

MIT LIBRARIES

DUPL



3 9080 02618 0114

CASEAEN





Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

<http://www.archive.org/details/searchobfuscatio00elli>

HB31
.M415
no.04-
27

DEWEY

Massachusetts Institute of Technology
Department of Economics
Working Paper Series

**SEARCH, OBFUSCATION, AND PRICE ELASTICITIES
ON THE INTERNET**

Glenn Ellison
Sara Ellison

Working Paper 04-27

June 2004

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
<http://ssrn.com/abstract=564742>

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

DEC 07 2004

LIBRARIES

Search, Obfuscation, and Price Elasticities on the Internet¹

Glenn Ellison
MIT and NBER

and

Sara Fisher Ellison
MIT

June 2004

¹We would like to thank Nathan Barczy, Nada Mora, Silke Januszewski, Caroline Smith and Andrew Sweeting for outstanding research assistance. We also thank Patrick Goebel for a valuable tip on Internet data collection, Steve Ellison for sharing substantial industry expertise, and Drew Fudenberg for his comments. This work was supported by NSF grants SBR-9818524 and SES-0219205. The first author's work was supported by fellowships from the Center for Advanced Study in the Behavioral Sciences and the Institute for Advanced Study. The second author's work was supported by fellowships from the Hoover Institute and the Institute for Advanced Study.

Massachusetts Institute of Technology
Department of Economics
Working Paper Series

**SEARCH, OBFUSCATION, AND PRICE ELASTICITIES
ON THE INTERNET**

Glenn Ellison
Sara Ellison

Working Paper 04-27

June 2004

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
<http://ssrn.com/abstract=564742>

Search, Obfuscation, and Price Elasticities on the Internet¹

Glenn Ellison
MIT and NBER

and

Sara Fisher Ellison
MIT

June 2004

¹We would like to thank Nathan Barczy, Nada Mora, Silke Januszewski, Caroline Smith and Andrew Sweeting for outstanding research assistance. We also thank Patrick Goebel for a valuable tip on Internet data collection, Steve Ellison for sharing substantial industry expertise, and Drew Fudenberg for his comments. This work was supported by NSF grants SBR-9818524 and SES-0219205. The first author's work was supported by fellowships from the Center for Advanced Study in the Behavioral Sciences and the Institute for Advanced Study. The second author's work was supported by fellowships from the Hoover Institute and the Institute for Advanced Study.

Abstract

We examine the competition between a group of Internet retailers that operate in an environment where a price search engine plays a dominant role. We show that for some products in this environment, the easy price search makes demand tremendously price-sensitive. Retailers, though, engage in obfuscation—practices that frustrate consumer search or make it less damaging to firms—resulting in much less price sensitivity on some other products. We discuss several models of obfuscation and examine its effects on demand and markups empirically. Observed markups are adequate to allow efficient online retailers to survive.

Keywords: search, obfuscation, Internet, retail, search engines, loss leaders, add-on pricing, website design, demand elasticities, frictionless commerce

Massachusetts Institute of Technology
Department of Economics
Working Paper Series

**SEARCH, OBFUSCATION, AND PRICE ELASTICITIES
ON THE INTERNET**

Glenn Ellison
Sara Ellison

Working Paper 04-27

June 2004

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
<http://ssrn.com/abstract=564742>

Search, Obfuscation, and Price Elasticities on the Internet¹

Glenn Ellison
MIT and NBER

and

Sara Fisher Ellison
MIT

June 2004

¹We would like to thank Nathan Barczy, Nada Mora, Silke Januszewski, Caroline Smith and Andrew Sweeting for outstanding research assistance. We also thank Patrick Goebel for a valuable tip on Internet data collection, Steve Ellison for sharing substantial industry expertise, and Drew Fudenberg for his comments. This work was supported by NSF grants SBR-9818524 and SES-0219205. The first author's work was supported by fellowships from the Center for Advanced Study in the Behavioral Sciences and the Institute for Advanced Study. The second author's work was supported by fellowships from the Hoover Institute and the Institute for Advanced Study.

Abstract

We examine the competition between a group of Internet retailers that operate in an environment where a price search engine plays a dominant role. We show that for some products in this environment, the easy price search makes demand tremendously price-sensitive. Retailers, though, engage in obfuscation—practices that frustrate consumer search or make it less damaging to firms—resulting in much less price sensitivity on some other products. We discuss several models of obfuscation and examine its effects on demand and markups empirically. Observed markups are adequate to allow efficient online retailers to survive.

Keywords: search, obfuscation, Internet, retail, search engines, loss leaders, add-on pricing, website design, demand elasticities, frictionless commerce

1 Introduction

When Internet commerce first emerged one heard a lot about the promise of “frictionless commerce.” Internet technologies would make it easy for consumers to find the exact product they wanted and to compare prices from competing retailers.¹ The most intriguing demonstration of this potential so far is price search engines. A consumer who knows the exact product he or she wants need only type the name into Dealtime, mySimon, Shopper, Pricewatch, Pricescan or Kelkoo to receive instantaneously and without charge a nicely sorted list of the prices charged by dozens of e-retailers. Once a curiosity, these sites now have tens of millions of users.² Internet search is also affecting many “brick and mortar” industries. For example, a recent J.D. Power survey indicated that about one quarter of all new car buyers reported that Internet research had a “big impact” on their make/model choice.³ The diffusion of Internet search technologies has been sufficiently slow, however, so that their potential impact is much greater than the impact we have seen so far.

The goal of this paper is to provide some insights into how online and traditional retail may be affected as price search technologies improve and become more popular. Our most basic message is that economists should think about “obfuscation” as well as search: as search technologies improve firms will have an incentive to put some “friction” back in the market by creating an environment in which price search is more difficult or at least less of a threat to profitability. We explore the interaction of search, price search engines, and obfuscation with a variety of methodologies: we present some informal descriptive evidence; we develop some formal models; and we carry out several econometric analyses of an online retail environment in which a search engine plays a central role.

¹For example, in 1999, *The Wall Street Journal* quoted an economist with Lehman Brothers who estimated that e-commerce would cut inflation by half a percentage point by 2002 because “the Internet eliminates local monopoly power.” See “New E-conomy: If E-commerce Helps Kill Inflation, Why Did Prices Just Spike?” *The Wall Street Journal*, October 18, 1999, A1.

²ComScore Media Metrix reported that Dealtime was visited by 13.8 million distinct U.S. users in August of 2003, which is about 9% of all U.S. Internet users. CNET’s 2002 annual report indicates that its shopping sites, which include mySimon and Shopper as well as sites offering product reviews, received referral fees from more than 100 million clickthroughs. Dealtime’s visitor total is up about 300% from 2000. See White (2000) for more on the impact of price search engines circa 2000.

³The J.D. Power survey result was announced in an October 1, 2003 press release. Another example is airline demand, which is surely affected by the popularity of sites like Expedia, Travelocity, and Orbitz. ComScore Media Metrix reports that each had over 10 million unique visitors per month as of October 2003.

We begin with a brief theoretical discussion of price search engines. We mention some potential paradoxes that any model of price search must resolve. For example, a price search engine could not make any money if it induced retailers to compete prices down to marginal cost. We argue that it is not hard to get around these problems and develop a theory of price search engines as long as one allows the search engines to have sufficient pricing instruments.

We use a simple reduced-form model to illustrate why search engines will often want to create a perfectly frictionless environment. This conflicts with the wishes of retailers, who earn zero profits in frictionless environments. We argue that it is therefore useful to think of search frictions as being determined in a balance-of-power game in which search engines' efforts to reduce frictions are counteracted by retailers' efforts to obfuscate. How technological improvements affect the equilibrium level of search frictions in such a model is clearly ambiguous.

The theory section concludes with a sketch of two specific models illustrating ways in which one can formalize the idea of obfuscating actions raising equilibrium profits. One of these uses standard search theory ideas. It is a theory in which obfuscation makes search more difficult and thereby reduces consumer learning. The other, detailed in Ellison (2003), is a competitive price discrimination model. It is a theory in which obfuscation doesn't actually reduce consumer learning in equilibrium, but instead reduces the degree to which consumer learning hurts firm profitability.

Firms have, of course, always tried to thwart price search. We feel that the Internet may, however, make obfuscation more important for two reasons. First, improved search tools make firms worse off if they do not engage in obfuscation. Second, the Internet may make obfuscation cheaper and easier in a number of ways. One is traditional retailers must hire a substantial number of articulate salespeople if they want to offer nonstandard sales contracts, to use legal "bait-and-switch" techniques, or to offer personalized prices, whereas Internet retailers can easily create automated sales pitches which cheaply implement all of these strategies.⁴ Another is that e-retail purchases are naturally bundled with shipping,

⁴The high costs of hiring articulate salespeople probably accounts for why such practices are usually restricted to retailers selling fairly expensive products like cars, appliances, and mattresses.

which allows products to be offered at many different prices. Another is that it is easier for e-retailers to gather information about consumers and personalize prices. Each of these practices can frustrate consumer price search.

Most people are still buying most items as they did before the Internet existed. To try to get insights into what may happen in the future if price search engines do become very popular we have chosen to examine an environment that is now very atypical. We call it the Pricewatch universe. Pricewatch is an Internet price search engine that is popular with relatively savvy people shopping for computer parts. Its universe is inhabited by hundreds of small, low-overhead retailers. They attract consumers largely just by keeping Pricewatch informed about their (low) prices.

Our informal evidence section describes various practices that we think of as forms of obfuscation. Some of these are as simple as making product descriptions complicated and creating multiple versions of products so that consumers have to examine the attributes and prices of a large number of products to know what is being offered. We particularly call attention to a practice that has become ubiquitous in the Pricewatch universe: firms offer an (inefficiently) low quality product at a very low price to attract customers and then try to talk consumers who visit the site into paying extra to get the product they really want. For example, the retailer we study lists very low prices on Pricewatch for “bare” CPUs, and then tries to convince consumers to instead purchase a CPU with an attached cooling fan. We will refer to this practice as being a “loss-leader strategy” even though it sometimes differs from the classic loss-leader strategy in two respects: it may involve getting consumers to upgrade to a superior product rather than getting them to buy both the loss-leader and a second physical good; and the “loss-leader” may be sold for a slight profit rather than at a loss. The model of Ellison (2003) we mentioned discusses why the adoption of such practices (referred to there as “add-on pricing”) can be a means of obfuscation.

The majority of this paper is devoted to formal empirical analyses of the Pricewatch universe. We address several questions via an analysis of demand and substitution patterns within eight specific product categories: four categories of computer memory modules and four categories of CPUs. We gather data from two sources. We obtained year-long hourly price series by using Go!Zilla to repeatedly conduct price searches on Pricewatch in each

of those eight product categories. We matched this to sales data obtained from a single private firm that operates several computer parts websites and derives most of its sales from Pricewatch referrals.

Our first empirical result is a striking confirmation of the hypothesis that price search on the Internet may dramatically reduce search frictions and lead to extremely elastic demand. We estimate that the firm faces a demand elasticity of between -25 and -40 for its lowest quality memory modules! These are some of the largest demand elasticities we have seen empirically estimated. For single product retailers they would lead to a “Bertrand paradox” where the equilibrium price would be so low as to prevent retailers from covering their fixed costs.

Our second main empirical result can be regarded as a contribution to the empirics of loss leaders. We show that charging a low price for a low quality product increases our retailer’s sales of medium- and high-quality products. The reason why this happens is that one cannot ask a search engine to find “decent-quality memory modules sold with reasonable shipping, return, warranty and other terms.” Hence, many consumers use Pricewatch to do what it can do – find the websites that offer the lowest prices for *any* memory module – and then search within a few of these websites to find other products that better fit their preferences. This result (as well as our findings that website design appears to be a critical determinant of the success of an e-retailer) may also be of interest to the marketing literature, which has discussed loss leaders but does not contain clear empirical evidence of their effectiveness.

