be achieved.

Interval Disclosing This method attempts to make use of the unique properties of the calculation of $c_{critical}$. It iterates over all the possible sizes of the sets of prices that can potentially be disclosed to the searcher, i.e., $\rho \leq k < n$. For each size k , it chooses an interval of prices (i.e., a sequence of consecutive prices) in the size of $k - 1$ (since the minimum price is necessarily disclosed), estimates the probability distribution $f(y)$, and calculates the critical cost. The required number of subsets to evaluate is therefore $\frac{(n-\rho+1)*(n-\rho+2)}{2}$. The rationale behind this method is quite simple: if many prices are concentrated within a small interval, then regardless of the distribution estimation method this interval and its surrounding are likely to be assigned with a substantial distribution mass. Consequentially, other intervals are likely to be assigned with small distribution masses. In particular, the values of $f(y)$ within the interval $(0, q)$, over which $\int_{y=0}^q (q-y)f_{N(w)}(y)dy$ in (1) is defined, are likely to be low, which yields a small critical cost.

In order to evaluate the above methods, we used the data of 4 products: Logitech Keyboard & Mouse, HP LaserJet Pro 400, HP 2311x screen and Sony WX50 camera, and sampled their prices from 5 well-known CSAs: PriceGrabber.com, Nextag.com, Bizrate.com, Amazon.com and Shopper.com. As observed by others in general (Baye and Morgan 2006; Serenko and Hayes 2010), none of these CSAs returns more prices than another for all four products, and there was not any significant difference between the number of sellers that each CSA presented. In order to estimate $N(1)$, i.e., the number of "new" prices the searcher is likely to obtain if querying a second CSA, we calculated the number of overlapping results between any two new CSAs for each product, resulting in average overlap of 12%. Applying this on the average number of sellers listed in the five CSAs resulted in $N(1)=18$. The distributions were estimated using the kernel density estimation method (KDE) (also called Parzen-Rosenblatt window estimation (Parzen 1962)), which is a non-parametric method to estimate the probability density function of a random variable. The estimation is based on a normal kernel function, using a window width that is a func-

tion of the number of samples.

Figure 1(a) depicts the performance of our methods in terms of the critical cost as a function of the number of subsets evaluated for HP LaserJet Pro 400. The initial number of prices available for the CSA is 30 prices, which were derived from the empirical distribution for that product with equal probability mass between any two consecutive prices. Finally, we use $N(1) = 18$ and $\rho = 10$. Each data point represents the average over 5,000 simulation runs. The figure includes also the critical cost of the original set of 30 prices, as a reference. As can be observed from the figure, both methods substantially improve the critical cost, where the improvement with the Interval method is achieved with fewer set-evaluations. Since $n = 30$ and $\rho = 10$, the Interval method’s performance becomes steady once it completes the evaluation of the 231 applicable continuous sets of prices, as no further sets need to be evaluated. Obviously, if having the option to evaluate a large enough set of subsets, the Monte-Carlo method is supposed to yield at least as good results as the Interval method (as it becomes close to brute force). Yet, even when we extended the simulation to 100,000 subsets, the Monte-Carlo method did not manage to outperform the Interval method on average, and the average critical cost achieved by the Interval method was better by 7.78%.

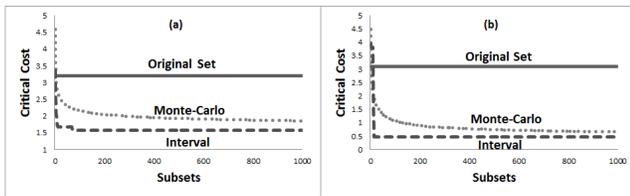


Figure 1: Critical cost as a function of the number of evaluated subsets: (a) KDE; (b) choosing among 17 fittings.