Our third main set of empirical results examine how it is that obfuscation affects profitability. Specifically, we explore predictions of the two specific obfuscation mechanisms discussed in our theory section. In the search theoretic model, obfuscation raises profits by making consumers less informed. In the competitive price-discrimination model, obfuscation raises profits by creating an adverse-selection effect makes firms less interested attracting consumers via price cuts. We find evidence of the importance of both mechanisms: we find that many consumers appear to have ended up with limited information; and we find that there is an important adverse selection problem that firms would want to consider when contemplating a price cut.

Finally, we examine an additional data source, cost data, for direct evidence that retailers obfuscation strategies have been successful in raising markups beyond the level that would otherwise be sustainable. Given the extreme price sensitivity of the demand for low-quality products, a naive application of single-good markup rules would suggest that equilibrium price-cost margins might be just 2.5% to 4%. We explain, using actual data in an example, why this is not what one should have expected in light of our findings about the adverse selection effect in demand. We find that the average markup on memory modules is actually about 12%. Pets.com could never have survived with a 12% markup, but this should be fine for our retailers. We conclude that the level of obfuscation on Pricewatch is sufficient to enable active competition between efficient retailers.

The most closely related study of price search engines and demand that we know of is Brynjolfsson and Smith (2001). They use a dataset containing the click sequences of tens of thousands of people who conducted price searches for books on Dealtime to estimate several discrete choice models of demand.⁵ They note that book retailers appear to be differentiated and identify price premia that consumers are on average willing to pay to buy from branded retailers (Amazon, Barnes & Noble, and Borders) and from retailers they have patronized in the past.⁶ The one study of e-retail we know of that reports demand elasticities is Chevalier and Goolsbee (2003), which estimates elasticities for book sales at Amazon and Barnes & Noble. They collected data on books' sales ranks at the two sites before, during, and after an Amazon pricing experiment, and then estimated demand elasticities by inferring sales from the data on sales ranks.

Two other studies have provided some evidence on the link between Internet search and price levels. Brown and Goolsbee note that Internet sites comparing life insurance premia appeared in 1996 and term life insurance premia fell 12% in 1996 and 8% in 1997, mostly falling for those with higher predicted Internet usage. Scott Morton, Zettelmeyer and Silva-Risso (2001, 2003) report that people who received dealer referrals from Autobyte.com (which does not itself provide any price comparison capabilities) paid 1% or 1.5% less for

⁵A disadvantage of their dataset is that they do not observe purchases and must use last clickthroughs as a proxy.

⁶They do not report price elasticities, but they could presumably easily be derived from their estimates and reported separately for branded and generic online bookstores.

their new cars than did other customers who bought comparable cars. Finally, a number of studies have examined price dispersion at e-retailers. See, for example, Clay, Krishnan and Wolff (2002) and Baye, Morgan and Scholten (2003).

2 Theory of price search engines; search and obfuscation

We begin this section with a high-level model that brings out a basic intuition for why search engines may want to reduce search frictions and retailers to increase them. We then discuss a couple ways in which one could model search and obfuscation more concretely.

Any model of price search engines must avoid two possible contradictions. The first is the Bertrand paradox. A price search engine that caused all retailers to go out of business would be of little use. The second is what we will call the search engine revenue paradox. If a price search engine creates Bertrand-like competition, then retailers cannot pay the search engine because they are making no profits, and consumers will not pay the search engine because if there is no price dispersion they can just go directly to any retailer. We note that it is easy to avoid these paradoxes if the search engine has adequate pricing instruments.

Consider a retail sector consisting of a large number of firms selling a single undifferentiated product. Suppose that the only way that retailers can reach consumers is via a monopoly price search engine. The search engine can monitor sales that are made as a result of its searches and charge retailers a referral fee of r for each sale made.⁷ Suppose that the outcome of the price competition game between the listed retailers depends on two parameters: the wholesale price w at which retailers acquire the good and a parameter s that we call the level of “search frictions.” Assume that aggregate sales in this equilibrium are $Q^*(w, s)$ and aggregate retailer profits are $\pi^r(w, s)$. Assume that these functions are differentiable and that for small s we have $\frac{\partial Q^*}{\partial w} < 0$, $\frac{\partial Q^*}{\partial s} < 0$, $\frac{\partial \pi^r}{\partial w} < 0$ and $\frac{\partial \pi^r}{\partial s} > 0$.⁸

⁷Yahoo! Shopping, for example, charges retailers 2% of their gross revenues on any sales made during the course of a browser session in which the consumer was referred to the website by Yahoo! Shopping.

⁸The one condition where one would expect “for small s ” to bind most quickly is the last one, but given the elasticities we report, we feel comfortable assuming that the firms in our data would prefer somewhat less efficient search. A main observation of Diamond’s (1971) search model is that profits can be discontinuous and jump to the monopoly price with a positive search cost. A large subsequent literature has explored variants of the model in which intermediate search costs lead to intermediate outcomes. See Varian (1980), Burdett and Judd (1983) and Stahl (1989). Bakos (1997) and Janssen and Moraga (2000) examine the effects of search costs on prices and welfare. Bakos examines a model with horizontally differentiated firms where consumers must pay a search cost to learn a firm’s location as well as its price and notes that lower

Assume that prices converge to marginal cost as s goes to zero so $Q^*(w, 0) = D(w)$ and $\pi^r(w, 0) = 0$.

Suppose for now that the price search engine can costlessly choose any level of search frictions s . Write c for the cost at which retailers acquire the good that they sell. The problem facing the search engine is now

$$\text{Max}_{r,s} rQ^*(c + r, s).$$

Given that Q^* is decreasing in s the optimal choice is to eliminate all search frictions (*i.e.*, set $s = 0$). The search engine's problem is then simply the standard monopoly pricing problem,

$$\text{Max}_p (p - c)D(p),$$

where $p \equiv c + r$. The search engine sets $r = p^m - c$ and gets the full monopoly profit. Retailers earn zero profits. This illustrates our first observation, that search engines would like to reduce frictions.

Our resolution to the search engine revenue paradox is that as long as search engines can charge per sale referral fees, firms pass the referral fees on to consumers and search engines can collect their revenues from retailers.⁹ The model does have the Bertrand-paradox problem if retailers have fixed costs. In this case a couple of solutions would be natural: the search engine could make fixed payments to the retailers to cover their fixed costs, or, if that is not feasible, the search engine could choose the minimal level of search frictions that would let the retailers recover their fixed costs.

Now consider a model in which search engines and retailers must make costly investments to increase or decrease search frictions. Search frictions will then typically not be eliminated, and the effect of technological progress on the level of search frictions is indeterminate. This is our second observation. For example, the simplest balance-of-power game would have the search engine choose an investment level x_{se} at cost $g(x_{se}; \theta)$ while search costs lead to lower prices and improved match quality. Janssen and Moraga examine a Stahl-style model with undifferentiated products and two types of consumers and emphasize that in some cases prices may be higher with lower search costs.

⁹As in the literature on vertical restraints, there may be many other contracts that could be used to extract the monopoly profits. For example, the search engine could refuse to post any price below the monopoly price and charge each retailer a fixed fee equal to its expected market share times the monopoly profit.

retailers simultaneously choose x_r at cost $h(x_r; \theta)$, resulting in the level of search frictions being $s_0 - x_{se} + x_r$. The parameter θ indexes the state of technology. Whether increases in θ increase or decrease equilibrium search frictions in a model like this is obviously indeterminate: it depends on whether the technology aids search-improving or search-obfuscating more.

A couple variants of the model are worth mentioning. First, suppose the the search engine could also charge fixed fees. It would then charge a fixed fee of $\pi^r(s^*) - h(x_r^*; \theta)$ to extract the retailers' profits. As long as the fixed and referral fees were chosen in a stage prior to when the x 's were chosen, the determination of equilibrium search frictions would be unchanged—the only difference would be that the retailers' efforts at obfuscation would just help them achieve a zero profit. Second, suppose that there were multiple search engines competing to attract consumers. In a model with differentiation between search engines à la Hotelling, the degree of differentiation determines the utility consumers receive. The search engines will want to provide this utility level in the most efficient way possible. In many models this would give search engines an incentive to reduce search frictions and also lead them to charge lower referral fees than in the monopoly model.

The only full models of price search engines we are aware of are those of Baye and Morgan (2001, 2003). They do not consider the possibility of charging referral fees. They nonetheless avoid the revenue paradox. The key insight is that differences from the Bertrand model that one might think are trivial—the presence of an outside option for retailers and/or positive listing fees—make the standard argument that the Bertrand game has no mixed strategy equilibria inapplicable. It turns out that the model has a symmetric mixed strategy equilibrium in which firms randomize both over whether to list and over the prices to choose if they do. Both retailers and consumers are willing to pay positive fixed fees to the search engine.

2.1 Incomplete consumer search

The model above treats “search frictions” very abstractly. How might one construct more concrete micro models in which search frictions are determined by competing investments by search engines and retailers?

One way would be to build on a standard search model. Stahl (1989), for example, considers a model with two types of consumers. A fraction μ enjoy searching for low prices and always learn the prices offered by all retailers (of which there are a finite number). A fraction $1 - \mu$ incur a cost of c for every price quote they obtain. Consumers are otherwise identical and have downward sloping demands. Stahl notes that this model has a unique symmetric equilibrium: retailers randomize over prices in some interval; informed consumers purchase from the firm offering the lowest price; and the other consumers purchase from the first store they visit provided that its price is below some threshold (which it always is in equilibrium). The solution is continuous in μ and varies in the manner assumed above if we define $s \equiv 1 - \mu$: retailer profits are lower and demand is higher (in a first order stochastic dominance sense) when the fraction of informed consumers is higher, and price converges to cost in the limit as μ goes to one.¹⁰

Building on this model, one could provide an explicit model of search frictions by assuming that consumers are heterogeneous in their ability to digest the information presented on a webpage. The search engine's investment x^{se} would be an investment in collecting information and presenting it in a manner that is easier to understand. Retailers' investments x^r would be investments in making pricing more complicated and difficult to describe. To fit the model above exactly, one could assume that the distribution of consumer abilities is such that after these investments are made, a fraction $\mu_0 + x^{se} - x^r$ of consumers are able to understand the webpage fully and thereby learn all firms' prices, whereas the rest of the consumers cannot understand the price information at all and must resort to visiting the listed websites sequentially (in a random order) to understand each retailer's price.¹¹ The analysis of the search and obfuscation game would then follow exactly as described above. In any interior equilibrium, the model will exhibit positive retail markups, price dispersion across retailers, and limited knowledge of prices by some consumers.

¹⁰The solution has similar comparative statics in the c parameter, so one could alternately take $s \equiv c$ as the measure of search frictions. There are a number of search models in the literature, but Stahl's model stands out for having intuitively appealing comparative statics.

¹¹Stahl's (1996) paper extends the model to allow for heterogeneity in search costs among the consumers with positive search costs, so that some of the partially informed consumers will learn many prices whereas others learn just one or a few.

2.2 Add-ons and adverse selection

Ellison (2003) describes an alternate approach to obfuscation in which obfuscation raises retailer profits even though consumers either observe or correctly anticipate every retailer's price. The model (as one would reinterpret it to fit the search and obfuscation application) involves two retailers, which are slightly differentiated à la Hotelling. Consumers are heterogeneous both in their preferences over the two firms and in their marginal utilities of income. The firms are assumed to sell a standard good that provides v utils (minus transportation costs) to each consumer. They are also able to sell a damaged good that provides only $v - w$ utils.¹² Critically, consumers can distinguish the damaged good from the original, but the search engine cannot. For example, if the search engine is not designed to collect and display shipping details the damaged good could be an offer to sell the standard good accompanied by fine print stating that the shipping time will be one month. Because the search engine cannot distinguish between the damaged and the standard good, in equilibrium all listings will be for the damaged good. Each retailer's website will also contain a separate higher price for the undamaged goods, and consumers will need to incur the cost of visiting each retailer's website to learn what these "add-on" prices are. As in Diamond (1971), the structure of the search costs makes the price of the add-on very high: retailers charge the monopoly price for the incremental value of the standard good over the damaged good. Consumers with rational expectations anticipate this and in equilibrium are fully informed: they directly observe each firm's price for the damaged good and correctly infer the prices that are not posted. Profits earned on the add-on may therefore be partially or completely competed away in the form of lower prices for the listed goods.

The interesting case of the model occurs when the amount of damage w and the heterogeneity in consumers' marginal utilities of income lead some consumers to upgrade to the undamaged good while others settle for the damaged good to save money. Notably, the firms' profits are higher when they compete by listing damaged good prices than they would be if the damaged good did not exist or if the search engine were able to distinguish damaged and undamaged goods. Moreover, the equilibrium profits are increasing in the

¹²See Deneckere and McAfee (1996) for more on damaged goods.

amount of damage w . An intuition for the basic result is that when the firms sell damaged goods, they create an adverse selection problem that makes firms hesitant to undercut each other: a price cut will disproportionately attract cheapskates who buy the damaged good at a low price (which is sometimes below cost).

One could provide a second explicit model of “search frictions” by defining $s \equiv w$ and considering a game in which the search engines invest in developing software that can recognize and inform consumers about different forms of damage, thereby reducing w , while retailers invest in inventing new and creative ways to damage products, thereby increasing w . This model would not exactly fit in the general framework described above for two reasons. First, it needs some differentiation so prices will not drop all the way to cost as search frictions vanish. Second, the higher prices that search frictions bring do not actually reduce the quantity sold because consumers are assumed to have unit demands up to a choke price that is not binding in equilibrium. The first should not matter for the search and obfuscation game. The second is important. A slight modification of the model would, however, both make it more realistic and make it fit within the general framework: one would just need to make the consumers’ demands downward sloping.¹³

3 The Pricewatch universe; memory modules and CPUs

We study a segment of e-retail largely mediated by a price search engine called Pricewatch. The Pricewatch universe is characterized by a large number of small, undifferentiated e-retailers selling memory upgrades, CPUs, and other computer parts. The customers include hobbyists building computers, IT support people, gamers, and other relatively savvy users upgrading old computers, and proprietors of small computer stores. The retailers tend to do little or no advertising, to have rudimentary websites, to receive no venture capital, and to run efficient, profit-maximizing operations. We assume that they receive a large fraction of their customers through Pricewatch.

One can use Pricewatch to locate a product in one of two ways. One can either type a

¹³The unit demands in Ellison’s model are primarily there for tractability and to provide a neat contrast by making add-on pricing completely irrelevant when there is less consumer heterogeneity. The best way to construct a tractable model without unit demands might be to build on the model around logit demands as in Verboven (1999) rather than around the Hotelling model.

technical product description, such as “Kingston PC2100 512MB,” into a search box, or one can run through a multilayered menu to select one of a number of predefined product categories, *e.g.*, clicking on “System Memory” and then on “PC133 128MB SDRAM DIMM.” We believe that the latter is much more common. Pricewatch returns a list of products sorted from cheapest to most expensive in a twelve-listings-per-page format. Note that running through the menu leads to predefined, not user-defined, product categories, and that some of these categories contain as many as 350 listings from 100 different websites. The categories inescapably encompass products of varying quality, often including products made by higher and lower quality manufacturers and always including offers in which the product is bundled with different levels of service. Typically, the first several pages of a list would contain only low quality offerings, the higher quality products at reasonable prices buried deep within the list of 350 products. Figure 1 contains the first page of a typical list, that for PC100 128MB memory modules from October 12, 2000. The low prices listed on Pricewatch are very low. For example, in the period we study, generic memory modules are typically offered for about one-half of the price that Dell charged for a Dell-branded product.

People who have not seen previous studies about price dispersion might think that it is remarkable that there are so many different prices listed for nearly identical items on Pricewatch. In comparison with previous studies, however, we think it is remarkable that prices are so close together. In the early part of our data when a 128MB PC100 memory module sold for about \$100, the average difference between the lowest listed price and the 12th lowest listed price was about \$8. Figure 2 illustrates the range over the course of the year.

Prices for most computer components have declined sharply. Retail prices for memory modules, for example, dropped from about \$100 per megabyte at the start of 1990 to about 10 cents per megabyte in September of 2003. Memory prices are also volatile. In Figure 2, for example, we see that although prices declined by about 70% over the course of the year, there are two periods of sharp price increases. Prices rose by about 50% between late May and early July 2000 and by about 25% in less than two weeks in November 2000. Prices for CPUs are less volatile and decline in a more orderly fashion.

A fascinating aspect of the Pricewatch lists (which greatly aids our econometric analysis) is the substantial turnover from day to day and even from hour to hour. On average, three of the twenty-four retailers on the first two pages of the above-mentioned list will change their prices in a given hour. Each price change typically moves several other retailers up or down one place on the list. Some websites are clearly big players that regularly occupy a position near the top of the Pricewatch list. From time to time one may observe a firm sitting in the first position for a week or more, but there is no rigid hierarchy.

Some of this turnover can be attributed to the technology used by Pricewatch and the various websites at the time. Pricewatch is a database-based system which relies on retailers' updating their own prices in its database. Our impression is that all (or almost all) of the retailers were setting prices in Pricewatch manually in the time period we study. A typical retailer has dozens or hundreds of products listed in the Pricewatch database, making it impractical to constantly monitor one's place in each predefined product category and the current wholesale prices for each product. Instead, a retailer might manually examine its position on the most important Pricewatch categories a few times a day and might look at current wholesale prices once or twice a day.

Our sales and cost data come from a firm that operates several websites selling computer components in the Pricewatch environment. We will examine the pricing of and demand for their products within eight product categories (each of which corresponds to a popular predefined search on Pricewatch): PC100 memory modules in both 128MB and 256MB; PC133 memory modules in both 128MB and 256MB; and AMD Athlon processors in 650 MHz, 700 MHz, 750 MHz, and 800 MHz. In each of these categories, our firm offers a number of products with varying quality levels, typically three.

The first part of each memory module description, PC100 vs. PC133, refers to the speed with which the memory communicates with the CPU. PC133 memory is better and costs about the same. Despite this, sales of the two speeds of memory are comparable, because the speed of a memory module must match the speed of a computer's CPU and motherboard. In the period we study, people who were upgrading old computers more often needed to buy PC100 memory. Hobbyists building new computers usually wanted to

build systems around AMD Athlon CPUs that required PC133 memory. The second part of the product description is the capacity of the memory in megabytes. Many consumers can substitute two 128MB modules for one 256MB modules, although some with older motherboards could not use the “high density” 256MB modules found at the top of the Pricewatch list and some consumers would be slot constrained. The 256MB modules are about twice as expensive. Typically, one spends about \$100 on memory when building an inexpensive computer. In the summer of 2000 when 256MB modules cost about \$250 most of the demand at Pricewatch was for 128MB modules. In the spring of 2001 when 256MB modules cost about \$60, the two were about equally popular.

The retailer that provided us with sales data sells three different products within each Pricewatch memory category on each of its websites, *e.g.*, each website sells three different PC100 128MB memory modules, differentiated by quality of the DRAM and the board as well as by contract terms. Figure 3 illustrates how a similar quality choice is presented to consumers on a website that copied site A’s design. Making quality comparisons within a website is much easier than making comparisons across websites: many memory modules are unbranded and sites often contain minimal technical specifications. In this regard shopping for a memory module can be similar to shopping for mattresses.

The other set of products we study are AMD Athlon processors. Since an AMD Athlon processor is a branded product, a “bare” Athlon with a given clock speed purchased from any website should be identical to a bare Athlon purchased from any other website. The primary purchasers of Athlon CPUs were hobbyists (or proprietors of small computer stores) building new computers. CPUs come in many clock speeds. In early June 2000 prices ranged from about \$120 for a 500 MHz Athlon to over \$1000 for a 1000 MHz Athlon (which was not yet widely available). A typical way to make the speed vs. price tradeoff would be to spend one-fourth of one’s total budget for a computer on the CPU. Pricewatch’s predefined categories recognize that one should compare prices for processors with the same clock speed. We will examine prices and demand within four of the categories that were most popular in the summer of 2000: the 650 Mhz, 700 MHz, 750 MHz, and 800 MHz. Typical prices for these CPUs ranged from \$140 for the 650 MHz Athlon to \$280 for the 800 MHz Athlon.

Despite the fact that the 650 MHz Athlon is a branded product, offerings within the 650 MHz category fail to be homogeneous for two reasons. First, different retailers offer different levels of service. Second, CPUs are often purchased together with a retailer-installed fan/heatsink. One can buy a decent quality fan/heatsink combination for about \$5 and an experienced installer can attach it to the CPU in less than a minute. The attachment can be tricky for a novice, though, and entails a nontrivial chance of ruining the CPU. Retailers may offer to attach any of several fans, and each CPU-fan combination should be viewed as a distinct product within the same Pricewatch category. For example, one of the websites from which we have data typically offers three products within each category: a bare CPU, a CPU attached to an AMD-approved fan, and a CPU together with a premium-quality fan (and also bundled with better service terms and an extended warranty). To reiterate, Pricewatch’s categories did not distinguish between CPUs with and without fans (or with better or worse fans), so all of these products would appear in the same Pricewatch list.

4 Observations of obfuscation

Over the last few years Pricewatch has made a number of enhancements to combat obfuscation. Practices that frustrate search nonetheless remain commonplace.

One of the most visible search-and-obfuscation battles was fought over shipping costs. In its early days Pricewatch did not collect information on shipping costs and sorted its lists purely on the basis of the item price. Shipping charges grew to the point that it was not uncommon for firms to list a price of \$1 for a memory module and inform consumers of a \$40 “shipping and handling” fee when they went to check out. Pricewatch fought this with a two-pronged approach: it mandated that all firms offer UPS ground shipping for a fee no greater than a Pricewatch-set amount (\$11 for memory modules); and it added a column to its listing pages that either displayed the shipping charge (if the retailer reported it) or a warning that customers should be wary of stores that do not report their shipping charges.¹⁴ Many retailers adopted an \$11 shipping fee for memory modules in response, but uncertainty about the cost of UPS ground shipping was not completely eliminated because a

¹⁴Our empirical work is based on data from the period when these policies were in effect.

number of retailers left the column blank or (more commonly) reported a range of shipping charges. The meaning of “UPS ground shipping” was also subject to manipulation: one company explicitly stated on its website that items ordered with the standard UPS ground shipping were given lower priority for packing than other orders and might take two weeks to arrive. More recently, Pricewatch mandated that retailers provide it with shipping charges and switched to sorting low-price lists based on shipping-inclusive prices (which are based on the maximum shipping charge if a range is given). This appears to have worked as desired, although it only provides the desired information for customers who prefer ground shipping: customers who wish to upgrade to 3rd-, 2nd-, or next-day air must look through the retailers’ websites manually.

One of the potential models of obfuscation we discussed involved firms’ trying to increase customers’ inspection costs and/or to reduce the fraction of customers who will buy from the firm on the top of the search engine’s list without doing any firm-by-firm searching. We observed several practices that might serve this purpose. The most effective seems to be bundling low-quality goods with unattractive contractual terms, like providing no warranty and charging a 20% restocking fee on all returns. Given the variety of terms we observed, it would seem unwise to purchase a product listed on Pricewatch without reading the fine print. Another is making advertised prices difficult to find. In 2001 it took us quite a bit of time to find prices listed on Pricewatch on several retailers’ sites. In a few cases, we never found the listed prices. Several other firms were explicit that Pricewatch prices were only available on telephone orders. Given that phone calls are more costly for the retailers, we can only assume that firms either wanted people to waste time on hold or intended to make them sit through sales pitches. Pricewatch has fought these practices in several ways. For example, it added a “buy now” button, which (at least in theory) takes customers directly to the advertised product.