In order to demonstrate that these results do not qualitatively depend on the estimation method, we repeated the simulation with a different estimation method. According to the new method we try to fit the data to 17 parametric probability distributions: Beta, Birnbaum-Saunders, Exponential, Extreme value, Gamma, Generalized extreme value, Generalized Pareto, Inverse Gaussian, Logistic, Log-logistic, Log-normal, Nakagami, Normal, Rayleigh, Rician, t location-scale, and Weibull. Based on the fitting results, we then choose the best distribution according to the Bayesian information criterion (Schwarz 1978). The result of the Monte-Carlo and Interval based methods when used with the above distribution estimation method are given in Figure 1(b). As observed from the figure, the methods exhibit a similar behavior even with the new distributions estimation method.

Similar analysis with the other three products reveals similar patterns thus was omitted from this section. Based on the results we conclude that the Interval method is highly effective with fully rational searchers and results in excellent performance (i.e., low critical cost) while requiring a relatively short running time.

Evaluation with People

While the methods described above are highly effective with fully rational agents, searchers in today’s markets are usually human, and it is well known that people do not always make rational decisions (Baumeister 2003). In particular, people often follow rules of thumb and tend to simplify the information they encounter. For example in our online shopping setting, people may ignore the high-range prices, as they are likely not to buy in those prices anyhow, rather than use them as part of the distribution modeling (Ellison and Ellison 2009), or may be effected by other psychological properties (Rao and Monroe 1989). In this section we report the results of an experimental evaluation of the Monte-Carlo based and Interval based disclosing heuristics when applied on human searchers. In addition we report the results of two complementary experiments. The first aims to evaluate the correlation between the number of results presented by the CSA and the (human) searcher’s tendency to query an additional CSA, partially explaining the findings related to the effectiveness of the disclosing heuristics with people. The second aims to evaluate a third selective price-disclosing heuristic which is more suitable for the case of human searchers.

Experimental Design

The experimental infrastructure developed for the experiments with people is a web-based application that emulates an online CSA’s website. Participants were recruited using Amazon Mechanical Turk (AMT). Once accessing the website, the participant obtained a list of sellers and their appropriate prices for a well-defined product (see a screenshot in Figure 2). The list is given (just as in all real CSAs) in an ascending order according to price, thus it is easy to identify the best price in the list or reason about the distribution of prices. At this point, the participant is awarded her show-up fee (i.e., the “hit” promised in Mechanical Turk) and a bonus of a few cents. The participant is offered to give up the bonus, in exchange to sampling $N(1)$ additional prices. If the second set of prices that will be obtained will include a better price, then the user will obtain the difference (i.e., the saving due to the better price) as a bonus. Therefore, each participant faced the same tradeoff captured by querying an additional CSA, where the bonus it needs to give up on is the equivalent to the search cost (e.g., the time it takes to query the additional CSA) and the alternative bonus in the form of the improvement achieved in the best price obtained is the saving on the product cost obtained from querying the additional CSA.

In order to adequately set the initial bonus participants were offered (i.e., the equivalent to the search cost) we experimentally measured the time it takes a common user to query a CSA. For this purpose we asked 30 undergraduate engineering students to browse to PriceGrabber.com and return the minimal price for a Brother HL-2240 printer. On average, this took 60.9 seconds. Since we use AMT as our main test bed, and the average hourly salary for a worker in AMT is \$4.8 (Ipeirotis 2010), we set the initial bonus accordingly to 8 cents.

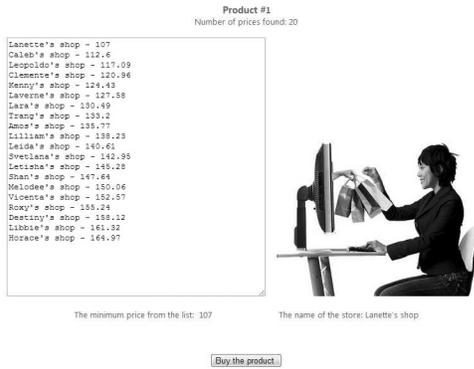


Figure 2: The first stage of the experiment

The price data used for the experiments with people was the same real data that we have used to evaluate the Monte-Carlo and Interval sampling with fully rational agents as detailed in the former section. Each scenario that we generated contained the minimal price as well as other prices, either the original ones or a subset according to the tested method. The number of participants used for experimenting with each setting ranged between 40 – 101 in order to obtain the appropriate statistical significance.