The second obfuscation mechanism we discussed is the adoption of a “loss-leader” or “add-on” pricing scheme: damaged goods are listed on the search engine at low prices and websites are designed to convince customers attracted by the low prices to upgrade to a higher quality product. Such practices are now ubiquitous on Pricewatch. Figure 3 is one example. Customers who try to order a generic memory module from Buyaib.com at the

price advertised on Pricewatch.com are directed to this page. It illustrates several ways in which the low priced product is inferior to other products the company sells (at higher markups). Figure 4 is another example. A consumer who tries to order a generic module from Tufshop.com is taken to this page, from which a number of complementary products, upgrades, and services. The figure presents the webpage as it initially appears. To avoid purchasing the various add-ons the consumer must read through the various options and unclick several boxes. After completing this page, a Tufshop.com customer is taken to another on which he or she must choose from a long list of shipping options. These include paying \$15.91 extra to upgrade from UPS ground to UPS 3-day, \$30.96 extra to upgrade to UPS 2-day, and \$45.96 extra to upgrade to UPS next day.¹⁵

Our impression is that the practices are also consistent with the add-on pricing model in terms of the low-priced goods being of inefficiently low quality. In Pricewatch's CPU categories all of the listings on the first few pages were bare CPUs. We are told that most of the "generic modules" at the top of Pricewatch's memory lists are very low quality products that are much more likely to have problems than are other modules that can sometimes be purchased wholesale for just one or two dollars more. We know that the wholesale price difference is occasionally so small as to induce the retailer from which we got our data to ship "medium quality" generic modules to customers who ordered low quality modules (without telling the customers) because it felt the time cost and hassle of dealing with returns was not worth the cost savings.

Obfuscation could presumably take many forms in addition to those we outlined in our theory section. One is that firms could try to confuse boundedly-rational consumers. Presumably, this would involve either tricking consumers into paying more for a product than it is worth to them or altering their utility functions in a way that raises equilibrium profits.¹⁶ Our impression is that many Pricewatch e-retailers sites' are intentionally confusing. For example, whereas several sites will provide consumers with product comparison

¹⁵The incremental costs to Tufshop of the upgraded delivery methods are about \$4, \$6, and \$20.

¹⁶Following our work on obfuscation, Gabaix and Laibson (2004) have developed a model of consumer confusion. They assume confused consumers behave exactly as if they had stronger idiosyncratic preferences for one firm relative to the other in a standard horizontal differentiation model. Hence, an analysis of pricing and incentives to invest in confusing consumers is identical to an analysis of firms that can invest in differentiation. Welfare consequences are different, however, because the idiosyncratic factors that make consumers prefer product A over product B are not preferences that contribute to utility.

lists like that in Figure 3, we did not see any that augmented such a comparison with a description of what “CAS latency” means to help consumers think about whether they should or shouldn’t care about it.

As mentioned above, Pricewatch is a database-based facility that requires that retailers enter their prices into a database. This scheme for running a price search website seems to be gaining relative to the alternate technology of using shopbots to maintain a price database or to search for prices in realtime whenever a consumer requests them. One explanation for this must be that the shopbot approach can result in substantial delays while the site searches. Another may be that the shopbot approach may be even more prone to obfuscation. In 2001, for example, Yahoo! Shopping search engine should have had a much easier time organizing products than a general search engine because it only searched sites hosted by Yahoo. Yahoo collected a royalty on all sales made by merchants through Yahoo! Shopping (as well as a monthly fee based on the number of items offered at a Yahoo! Store), so there must have been some standardization of listing and ordering mechanics. Nonetheless, when we typed “128MB PC100 SDRAM DIMM” into the search box and sorted on price, the five lowest listed prices were from merchants who had figured out how to get Yahoo! Shopping’s search engine to think the price is zero even though a human who clicks over to the retailer can easily see the price (and see that it is 50-100% above the Pricewatch price). The next hundred or so cheapest items on Yahoo’s search results were also either products for which Yahoo’s search engine had misinterpreted the price or items which were inexpensive because they were not the desired product.¹⁷

5 Data

Our price data were downloaded from Pricewatch.com. They contain information on the twelve (sometimes twenty-four) lowest price offerings within each of the eight predefined

¹⁷In 2001 it was possible to use Yahoo! Shopping’s search engine to get a well-sorted list and find low prices if one was experienced with it. In 2003 Yahoo! Shopping’s search engine conducts a broad search not restricted to Yahoo! Shops and affiliates. The list returned in response to a corresponding search no longer begins with a large number of sites that Yahoo thinks have zero prices, but it does begin with a large number of sites for which it has incorrectly parsed the HTML and inferred the price to be less than it truly is. Most of these mistakes appear to be accidental rather than due to a plan by the retailers to fool the Yahoo! Shopping search engine.

categories mentioned above.¹⁸ These data were collected at hourly frequency from May 2000 to May 2001.¹⁹ Recall that Pricewatch’s predefined categories are somewhat coarse. The vast majority of the products for which we downloaded prices are what we could call “low-quality” products.

In addition to the price data for these “low-quality” products, we obtained price and quantity data from one Internet retailer. The data contain the prices and quantities sold for all products that fit within the eight Pricewatch categories, often three different quality products in each category. The firm operates several different (but similar) websites, which typically have different prices for the products studied.²⁰ We use data from three websites, which we will simply call A, B, and C.

The raw data are at the level of the individual order and indicate from which website the customer made the purchase. We currently have approximately a year of sales data starting in May of 2000. We aggregate the individual orders to produce daily sales totals for each product at each website.²¹ Our primary price variables are the average transaction prices for sales of a given product on a given day.²² We also record the daily average position of each website on Pricewatch’s price ranked list.

The same Internet retailer also provided us with data on wholesale acquisition costs for each product. The firm maintains very low inventories. Some products are purchased from wholesalers every few days or once a week. Others are purchased almost every day. The firm often has a “negative” inventory of the most popular memory modules, buying them each afternoon to fill the backlog of retail orders taken that morning and the previous

¹⁸We collected the twenty four lowest prices for the 128MB PC100 and 128MB PC133 memory modules and the twelve lowest prices in each of the other six categories.

¹⁹We used Go!Zilla to carry out the downloads.

²⁰Among the motivations for having several websites are that different websites may be given different looks and consumers may have heterogeneous reactions, that it allows the websites to be more specialized (which seems to be attractive to some consumers), that it facilitates experimentation, that it may help promote private-label branded products, and that it lets the firm occupy multiple places on the Pricewatch screen.

²¹Here, a “product” includes also the quality level, *e.g.*, a high-quality 128MB PC100 module.

²²Transaction prices are unavailable for products which have zero sales on a given day. If low-quality products ever have zero sales, then we fill the prices in using the listed prices from Pricewatch, since these prices are quite volatile. Prices for the other products change less often and can be filled in fairly completely by just assuming that prices are constant in the gaps between times when we observe identical transaction price observations. We also fill in a few other values by looking at when low-quality prices and prices at other websites operated by the firm changed.

evening. The data provided to us contain wholesale prices whenever the firm purchased a particular product.²³

Our analysis of memory module demand uses data from websites A and B. These two websites have identical product lineups: they sell three products within each memory module category, which we refer to as the low-, the medium-, and the high-quality module. Our dataset on memory modules contains between 575 and 683 observations in each category.²⁴ Summary statistics for each of the four categories are presented in Table 1. Note that the data are at the level of the website-day, so the number of days covered is approximately half of the number of observations. *LowestPrice* is the lowest price listed on Pricewatch (which is presumably for a low-quality memory module.)²⁵ *Range1-12* is the difference between the twelfth lowest listed price and the lowest listed price. Note that the price distribution is fairly tight. *PLow*, *PMid* and *PHi* are the prices for the three qualities of memory modules at the two websites. *PLowRank* is rank of the website's first entry in Pricewatch's sorted list of prices within the category in question.²⁶ This variable turns out to allow us to predict sales much better than we can with simple functions of the cardinal price variables. Note that the websites we study are consistently near the top of the Pricewatch list: the average prices for the low-quality products they sell are within 10% of the average lowest price on Pricewatch in all four categories. *QLow*, *QMid* and *QHi* are the average daily quantities of each quality of module sold by each website. The majority of the sales are the low-quality modules.

We have not broken the summary statistics down by website. Website A's prices are usually lower than website B's, but there is no rigid relationship. For example, in the 128MB

²³The only products purchased infrequently are CPU fans. This does not cause any substantial incompleteness of the data because prices of CPU fans rarely change.

²⁴Data are occasionally missing due to failures of the program we used to collect data and missing data in the files the firm provided. We left gaps in the price curves in Figure 2 to indicate the frequency with which we were missing data. The 256MB prices are missing for most of the last six weeks, so we chose to use mid-March rather than May as the end of the 256MB samples.

²⁵The Pricewatch data are hourly. Daily variables are constructed by taking a weighted average across hours using weights that reflect the average hourly sales volumes of the websites we study.

²⁶We only know a sites' Pricewatch rank if it is among the 24 lowest priced websites for 128MB modules or the 12 lowest priced websites for 256MB modules or Athlon processors. When a site's price is too high and it does not appear on the lists we downloaded from Pricewatch we set *PLowRank* equal to one more than the number of firms appearing on the lists we collected, i.e. 25 for 128MB modules and 13 for the other products. In the 128MB category this happens for fewer than 1% of the observations. In the 256MB category this happens for 3% of the site A observations and 14% of the site B observations.

PC100 memory module category, website A's low-quality price is lower on 251 days, the same on 27 days, and higher on 60 days. In this category the price differences are usually quite small, e.g. they are two dollars or less more than 70% of the time. In the 256MB categories the price differences are somewhat larger: the median is five dollars. As one would expect, Site A also has higher sales volumes. For example, it is responsible for about 70% of the sales of low-quality 128MB PC100 memory modules.

Our analysis of demand patterns for AMD Athlon processors uses data from websites B and C. Athlon sales drop off sharply in the fall of 2000 following AMD's introduction of the improved Athlon Thunderbird. Accordingly, we analyze data from a shorter period, May to November of 2000. An additional factor complicating the analysis is that sites B and C do not have identical product lineups. Site C offers three products within each processor-speed category: the "low quality" product is usually a bare CPU (although the firm experimented briefly with selling a CPU with a nonapproved fan instead); the "medium quality" product has an AMD-approved fan/heatsink attached; and the "high quality" product is a CPU with a premium quality dual fan. Site B offers two products in each category: the bare CPU and the CPU with a dual fan. The fact that sites B and C have separate sets of offerings forces us to analyze site B's sales and site C's sales in separate regressions. The combination of this and the shorter time period led us to pool data from the four clock-speed categories and restrict the coefficients to be identical across clock speeds except for additive clock-speed dummies. Table 2 contains summary statistics for the Athlon processor data broken down as it is used in the regressions. The first panel refers to the data on site B's sales and the second to site C's sales. The unit of observation is the clock speed-day. Hence, each panel summarizes a dataset with four observations per day, and the mean prices reported are averages across clock speeds. The minimum price of \$105.25 for *PLow* on site B is what a bare 650MHz Athlon cost near the end of the sample. The maximum price of \$302.45 is what a bare 800MHz Athlon cost near the start of the sample. The mean prices for *PLow* and *PHi* in the website B data reflect that customers of website B would on average need to pay about \$39 more to get a CPU with a fan attached. Unit sales of CPUs are lower than unit sales of memory modules. This will also reduce the precision of our estimates. The fraction of consumers who buy upgraded products is a little higher than in the memory

data: 39% of the customers of website B and 56% of the customers of website C buy a medium- or high-quality product.

6 Demand patterns

In this section we estimate the demand elasticities Pricewatch e-retailers face and examine how consumers substitute between low-, medium-, and high-quality products. We do this both to provide new descriptive evidence on the functioning of online markets and to provide empirical evidence on the theories of obfuscation discussed above.

6.1 Methodology for demand estimation

Our primary interest is in how website w 's sales of low-, medium-, and high-quality products depend on its prices for the three products and on its competitors' prices. Specifically, we suppose that within each product category c , the quantity of quality q products purchased from website w on day t is

$$Q_{wcqt} = e^{X_{wct}\beta_{cq} + u_{wcqt}},$$

with

$$\begin{aligned} X_{wct}\beta_{cq} = & \beta_{cq0} + \beta_{cq1}\log(1 + PLowRank_{wct}) + \beta_{cq2}\log(PMid_{wct}) + \beta_{cq3}\log(PHi_{wct}) \\ & + \beta_{cq4}\log(LowestPrice_{ct}) + \beta_{cq5}Weekend_t + \beta_{cq6}SiteB_w + \sum_{s=1}^{12} \beta_{cq6+s}TimeTrend_{st}. \end{aligned}$$

In this equation $PLowRank_{wct}$ is the rank of website w 's low-quality offering on Pricewatch's list of the lowest prices in category c . It is the primary variable that we use to capture how website w 's low-quality price differs from those of the other firms listed on Pricewatch. $PMid_{wct}$ and PHi_{wct} are the prices of website w 's medium- and high-quality offerings.²⁷ The $TimeTrend$ variables allow for a piecewise linear time trend with a slope that changes every 30 days. We think of the fourth price variable, $LowestPrice_{ct}$, mainly as an additional control variable.²⁸

²⁷Medium- and high-quality cannot be defined consistently across retailers. Accordingly, we have no data on competitors' prices for comparable products to include in the regressions.

²⁸We do not treat the variable as providing an estimate of the elasticity of the aggregate demand at all Pricewatch retailers as a function of their price because we lack data on prices at traditional retailers. Interpretation would be even more difficult in the medium- and high-quality demand regressions because the coefficient on $LowestPrice$ will also reflect how decisions depend on medium- and high-quality prices relative to low-quality prices.

We estimate the demand equations via GMM. Specifically, for most of our estimates we assume that the error term u_{wct} satisfies $E(e^{u_{wct}}|X_{wct}) = 1$ so that we can estimate the models using the moment condition

$$E(Q_{wct}e^{-X_{wct}\beta_{cqt}} - 1|X_{wct}) = 0$$

These estimates are done separately for each product category and each quality level. Standard errors were computed with a Newey-West style approach.

The coefficients on the logs of the prices of the medium- and high-quality modules can be directly interpreted as elasticities. To construct elasticities of demand with respect to a website's low-quality price, we treat $\log(1 + PLowRank)$ as a continuous variable and compute estimated elasticities when all variables are at their means by setting the derivative of $PLowRank$ with respect to a change in $PLow$ equal to the inverse of the average distance between the twelve lowest prices.

This approach presumes that the price variables are not endogenous. Although we present estimates in a later section that do not assume exogeneity of prices, these non-instrumented estimates are our preferred ones for two reasons. First and most importantly, we think that little of the variation in our e-retailers' prices is due to its using information about demand. The firm appears to have little information about whether particular days or times might be good ones to be higher or lower on Pricewatch's lists. Most of the changes in our e-retailer's ranks, in fact, are not due to any decisions: the e-retailer has limited attention to devote to each of the dozens of Pricewatch lists on which its products appear, and most rank changes are due to other firms changing their prices.

The person who sets prices noted in discussions with us that he would typically change his low-quality prices for one of three reasons. First, a few times a day he checks some of the Pricewatch lists and may raise or lower a price if a rank has drifted too far from where he typically leaves it. Second, once a day or so he receives updated wholesale price lists and may change a price in response to a wholesale price change. Third, he occasionally raises prices if order-processing or shipping employees have failed to show up for work. Less attention is paid to medium- and high-quality prices. These are typically left fixed for several weeks at a time.

Second, given our assumed functional form for demand, economic theory implies that prices should not be endogenous. In our demand model, shocks increase or decrease demand multiplicatively. In such an environment, the profit-maximizing price of a firm with constant marginal costs is independent of the realization of the demand shock. Hence, prices set to maximize profits would not be endogenous even if the firm observed the demand shock before setting its price.

Despite these arguments, we do also estimate our model using two distinct sets of instruments for *PLowRank*, *PMid*, and *PHi*. These estimates are obtained from a GMM estimation of the moment condition

$$E(Q_{wct}e^{-X_{wct}\beta_{cq}} - 1|W_{wct}) = 0,$$

where W_{wct} is a vector containing the instruments instead of the prices. We discuss this estimation in section 6.5.

6.2 Basic results on demand

Table 3 presents estimates of demand from one of our product categories: 128MB PC100 memory modules. The first column of the table contains estimates of the demand equation for low-quality 128MB PC100 modules. The second and third columns contain estimates of the demand for medium- and high-quality 128MB PC100 modules, respectively. Coefficient estimates are presented with t -statistics in parentheses below them.

Our first main empirical result is that the demand for low-quality modules a website faces is extremely price sensitive. This is due to a combination of two factors. First, a price change of a few dollars can move a website up or down several places on the Pricewatch list. Second, most people buying low-quality memory modules buy them from one of the first few firms on the Pricewatch list. The coefficient estimate on the $\log(1 + PLowRank)$ variable in the first column implies that moving from from first to second on the list reduces a website’s sales of low-quality modules by 45% and dropping from first to seventh reduces sales by 87%.²⁹ This feature of demand is remarkably clear in the data—we get a t -statistic of 15.1 in a regression with only 683 observations.

²⁹A move from first to second increases $\log(1 + PLowRank)$ from from $\log(2) \approx 0.69$ to $\log(3) \approx 1.1$, thereby reduces log-quantity by $1.48 * 0.41 \approx 0.6$.

Table 4 presents demand elasticities. The upper left number in the upper left panel indicates that the 1.48 coefficient estimate in the low-quality 128MB PC100 demand equation corresponds to an own-price elasticity of -25.0. In each of the other memory module categories the corresponding elasticity is even larger. The estimates range from -33.0 in the 128MB PC133 category to -41.4 in the 256MB PC133 category. These elasticities are a powerful illustration of the potential that price search has to create a real Bertrand paradox. The familiar Lerner index formula for equilibrium markups with competing single-product retailers, $(P - MC)/P = -1/\eta$, implies that if retailers sold only low-quality products, their markups over marginal cost would be 2.5% to 4%. Even efficient, non-advertising retailers would find it difficult to survive with such markups.

A second striking empirical result in Table 3 is that low quality memory is an effective loss leader. In the second column of the table, note that the coefficient on $\log(1 + PLowRank)$ is negative and highly significant. This means that controlling for the prices a website charges for medium- and high-quality memory and for the lowest price at which low-quality memory is available, it sells more medium-quality memory if its price for low-quality memory is lower. Ordinarily one would expect the opposite relationship because low- and medium-quality memory are substitutes. Apparently, however, what many consumers are doing is using Pricewatch to find a set of websites offering the lowest prices for low-quality memory, clicking through to those websites to examine their offerings further, and then sometimes purchasing a higher quality product. The higher a firm is on the Pricewatch list, the more such consumers it attracts. The loss leader effect is very large. The -0.76 coefficient estimate indicates that moving from first to seventh on the Pricewatch list for (low-quality) 128MB PC100 memory reduces a website's sales of medium-quality 128MB PC100 memory by 65%. The first coefficient in the third column indicates that the loss-leader benefits include sales of high-quality memory as well. The -0.41 coefficient indicates that moving from first to seventh on the Pricewatch list for low-quality memory would reduce high-quality memory sales by 43%.

Although it is common in marketing to talk about loss leaders, the empirical marketing literature on the effectiveness of loss leaders has produced mixed results (Walters, 1988; Walters and McKenzie, 1988; Chevalier, Kashyap and Rossi, 2000). We are not aware of

any evidence nearly as clear our results.

A third noteworthy result is the fact that the coefficients on the site B dummy are negative and significant in all three regressions. They indicate that when the two sites charge the same prices, site B sells substantially fewer memory modules. It is especially less successful at selling high-quality memory modules. We find this to be a striking demonstration of the importance and difficulty of effective website design. Website A and website B are owned by the same firm. They were designed by the same people. They share the same telephone operators and packing employees and hence provide exactly the same level of service quality. A few attributes should make website B more attractive to some customers: it had slightly lower shipping charges for part of the sample; it offers a broader variety of additional products that may be purchased at the same time; and at the time it had a higher customer feedback rating at ResellerRatings.com. Nonetheless, it is not as effective at selling memory. One hypothesis the firm suggested to us is that some customers may prefer to buy memory from a website (like site A) that specializes in memory.

The significance levels in this table are similar to what we find in the other memory categories. The $\log(1 + PLowRank)$, *Weekend*, and *SiteB* variables are usually highly significant. The other variables are usually insignificant.

The one way in which the results for the 128MB PC100 category are unusual is that the own-price elasticities of medium- and high-quality memory are precisely estimated. A difficulty with our estimation strategy is that prices for medium- and high-quality memory modules change infrequently. As a result, they are highly collinear with the flexible time trends that we have included in the model. This problem is particularly severe in the 256MB categories because the effective sample size is reduced by the fact that most of the memory is sold toward the end of the data period.

We report elasticity matrices for the other memory categories in Table 4, but to save space we have not included full tables of demand estimates. The elasticity tables reveal that we find large own-price elasticities and large loss-leader benefits consistently across categories. The finding of a negative coefficient on *SiteB* is consistent, although the magnitude varies somewhat. It is larger in some of the 256MB regressions.

Table 8 presents estimates of the demand for AMD Athlon CPUs sold through website

C. The first column gives parameter estimates of the demand for low quality (bare) CPUs. The second column reports on the demand for the website’s medium-quality offering: a CPU with an AMD-approved fan attached. The third column reports on demand for the website’s high-quality offering: a CPU with a premium quality dual fan. The demand for low-, medium-, and high-quality CPUs is quite similar to the demand for low-, medium-, and high-quality memory modules in a number of ways. The demand for the low-quality product is highly dependent on the website’s position on the Pricewatch list. (Moving from first to seventh is predicted to lower demand by about 80%. This corresponds to an own-price elasticity of -27.7 for the low-quality product.) There is a significant loss-leader benefit: when the firm’s price for a bare CPU puts it near the top of the Pricewatch list it sells many more of the medium-quality CPU/fan combinations. For part of our sample the firm sold a CPU with a non-AMD-approved fan rather than a bare CPU as its lowest-quality product. It abandoned this experiment after just a few weeks because too many customers were buying the low-quality product instead of upgrading. The coefficient estimate on the *LowHasFan* variable in the third column indicates that there is clear statistical evidence that the alternate product lineup led to many fewer sales of CPUs with premium-quality fans. As in many of the product categories the own-price elasticities of medium- and high-quality products cannot be precisely estimated.

To save space we do not report estimates of the demand for CPUs through site B. The estimates are qualitatively similar. The own price elasticity of site B’s low-quality offering (a bare CPU) is estimated to be -28.3.

6.3 The mechanics of obfuscation: incomplete consumer search

The first explicit model of obfuscation we discussed in section 2 was Stahl’s model of search with heterogeneous consumers. In Stahl’s model, consumers do not learn all prices. We noted that any action that raised the cost of learning about each firm’s offerings or that forced more consumers to conduct firm-by-firm searches could be regarded as a form of obfuscation. In this section we exploit a feature of our dataset to provide additional evidence that consumers are not fully informed about the prices charged by the Pricewatch retailers.