Experimental Results

We started by testing whether the termination probability increases as a function of the number of sellers that the CSA presents. For this purpose, we extracted the distribution of prices for each of the four products using KDE, based on the real set of prices listed by the different CSAs. Then, we generated four subsets of 10, 20, 30 and 40 prices where on each subset the minimum price is the minimum in the original set and the remaining prices are generated in a way that divide the distribution function to equal probability mass intervals (i.e., the i th price was selected such that $F(q_i) - F(q_{i-1}) = 1/k$, where k is the number of prices in the subset and q_0 is the minimum price in the original set). This way, all four subsets for the same product, although containing different prices, similarly represented the same price distribution and had the same minimum price. For each subset of each product (i.e., a total of 16 subsets), we had different subjects offered to obtain an additional sample of $N(1)$ prices in exchange to the initial bonus, promising to pay a sum equal to the improvement in the best price if a better price will be obtained in the new sample. Figure 3 summarizes the results of this experiment, depicting the percentage of participants that chose to terminate the search and avoid querying another CSA in each setting, i.e, the termination probability. As expected, the termination probability is monotonically increasing as a function of the number of prices displayed up to a certain point (40 prices). For all four products the difference between 10 and 30 prices was found to be statistically significant ($p < 0.05$), despite the fact that neither the underlying distribution of prices nor the minimum price displayed have changed. The performance degradation with 40 products can be explained by prior work where it was shown that listing too many options in a CSA's

results leads to lower-quality choices, decreases the selectivity with which consumers process options (Diehl 2005).

The effect of the number of prices displayed per-se is unique to people, as fully rational searchers' CSA-querying decisions are only affected by the resulting estimated probability and the minimal price. The fact that with less prices displayed the tendency of people to query additional CSAs substantially increases poses a great challenge to our selective price disclosure approach, which essentially reduce the number of prices listed as a response to the searcher's query. As we show in the following paragraphs, this effect definitely affects the result of the methods proposed for fully-rational searchers when applied with people. Yet, even with human searchers, an effective selective price-disclosure heuristic can be designed.

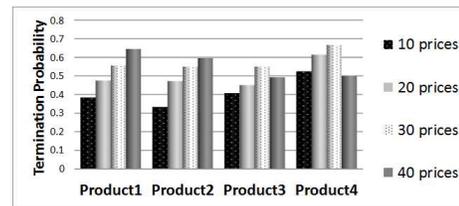


Figure 3: Termination probability with different sets sizes.

To test the performance of Interval and Monte-carlo based sampling with people, we fixed the number of prices to 30, using the same prices that were generated for the experiment summarized by Figure 3. This choice of the number of prices to be presented favors full price disclosure, as it was found to improve people's termination probability compared to any lower number of prices in all four products. Therefore it is likely to be more challenging for the price disclosure approach to present an improvement. The number of subsets evaluated with the Monte-Carlo sampling method in this experiments was set to 10,000. Figure 4(a) summarizes the results of applying the Interval and Monte-Carlo based sampling on the set of 30 prices, depicting the percentage of participants that chose to terminate the exploration with the use of each method. Here, again, we had different subjects presented with the prices of each setting. From the figure we observe that both methods did not succeed in increasing the termination probability (with appropriate statistical significance), compared to full disclosure.

One possible explanation for the failure of the methods with people is that people are highly affected by the number of prices they are presented with, as evidenced in Figure 3. Therefore, a second set of experiments was carried out, this time however, constraining the number of prices the method must disclose to 10 and 20. The results of the Interval sampling and the Monte-Carlo based methods in this case are depicted in Figures 4(b) and 4(c), respectively. From the figures we cannot observe any consistent behavior. Indeed with the 20 prices constraint the performance of both methods were improved (except for the non-significant case of product 2 with the intervals method). Yet, the performance with the 20 prices constraint are not generally better than those achieved with the 10 prices constraint. A possible explanation for the failure of the Interval sampling method with peo-

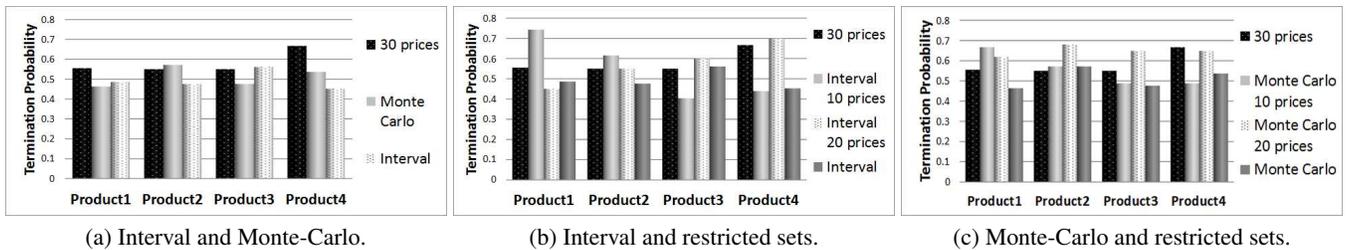


Figure 4: termination probability with Interval and Monte-Carlo, with restricted sets.

ple is that it produces price sets with a large gap between the minimum price and the rest of the prices. Possibly, this makes human searchers believe that there are lower prices that the CSA failed to query, therefore encouraging an additional CSA exploration. As for the Monte-Carlo based sampling, here since any random subset selection is allowed, it is more likely that the subset that will be eventually picked by the method will be one that implies a complex distribution function, which is more difficult for people to fit to.

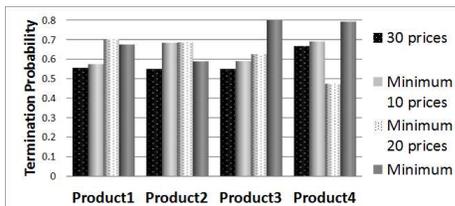


Figure 5: Termination probability with the minimum prices method.

Therefore, we suggest an additional selective price disclosing method that is more suitable for human searchers. The new method is a variation of the Intervals method that takes the subset of the lowest k prices (where the value of k is set according to the critical cost calculation). The method thus requires evaluating only $n - \rho + 1$ subsets. The results of the new method when tested with people are depicted in Figure 5 alongside the performance of the method when constraint to disclose 10 and 20 prices and when using full disclosure. From the figure we observe that the new method managed to improve the termination probability compared to full disclosure in all four products.

To summarize, the empirical results obtained in our experiments with people show that CSAs should act differently when dealing with fully rational agents and human searchers. Moreover, people’s decision to terminate their search is affected by the number of prices that are presented by the CSA. Even though, we show that with a simplistic selection rule for the prices to be disclosed, a substantial improvement can be achieved in the termination probability.

Discussion and Conclusions

The encouraging results reported in the former two sections support the hypothesis that selective disclosure of findings enables a CSA to substantially discourage people from

querying additional CSAs, thus improving its expected revenue. As discussed in the introduction, the method does not conflict with the initial tendency to increase the number of sellers the CSA queries, depending on the available resources. Thus it is suggested that the CSA will obtain the price of as many sellers as possible, benefiting from the potential decrease in the expected minimum price found, and then disclose a subset of the remaining prices using the methods presented in this paper, depending on whether the searcher is a fully rational agent or a person. The methods presented in the paper for selecting the subset of prices to be delisted are characterized with a polynomial computational complexity and are demonstrated to be effective using real data.

Our empirical findings related to the differences between the effectiveness of the different price-disclosing methods when applied with human and fully rational agents are not surprising. Prior research in other domains has provided much evidence for the benefit in being able to distinguish between these two populations in mechanism design. Still, the results reported in the previous section make several important contributions. First, they provide a simple heuristic for selective-price disclosure that substantially improves CSA’s performance with people and requires minimal computation. Second, we empirically show that for the typical range of the number of prices CSAs present nowadays, presenting more prices is generally more beneficial. This latter result strengthens the significance of the price-disclosing idea, as it suggests that the improvement achieved in people’s tendency to terminate their search is way greater than the inherent resulting discouragement they experience due to the decrease in the number of listings they receive from the CSA.