Our data do not provide any easy way to see whether consumer learning about the low-quality offerings is incomplete. The market share advantage of firms ranked more highly on the Pricewatch list could reflect that some consumers have not examined the offers by firms farther from the top of the list, but it could also be attributable purely to price differences and idiosyncratic consumer preferences. The structure of our dataset does, however, provides a nice opportunity to examine whether consumer learning about higher quality products and prices is incomplete. We have data on two websites that offer identical products. If all consumers learned about each site’s prices for medium- and high-quality memory, then we would expect that most of consumers buying medium-quality memory from one of the two sites would buy it from the cheaper site. Differences between low-quality prices should not help predict which site sells more medium-quality memory.³⁰

To provide a straightforward analysis of how consumers medium- and high-quality memory choose between site A and site B we estimate simple logit models on the consumer-level data using a dummy for whether each consumer chose to buy from site A (versus site B) as the dependent variable. Formally, we assume each customer i choosing to buy a category c quality q product from site A or site B at time t makes this choice to maximize his utility u_{iwcqt} over the website w with

$$u_{iwcqt} = \beta_1 \log(1 + PLowRank_{wcqt}) + \beta_2 \log(P_{wcqt}) + \beta_3 TimeTrend_t + \eta_{cq} + \epsilon_{iw},$$

with η_{cq} a quality-category fixed effect and ϵ_{iw} a standard logit error. The first right-hand side variable reflects a site’s position on the Pricewatch list for the category. The second is the site’s price for the product that being purchased. Note that the estimated coefficients for different product categories and different quality levels are constrained to be equal (for parsimony). If consumers are fully informed, one would expect that β_1 would be zero and that β_2 would be negative.

The first column of Table 7 reports estimates from a regression run on a dataset containing observations from all four memory categories and on both medium- and high-quality

³⁰One could imagine that the coefficient on $\log(1 + PLowRank)$ was significant in the medium-quality regressions we ran earlier, not because consumers were less than fully informed about prices, but only because a site’s rank on the low-quality product is correlated with its rank on the medium-quality product (which is not well-defined and, therefore, not included in the regression). This alternate interpretation is no longer possible when we study the relative shares of the two websites with identical products sold at prices we know.

modules. The coefficients on $\log(1 + PLowRank)$ and $\log(P)$ are both significant. The standard deviation of the difference in $\log(1 + PLowRank)$ across websites is more than ten times larger than the standard deviation of the difference in $\log(P)$, so the coefficient estimates can be thought of as indicating the a firm’s position on the low-quality Pricewatch list has a slightly stronger effect than does its price for the item being purchased. This suggests that there are a substantial number of consumers who are informed about both sites’ prices and a substantial number who are not.

The second and third columns report estimates from the same regressions run separately on medium-quality and high-quality memory modules. In both cases the results indicate clearly that the low-quality module’s position on the Pricewatch list is an important determinant of sales of medium- and high-quality modules. In the medium-quality case, we are unable to find significant evidence that the relative prices of medium-quality modules on the two sites affects where consumers buy them. In the high-quality case, the sites’ prices for the high-quality products appear to be playing a more important role.

6.4 The mechanics of obfuscation: add-ons and adverse selection

The second explicit model of obfuscation we discussed in Section 2 was a variant on the rational expectations add-on pricing model of Ellison (2003). In this model the practice of posting prices for damaged goods and requiring consumers to visit a site to learn its prices for higher quality goods raises prices even if consumers (in equilibrium) correctly anticipate what each retailer’s prices for the higher quality goods will be. The main reason for this is that adopting add-on pricing creates an adverse selection problem that discourages price cutting: a firm that cuts prices will disproportionately attract cheapskates who are less likely than other customers to buy higher-quality products.

In the section on basic results, we verified one implication of this model, that the low-quality products were attracting customers who upgraded to the higher quality products, or, in other words, that the low-quality products were effective loss leaders. This was reflected in negative coefficients on $\log(1 + PLowRank)$ in the medium- and high-quality regressions. In this section, we examine the evidence on the theory’s more important prediction, that the use of the add-on pricing creates an adverse selection problem (which would be expected to

lead to higher prices). In our demand system, such an adverse selection problem would be reflected in the coefficient on $\log(1 + PLowRank)$ being larger in the low-quality regression than in the medium- and high-quality regressions. The larger coefficient indicates that lowering the price of the low-quality product would lead to a larger proportional increase in low-quality sales than in medium- or high-quality sales. Table 6 reports the coefficient estimates on the $\log(1 + PLowRank)$ variable in the seventeen demand regressions we ran: three quality levels in each of the four memory module categories; three quality levels in the regression examining CPUs sold through site C; and two quality levels in the regression examining CPUs sold through site B.

In all cases, the pattern of the coefficient estimates is consistent with the presence of an adverse selection effect. In every category (*i.e.* in every column of the table) the coefficient estimate from the low-quality regression is larger than the coefficient estimates from the medium- and high-quality regressions. Most of the differences are statistically significant. To offer a sense of the magnitudes suggested by these coefficient estimates, consider the estimates from the 128MB PC100 column. A move from first to third on the Pricewatch list would result in 64%, 41%, and 25% decreases in unit sales of low-, medium-, and high-quality modules, respectively. An alternative way for us to offer intuition for the magnitude of this adverse selection effect (and a way to do so avoiding any worries about results being driven by functional form assumptions) is to look at the firm's quality mix using simple sample means. For example, when site A or site B is first on one of the Pricewatch lists for 256MB memory, 63% of its unit sales are low-quality memory. On days when one of them is in tenth place, only 35% of the unit sales are low-quality memory.³¹

Finally, we also find Table 6 striking for the consistency of the estimates across product categories. We cannot reject the hypothesis that the six coefficient estimates in the first row are identical. Nor can we reject the equality of the five estimates in the second row, nor of the five estimates in the third row. Given that the products in the different categories differ in price, that moving up or down one rank involves changing the item price by a different

³¹Medium quality sales increase from 16% of total unit sales to 21% and high-quality sales from 21% to 40%. Total sales, of course, are much lower when the firm is in the tenth position so these "increases" are decreases in unit sales of the higher quality products being offset by even larger decreases of the low-quality product

number of dollars in the different categories, and that the attributes that distinguish the quality levels differ between memory modules and CPUs, we find this regularity intriguing.

6.5 Instrumental variables estimates

We have assumed so far that variation in a website's rank on the Pricewatch list and in its medium- and high-quality prices is exogenous. We have argued that this is a reasonable assumption, but in this section will investigate the robustness of our conclusions to two instrumental variables strategies.

We feel that there are two primary sources of exogenous variation in our websites' prices (relative to the prices of competing websites). First, because monitoring other firms' prices and changing one's prices involves a time cost, the websites' ranks on the Pricewatch lists drift as other sites change prices during periods in which our firm is inattentive. Second, the firm will raise or lower its prices in response to cost shocks. One source of such shocks is unexpected employee absences that make it more difficult to process a normal day's order flow. Another is variation in the firm's wholesale acquisition costs. The latter are more common, but less clearly suitable for use as instruments because (1) they will not lead to variation in a firm's Pricewatch rank if wholesale prices for low-quality modules are common to all firms; and (2) prices of other firms' higher quality products are an omitted variable that may be correlated with wholesale prices of medium- and high-quality memory (especially in the medium- and high-quality demand regressions).

We have constructed two sets of instruments we will use to estimate demand for 128MB PC100 memory modules, which we call "other speed" instruments and "cost" instruments. The "other speed" instruments are the contemporaneous values of $\log(1 + PLowRank)$, $\log(PMid)$, and $\log(PHi)$ for the same website but for a different product, 128MB PC133 memory. These variables may reflect both sources of exogenous variation in PC100 memory prices: the website's ranks in the two categories could drift in unison during periods of inattention; and if employees do not show up for work, the firm may raise its prices in all product categories.³² If the demands for PC100 and PC133 memory are independent (recall

³²Memory prices have a strong downward trend so a firm's rank in each category will drift up during long periods of inattention. In practice, the correlation is strong but far from perfect. In the 128MB PC100 category, for example, the R-squared of a regression of $\log(1 + PLowRank_{100})$ on $\log(1 + PLowRank_{133})$,

that the products are not substitutes for most consumers), then the PC133 ranks would be exogenous. The “cost” instruments for $\log(1 + PLowRank)$, $\log(PMid)$, and $\log(PHi)$ are derived from our data on the firm’s acquisition costs for low-, medium-, and high-quality 128MB PC100 memory modules.³³ If the wholesale prices for low-quality memory our firm faces are not perfectly correlated with the prices other firms face, then the low-quality wholesale price should be correlated with the firm’s Pricewatch ranks.³⁴

Table 8 reports estimates of the demand equations for 128MB PC100 memory modules (comparable to those in Table 3) obtained using the two sets of instruments. In the first low quality regression (in the first column of the table), the coefficient estimate on the $\log(1 + PLowRank)$ variable is unchanged and the variable remains highly significant. Again we find that the demand for low-quality products is extremely price-sensitive. The coefficient estimates on the $\log(1 + PLowRank)$ variable in the next two columns are reduced from their values in the uninstrumented variables by about one-half. The standard errors are substantially increased, failing to provide significant evidence either that loss leader effects exist or that our previous estimates of the loss leader benefit were too large. We can, however, reject the hypothesis that the coefficient estimates in the medium- and high-quality regression are as large as the coefficient estimate from the low-quality regression. Hence, the results confirm that the loss leader pricing creates an adverse selection problem. The standard errors on the $\log(PMid)$ and $\log(PHi)$ variables are also increased.

Again, with the “cost” instruments, the $\log(1 + PLowRank)$ variable remains significant in the low-quality regression confirming our finding about extreme price-sensitivity. Again,

the other rank instruments, and the exogenous variables in the demand equation is 0.58. The coefficient on $\log(1 + PLowRank_{133})$ has a t-statistic of 13. There is much less variation in $\log(PMid)$ and $\log(PHi)$ and they can be predicted almost perfectly in such regressions, e.g. we get an R-squared of 0.995 in the 128MB PC100 data.

³³These data are only available on days on which the firm made wholesale purchases. Gaps of one or a few days are common. We fill in the missing data by interpolating linearly. The instruments for $\log(PMid)$ and $\log(PHi)$ are simply the log of the cost of the product in question. Our instrument for $PLowRank$ is the log of a predicted rank obtained by adding a constant to what the firm’s rank would be if it priced at cost.

³⁴In practice, the cost-based instrument for $PLowRank$ has predictive power, but less than the speed-based instrument. In the 128MB PC100 category, for example, the R-squared of a regression of $\log(1 + PLowRank)$ on the cost-based instruments and exogenous variables has an R-squared of 0.48 with the t-statistic on the instrument being about five. The instruments for $\log(PMid)$ and $\log(PHi)$ again have very strong predictive power, e.g. in the 128MB PC100 category we again get R-squareds above 0.99 in first-stage regressions.

the coefficients on this variable become insignificant in the other two columns leaving us unable to say that the IV regressions provide additional support for our loss leader results. The standard errors on the $\log(PMid)$ and $\log(PHi)$ variables become much larger. We believe that this reflects that the instruments eliminate most of the true exogenous variation in these variables (which comes from these prices gradually becoming higher relative to the low-quality price in between adjustments and then being much smaller immediately after adjustments).

7 Markups

This section examines price-cost margins. It has two motivations. First, the future of e-retail depends on the balance between search and obfuscation. At one extreme, retail firms would struggle to survive. At the other, search engines wouldn't be worth using. Markups provide descriptive information about how the battle is playing out in the Pricewatch universe. Second, analyzing markups can provide a more complete understanding of the mechanisms by which obfuscation affects prices.

Table 9 presents our estimates of the actual markups. The table presents revenue-weighted average percentage markups for each of the four categories of memory modules and for CPUs sold through each of the two sites.³⁵ The dollar markups were obtained by adding the standard shipping and handling charge to the advertised item price, and then subtracting the wholesale acquisition cost, credit card fees, an approximate shipping cost, an estimate of marginal labor costs for order processing, packing, and returns, and an allowance for losses due to fraud.³⁶ In each category the sample period is that used in the demand estimation, omitting observations from December 2000 (for which we do not have cost data).

In the two 128 MB memory categories, the markups for low-quality products are slightly negative. Prices have not, however, been pushed far below cost by the hope of attracting customers who can be talked into upgrading. Markups are about 16% for medium-quality modules, and about 27% for high-quality modules. Averaging across all three quality levels,

³⁵The percentage markup is the percentage of the sale price, i.e. $100(p - mc)/p$.