We see various directions for future research evolving from the results given in this paper, among which a more detailed investigation of the source of difference in the decision to resume exploration between agents, people (and also possibly bounded rational agents that were developed by people). Another interesting direction would be the integration of complementary considerations into the selection of the subset of prices to be disclosed, e.g., additional preferences the searcher may have (other than price).

Acknowledgment

This work was partially supported by BSF grant 2008-404.

References

- Azaria, A.; Rabinovich, Z.; Kraus, S.; Goldman, C.; and Gal, Y. 2012. Strategic advice provision in repeated human-agent interactions. In *Proc. of National Conference on Artificial Intelligence (AAAI)*.
- Bakos, J. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Mng. Sci.* 43(12):1676–1692.
- Baumeister, R. 2003. The psychology of irrationality: Why people make foolish, self-defeating choices. In Brocas, I., and Carrillo, J., eds., *The Psychology of Economic Decisions: Rationality and Well-Being*. Oxford University Press.
- Baye, M. R., and Morgan, J. 2006. Persistent price dispersion in online markets. In Jansen, D., ed., *The New Economy and Beyond*. Edward Elgar Press.
- Baye, M. R.; Morgan, J.; and Scholten, P. 2004. Temporal price dispersion: Evidence from an online consumer electronics market. *Journal of Interactive Marketing* 18(4):101–115.
- Bikhchandani, S., and Sharma, S. 1996. Optimal search with learning. *Journal of Economic Dynamics and Control* 20(1):333–359.
- Brynjolfsson, E.; Hu, Y. J.; and Smith, M. D. 2003. Consumer surplus in the digital economy: Estimating the value of increased product variety at online bookseller. *Management Science* 49(11):1580–1596.
- Clay, K.; Krishnan, R.; Wolff, E.; and Fernandes, D. 2002. Retail strategies on the web: Price and non-price competition in the online book industry. *Journal of Industrial Economics* 50:351–367.
- Decker, K.; Sycara, K.; and Williamson, M. 1997. Middle-agents for the internet. In *Proc. of IJCAI*, 578583.
- Diehl, K. 2005. When two rights make a wrong: Searching too much in ordered environments. *Journal of marketing research*.
- Ellison, G., and Ellison, S. 2009. Search, obfuscation, and price elasticities on the internet. *Econometrica* 77(2):427–452.
- Gorman, M. F.; Salisbury, D.; and Brannon, I. 2009. Who wins when price information is more ubiquitous? an experiment to assess how infomediaries influence price. *Electronic Markets* 19(2-3):151–162.
- He, M.; Jennings, N. R.; and Leung, H. 2003. On agent-mediated electronic commerce. *IEEE Transaction on Knowledge and Data Engineering* 15(4):985–1003.
- Ipeirotis, P. 2010. Analyzing the amazon mechanical turk marketplace. *XRDS: Crossroads, The ACM Magazine for Students* 17(2):16–21.
- Janssen, M. C. W., and Moraga-Gonzalez, J. L. 2004. Strategic pricing, consumer search and the number of firms. *Review of Economic Studies* 71(4):1089–1118.
- Janssen, M. C. W.; Moraga-Gonzalez, J. L.; and Wildenbeest, M. R. 2005. Truly costly sequential search and oligopolistic pricing. *International Journal of Industrial Organization* 23(5-6):451–466.
- Johnson, E. J.; Moe, W. W.; Fader, P. S.; Bellman, S.; and Lohse, G. L. 2004. On the depth and dynamics of online search behavior. *Management Science* 50(3):299–308.
- Karat, C. M.; Blom, J. O.; and Karat, J. 2004. *Designing personalized user experiences in eCommerce*. Human-Computer Interaction Series. Kluwer Academic.
- Kephart, J. O., and Greenwald, A. R. 2002. Shopbot economics. *JAAMAS* 5(3):255–287.