³⁶The labor and shipping costs were chosen after discussions with the firm, but are obviously subject to some error.

markups are about 8% and 12% in the two categories. This corresponds to about five dollars for a PC100 module and ten dollars for a PC133 module.

The firm's average markups in the 256MB memory categories were higher: 13% and 16% in the two categories. Part of the difference is due to the fact that a higher fraction of consumers buy premium quality products, but the largest part comes from the markups on low-quality memory being substantially higher.³⁷

Finally, the average percentage markups on Athlon CPUs are much lower: 3% at website B and 6% at website C. The difference in markups is not as dramatic if one looks at them in dollar terms. The average markups on CPUs at the two sites are about six and twelve dollars, respectively. The largest part of the difference in average percentage markups is due to markups on the upgraded CPU-fan combinations being lower than markups on high-quality memory modules. Markups on low quality CPUs are also lower than markups on low-quality memory modules.

An obvious question to ask is whether the markups reported above are what one would expect under static Nash equilibrium pricing given the strength of the adverse selection effect we identified in the previous section. (There are, of course, many reasons for prices to be higher or lower, such as repeated-game collusion or desire to build short-run market share.) This, however, is something that one can not do with our demand estimates. One reason is that we have not been able to obtain precise estimates of own-price elasticities for medium- and high-quality products. A second is that the functional forms for demand we are using are not well-behaved in some of the counterfactuals that must be considered when doing a Nash equilibrium analysis. This is partially due to data limitations and partially due to functional form choices we made in order to make the demand results easier to interpret.³⁸

³⁷Part is also due to the sample period being different. The revenue weighted estimates put a great deal of weight on the latter part of the data. The 256MB data end 6 weeks before the 128MB data. Markups on 128MB modules were relatively low in the last six weeks.

³⁸Examples of unavoidable data limitations are that it is hard to separately identify the effect of the firm's low quality price and its Pricewatch rank on demand since they are highly correlated and that there is little variation in medium- or high- quality prices that is not collinear with the time trends. Examples of choices we made include our use of a firm's Pricewatch rank as the primary explanatory variable, which is not appropriate if one wants to predict what would happen if the last firm on the list raised its price toward infinity, and the specification of $\log(QLow)$ and $\log(QMid)$ as linear functions of $\log(PHi)$, which implies that a firm could also earn unbounded profits by setting $PLow$ and $PMid$ above cost and sending

Although we cannot provide a structural analysis of markups, we think that we can provide evidence that the observed markups are roughly consistent with firms' approximately maximizing static profits. We think that our demand system accurately reflects demand at the prices one typically observes in the data. We therefore feel comfortable computing projected profits at prices charged in the data and can examine whether there appear to be strong incentives to deviate from the actual prices to other prices in the same range.

Consider, for example, the prices from October 12, 2000 in Figure 1. The markups on this day are fairly typical. Our retailer bought low-quality modules wholesale on the previous day for \$72.³⁹ Shipping and handling charges are about \$1 below the sum of actual shipping costs and other marginal costs, so the lowest-priced firm, listing at \$68, is selling low-quality modules at about \$5 below cost and the sixth firm, listing at \$74, is at about \$1 above cost. Given that the low-priced firms make many more sales, the quantity-weighted average markup for low-quality modules is probably about -2%, which is just a little below the sample mean. Site B's markups on medium- and high-quality memory on that day were also typical: the wholesale prices were \$77 and \$94, and it charged \$103.49 and \$148.50.

To illustrate the effect on a firm's profits of moving positions on the Pricewatch list, Table 10 contains results from a series of scenarios involving site B. In particular, it has predicted quantities and predicted profits for site B assuming site B had occupied each of the twelve places on the list in Figure 1. Predicted sales of low-quality memory drop sharply from 54.5 units if site B occupies the first position to 3.4 units if it occupies the twelfth position. Sales of high-quality modules decline much more gradually, from 2.8 units in the first position to 1.3 units in the twelfth. The sixth column reports estimates of site B's expected gross profit from being in each position under the assumption that site B would have used the wholesale prices from the previous day to calculate its costs. Site B would most have liked to be in the fifth position, where it would have expected to earn \$154. Note that profits are fairly flat over a wide range of markups. For example, the firm's profits in the third and tenth positions are nearly equal. The one position that is clearly worse than the others is the first.

PHi toward infinity.

³⁹Given that the prices are from 9:01am Pacific time, the \$72 figure is likely to be the latest information most firms had on costs before they set the observed prices.

What would an ϵ -Nash equilibrium look like? The requirement is that no firm can gain more than ϵ by changing its price and moving to another location on the list. If markups are too high this will fail because everyone will want to move to the top of the list, and if they are too low it will fail because everyone will want to move to the bottom. At some intermediate level, the losses that the top few firms incur selling low-quality modules below cost could roughly offset the profits they derive from their greater sales of premium-quality modules and thus we could have an ϵ -Nash equilibrium with price dispersion.

Although site B is not indifferent between the various places on the list (and would find the first position quite unattractive), we think of the profit estimates in Table 10 as suggesting that play is close to an ϵ -Nash equilibrium with a reasonable ϵ for a few reasons.

First, profits need not be exactly equal in equilibrium. They need only be close enough to prevent deviations.⁴⁰ Site B cannot switch its price to \$74 and take the fifth spot. It can claim this spot only by charging \$73, which would result in a profit of \$144 rather than \$154.

Second, heterogeneity in website design will lead to firms' having different preferences over positions on the Pricewatch list. Recall that site A makes higher sales than site B at the same prices, and is more successful in inducing customers to buy medium- and high-quality memory. The seventh column of Table 10 reports estimates of site A's gross profits. The fifth position is also the most attractive for firm A, but its preferences over the other positions are somewhat different.⁴¹ In particular, it does relatively better in the top few slots. This is consistent with our earlier observation that site A tends to set lower prices than site B. Sites A and B have similar designs. Other websites that employ different tactics, *e.g.*, requiring that consumers call on the phone and talk to salespeople, may face quite different demand functions. It does not seem unreasonable to suppose that such heterogeneity might be able to account for why Computer Craft is willing to occupy the top slot.

Third, cost heterogeneity will also contribute to firms having different preferences over

⁴⁰Ellison and Fudenberg (2003) provide several examples of how such "market impact effects" can create a multiplicity of asymmetric pure strategy equilibria.

⁴¹The attractiveness of the fifth slot is due to an unusual feature of the prices in Figure 1. The fifth-ranked firm's markup is a full four dollars above the markup of the third-ranked firm. This may be a situation where our rank-based demand system is overestimating profits.

positions on the Pricewatch list. In addition to the standard cost heterogeneities, Pricewatch retailers may differ significantly in their expectations about wholesale prices. For example, the most recent time prior to October 11th that our firm made a wholesale purchase of low-quality memory was on Monday October 9th when it paid \$78. Some of the listed firms might not have checked wholesale prices since then, or might have made the mistake of using acquisition costs rather than replacement costs when pricing inventory acquired three days earlier. The eighth column of Table 10 lists the prices that site B would have expected if it used the three-day-old costs. The ninth position is now the most attractive. It is also true that many retailers would not ship orders received on the 12th until the 13th or would wait until the 13th to replace shipped modules in inventory. It turns out that before the end of the day on the 13th, low-quality modules could be purchased wholesale for \$68. The final column of the table lists the profits site B would have expected if it had anticipated these lower costs. The first position is now the most attractive. Another potential explanation for the top-ranked firm's behavior is thus that it might have had different information about costs.

Fourth, the ϵ in ϵ -Nash equilibrium would include monitoring and decision costs and should not be thought of as extremely small. Our retailer, for example, has just one manager. Monitoring dozens of Pricewatch lists is just a small part of his duties and he does not have nearly enough time to make minor adjustments to every price every hour. Hence, we should not expect to see all firms maximizing the profits we calculated at all times. Decision costs of some kind are necessary to account for the price inertia that is present in the data. On October 12th, for example, the top six firms on the list in Figure 1 left their prices fixed throughout the afternoon. The first change to the top half of the list came later in the day: Augustus Technology cut its price to \$73 and took over the fifth spot.⁴²

Although we think that the observed markups are roughly consistent with a static Nash model of competition, two features of the data stand out as very hard to reconcile to such a model. First, the markups for low-quality 256MB memory are higher than one would

⁴²There were five other changes among firms lower on the list. The eighth-place firm raised its price and left the top 12. Two other firms cut their prices to \$76 to enter the first page in the ninth and tenth positions. Two other firms cut their prices to \$77, one to enter the first page and one to stay on it.

predict even if one ignored the loss leader benefits. Second, there is some variation in markups over time. The two features are not completely unrelated. The fact that average markups for low-quality memory are positive is entirely attributable to two subperiods. The firm sold a large number of low-quality 256MB modules at about ten dollars above cost in September-October 2000 and a large number at about five dollars above cost in February-March 2001. We think we understand what happened in the former period. A small number of retailers found an obscure supplier willing to sell them 256MB modules at a price far below the price offered by the standard wholesale distributors.⁴³ As a result, there were effectively six or fewer retailers competing in these two months rather than dozens.⁴⁴ We do not know why higher markups prevailed in February and March of 2001. In this period markups were relatively high in all four memory categories.

8 Conclusion

In this paper we have noted that the extent to which the Internet will reduce consumer search costs is not clear. While the Internet clearly facilitates search, it also allows firms to adopt a number of strategies that make search more difficult. In the Pricewatch universe, we see that demand is sometimes remarkably elastic, but that this is not always what happens.

The most popular obfuscation strategy for the products we study is to intentionally create an inferior quality good that can be offered at a very low price. Retailers could, of course, avoid the negative impact of search engines simply by refusing to let the search engines have access to their prices. This easy solution, however, has a free rider problem—if other firms are listing, I will suffer from not being listed. What may help make the

⁴³The first retailer to have found the supplier appears to have found it on July 10. On that day, when the firm that supplied us with data bought modules wholesale for \$270, PC Cost cuts its retail price to \$218—a full \$51 below the next lowest price. The firm remained all alone with a price at least \$37 less than that of any other firm for three weeks. We infer from this that its acquisition costs must have been very low.

⁴⁴We cannot provide any estimates about how the markups of the low-cost retailers varied in August as the number of low-cost retailers increased from one to (we think) six, because the retailer from which we got the data appears to be the last of the six to have found the low-priced supplier. At that point, in early September, low-quality memory modules were selling for about \$10 above cost. Markups remained steady and no other retailers appear to have found the supplier for the next two months. At the end of October wholesale prices fell dramatically, apparently eliminating the advantage. At this point several retailers that had been absent for months jumped to the top of the Pricewatch list and margins on low-quality memory quickly dropped to approximately zero.

obfuscation strategy we observe popular is that it is hard not to copy it—if a retailer tries to advertise a decent quality product with reasonable contractual terms at a fair price it will be buried behind dozens of lower price offers on the search engine’s list. The endogenous-quality aspect of the practice makes it somewhat different from previous bait-and-switch and loss-leader models, and it seems that it would be a worthwhile topic for research.⁴⁵ We would also be interested to see more work integrating search engines into models with search frictions, exploring other obfuscation techniques (such as individualized prices), and trying to understand why adoption of price search engines has been slow.

Large online retailers have very high expenses. Although they can avoid the expense of setting up hundreds of retail locations, they have to handle every item they ship individually. It is much cheaper to ship a truckload of merchandise to a store than to ship a similar number of individual packages with UPS. It is cheaper to have a customer take packages off the shelf and bring them to a cashier than it is to hire someone to pick packages off shelves and put them in boxes. Large online retailers also have very high expenses of other types: they have been very intensive users of managerial talent; they have spent a lot on advertising; and they have enormous website development costs.⁴⁶

Given these expenses, large e-retailers cannot be competitive with Walmart’s prices. This is not necessarily worrisome. Many people may be willing to pay a premium to avoid going to the mall and to take advantage of superior product variety. A more serious problem is that managerial, advertising, and website development costs seem more like fixed costs than marginal costs. If price search engines become more popular, and better sources of information develop on the reliability of retailers, then vigorous price competition may push online prices below the level necessary to cover these fixed costs.