- Kephart, J. O.; Hanson, J. E.; and Greenwald, A. R. 2000. Dynamic pricing by software agents. *Computer Networks* 32:731–752.
- Knight, E. 2010. The Use of Price Comparison Sites in the UK General Insurance Market.
- Krulwich, B. 1996. The bargainfinder agent: Comparison price shopping on the internet. In Williams, J., ed., *Bots and Other Internet Beasts*. Sams Publishing. chapter 13, 257–263.
- Markopoulos, P., and Kephart, J. 2002. How valuable are shopbots? In *Proc. of AAMAS*, 1009–1016. ACM.
- Markopoulos, P., and Ungar, L. 2001. Pricing price information in e-commerce. In *Proc. of ACM-EC*, 260–263.
- Markopoulos, P., and Ungar, L. 2002. Shopbots and pricebots in electronic service markets. *Game Theory and Decision Theory in Agent-Based Systems* 177–195.
- Moraga-Gonzalez, J. L., and Wildenbeest, M. 2012. Comparison sites. In Peitz, M., and Waldfogel, J., eds., *The Oxford Handbook of the Digital Economy*. Oxford University Press.
- Nermuth, M.; Pasini, G.; Pin, P.; and Weidenholzer, S. 2009. Price dispersion, search externalities, and the digital divide. Vienna Economics Papers 0916, Department of Economics.
- Parzen, E. 1962. On estimation of a probability density function and mode. *The annals of mathematical statistics* 33(3):1065–1076.
- Pathak, B. 2010. A survey of the comparison shopping agent-based decisions support systems. *Journal of Electronic Commerce Research* 11(3):177–192.
- Pereira, P. 2005. Do lower search costs reduce prices and price dispersion? *Information Economics and Policy* 17(1):61–72.
- Rao, A., and Monroe, K. 1989. The effect of price, brand name, and store name on buyers' perceptions of product quality: An integrative review. *Journal of marketing Research* 351–357.
- Sarne, D.; Elmalech, A.; Grosz, B. J.; and Geva, M. 2011. Less is more: restructuring decisions to improve agent search. In *Proc. of AAMAS*, 431–438. International Foundation for Autonomous Agents and Multiagent Systems.
- Sarne, D.; Kraus, S.; and Ito, T. 2007. Scaling-up shopbots: a dynamic allocation-based approach. In *Proc. of AAMAS*, 256–263.
- Schwarz, G. 1978. Estimating the dimension of a model. *The annals of statistics* 6(2):461–464.
- Serenko, A., and Hayes, J. 2010. Investigating the functionality and performance of online shopping bots for electronic commerce: a follow-up study. *International Journal of Electronic Business* 8(1):1–15.
- Stahl, D. O. 1989. Oligopolistic pricing with sequential consumer search. *American Economic Review* 79(4):700–712.
- Tan, C. H.; Goh, K. Y.; and Teo, H. H. 2010. Effects of comparison shopping websites on market performance: Does market structure matter? *Journal of Electronic Commerce Research* 11(3):193–219.
- Tang, Z.; Smith, M. D.; and Montgomery, A. 2010. The impact of shopbot use on prices and price dispersion: Evidence from online book retailing. *International Journal of Industrial Organization* 28(6):579–590.
- Varian, H. R. 1980. A model of sales. *American Economic Review* 70(4):651–59.
- Waldeck, R. 2008. Search and price competition. *Journal of Economic Behavior and Organization* 66(2):347–357.
- Wan, Y., and Peng, G. 2010. What's next for shopbots? *IEEE Computer* 43:20–26.
- Wan, Y.; Menon, S.; and Ramaprasad, A. 2009. The paradoxical nature of electronic decision aids on comparison-shopping: The experiments and analysis. *Journal of Theoretical and Applied Electronic Commerce Research* 80–96.
- Xiao, B., and Benbasat, I. 2007. E-commerce product recommendation agents: Use, characteristics, and impact. *MIS Quarterly* 31(1):137–209.