What will happen to e-retail if price search flourishes? One thing we have seen can happen is that demand can become incredibly elastic, so elastic that retailers could never break even. While this is a frightening possibility to retailers, it does not seem very likely to us. It seems likely that retailers would make significant efforts to ensure that price search

⁴⁵Simester’s (1995) model seems the most similar to the practice. We would imagine, however, that what makes the low prices on Pricewatch have advertising value is that the offerings are sufficiently attractive so as to force a retailer to set low prices for its other offerings to avoid having everyone buy the advertised product.

⁴⁶Amazon, for example, reported more than \$250 million in R&D costs in 2000.

engines never work too well, and markups will end up at least as high as they are in the (unbranded) Pricewatch universe. Similar markups would force large retailers to cut back substantially on management, advertising, and website costs, but it is certainly plausible that they could do so. Another possibility is that e-retail could end up with a very different market structure than traditional retail. Small e-retailers without venture capital realized long ago that the Internet allows one to operate a retail store with virtually no advertising or website development costs. At present, small e-retailers are far more efficient than their more famous competitors. Whether they will remain more efficient and come to dominate price-search mediated markets or whether they will be pushed aside as larger firms copy their model is another interesting question.

References

- Bakos, J. Yannis (1997): "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science* 43, 1976 - 1993.
- Baye, Michael R., and John Morgan (2001): "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review* 91, 454-474.
- Baye, Michael R., and John Morgan (2003): "Information Gatekeepers and Price Discrimination on the Internet," *Economics Letters* 76, 47-51.
- Baye, Michael R., John Morgan, and Patrick Scholten (2003): "Price Dispersion in the Small and the Large: Evidence from an Internet Price Comparison Site," *Journal of Industrial Economics*, forthcoming.
- Brown, Jeffrey R., and Austan Goolsbee (2002): "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy* 110, 481-507.
- Brynjolfsson, Erik, and Michael Smith (2001): "Consumer Decision-making at an Internet Shopbot: Brand Still Matters," *Journal of Industrial Economics* XLIX, 541-558.
- Burdett, Kenneth, and Kenneth L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica* 51, 955-969.
- Chevalier, Judith, and Austan Goolsbee (2003): "Measuring Prices and Price Competition Online: Amazon vs. Barnes and Noble," *Quantitative Marketing and Economics* 1, 203-222.
- Chevalier, Judith, Anil Kashyap, and Peter Rossi (2000): "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," University of Chicago Graduate School of Business, mimeo.
- Clay, Karen, Ramayya Krishnan, and Eric Wolff (2001): "Prices and Price Dispersion on the Web: Evidence from the Online Book Industry," *Journal of Industrial Economics* 49, 521-540.
- Deneckere, Raymond, and R. Preston McAfee (1996): "Damaged Goods," *Journal of Economics and Management Strategy* 5, 149-174.
- Diamond, Peter (1971): "A Model of Price Adjustment," *Journal of Economic Theory* 3, 156-168.
- Ellison, Glenn (2003): "A Model of Add-on Pricing," mimeo, MIT.
- Gabaix, Xavier and David Laibson (2004): "Competition and Consumer Confusion," mimeo, Harvard University and MIT.
- Janssen, Maarten, and Jose Luis Moraga (2000): "Pricing, Consumer Search and the Size of Internet Markets," Tinbergen Institute Discussion Paper 2000-0042.

Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso (2001): "Internet Car Retailing," *Journal of Industrial Economics* 49, 501-519.

Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso (2003): "Cowboys or Cowards: Why are Internet Car Prices Lower?," mimeo.

Simester, Duncan (1995): "Signalling Price Image Using Advertised Prices," *Marketing Science* 14, 166-188.

Stahl, Dale O. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review* 79, 700-712.

Stahl, Dale O. (1996): "Oligopolistic Pricing with Sequential Consumer Search and Heterogeneous Search Costs," *International Journal of Industrial Organization* 14, 243-268.

Varian, Hal (1980): "A Model of Sales," *American Economic Review* 70, 651-659.

Walters, Rochney (1988): "Retail Promotions and Retail Store Performance: A Test of Some Key Hypotheses," *Journal of Retailing* 64, 153-180.

Walters, Rochney, and Scott McKenzie (1988): "A Structural Equation Analysis of the Impact of Price Promotions on Store Performance," *Journal of Marketing Research* 25, 51-63.

White, E. (2000): "No Comparison. Shopping 'Bots' Were Supposed to Unleash Brutal Price Wars. Why Haven't They?" *The Wall Street Journal*, October 23, 2000, R18.

BRAND	PRODUCT	DESCRIPTION	PRICE	SHIP	DATE/HR	DEALER/PHONE	ST	PART#
Generic	PRICE FOR ONLINE ORDERS ONLY - 128MB PC100 SDRAM DIMM - 8ns Gold leads	- * LIMIT ONE - Easy installation - in stock	\$ 68	9.69 INSURED	10/12/00 12:40:05 AM CST	Computer Craft Inc. 800-487-4910 727-327-7559 Online Ordering	FL	MEM-128-100PCT
Generic	ONLINE ORDERS ONLY - 128MB SDRAM PC100 16x64 168pin	- * LIMIT ONE	\$ 69	INSURED\$9.95	10/11/00 10:59:56 PM CST	Connect Computers 888-277-6287 949-367-0703 Online Ordering	CA	-
Generic	PRICE FOR ONLINE ORDER - 128MB PC100 SDRAM DIMM	- * LIMIT ONE - - InStock, 16x64-Gold Leads	\$ 70	10.75	10/11/00 2:11:00 PM CST	1st Choice Memory 949-888-3810 -- P.O.'s accepted Online Ordering	CA	-
Generic	PRICE FOR ONLINE ORDER - 128mb True PC100 SDRAM EEPROM DIMM16x64 168pin 6ns/7ns/8ns Gold Leads	- * LIMIT ONE - in stock - with Lifetime Warranty	\$ 72	9.85	10/10/00 11:30:39 AM CST	pcboost.com 800-382-6678 -- P.O.'s accepted Online Ordering	CA	-
Generic	IN STOCK, 128MB PC100 3.3volt unbuffered SDRAM Gold Lead 168 Pin, 7/8ns - with Lifetime warranty	- * LIMIT ONE Not compatible with E Machine	\$ 74	10.95- UPS INSURED	10/11/00 12:44:00 PM CST	Memplus.com 877-918-6767 626-918-6767	CA	- 880060
Generic	PRICE FOR ONLINE ORDERS ONLY - 128MB True PC100 SDRAM DIMM - 8ns Gold - warranty	- * LIMIT ONE	\$ 74	10.25	10/9/00 6:53:25 PM CST	Portatech 800-487-1327	CA	-
House Brand	128MB PC100 3.3volt SDRAM 168 Pin, 7/8ns - with LIFETIME WARRANTY	- * LIMIT ONE	\$ 74	10.50 FedEx	10/11/00 10:20:23 AM CST	1st Compu Choice 800-345-8880 800-345-8880	OH	-
Generic	128MB 168Pin TRUE PC100 SDRAM - OEM 16X64	DIMM16x64 168pin 6ns/7ns/8ns Gold Leads	\$ 75	\$10	10/11/00 2:37:00 PM CST	Sunset Marketing, Inc. 800-397-5050 410-626-0211 -- P.O.'s accepted	MD	-
Generic	128MB 16x64 PC100 8ns SDRAM.	- * LIMIT ONE	\$ 77	10.90	10/12/00 9:37:45 AM CST	PC COST 800-877-9442 847-690-0103 Online Ordering	IL	-
Generic	IN STOCK, PC100, 128MB, 168pins DIMM NonECC, - with Lifetime warranty	- * LIMIT 5	\$ 77	\$10.95 & UP For UPS Ground	10/9/00 5:11:10 PM CST	Augustus Technology, Inc 877-468-5181 909-468-1883 Online Ordering	CA	-
Generic	128MB PC100 8NS 16x64 SDRAM - one year warranty	- * LIMIT ONE	\$ 78	Ups Ground \$10.62	10/11/00 5:16:36 PM CST	Computer Super Sale 800-305-4930 847-640-9710 Online Ordering	IL	-
Generic	PRICE FOR ONLINE ORDERS ONLY - PC100 128MB NonBuffered, NonECC 16x64 SDRAM DIMM 3.3V 8ns mem module	- * LIMIT ONE - with lifetime warranty	\$ 78	10.95	10/5/00 6:29:59 PM CST	Jazz Technology USA, LLC 888-485-8872 909-869-8859	CA	ME-GBP100128

Figure 1: A sample Pricewatch search list: 128MB PC100 memory modules at 12:01pm ET on October 12, 2000.

Prices for 128MB PC100 Memory Modules

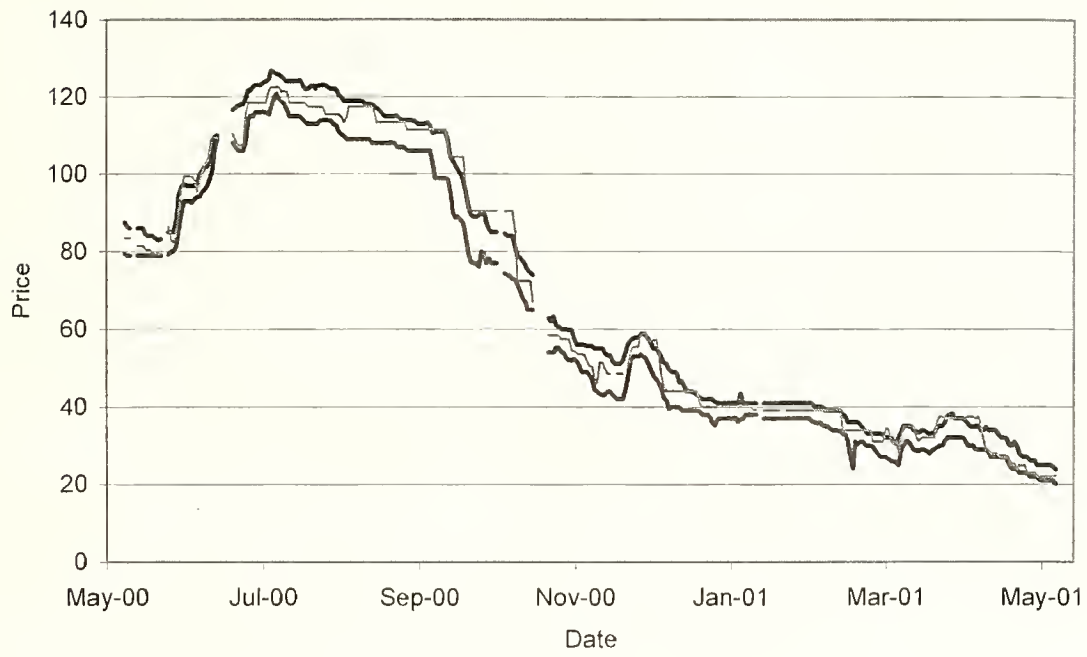


Figure 2: Prices for 128MB PC100 memory modules: the lowest and 12th lowest prices on Pricewatch and website B's price

Memory Spec. Chart - PC3200 DDR 512MB (Select Your Memory Module)		
<p>Samsung/Micron or Major 512MB PC3200 [ADD \$25]</p> <ul style="list-style-type: none"> ● CAS 2.5 Latency ● Hand Picked 5ns ● 6 Layer Low Noise Shielded PCB Board ● 32x8 DRAM Type ● Samsung/Micron or Major Brands ● Return Shipping Paid ● No Restocking Fee ● Satisfaction & Compatibility Guaranteed ● Lifetime Warranty ● 15 Days Full Refund ● Memory Tested Before Ship Out ● Copper Heat Sink - Cool Down the Memory up to 40% 	<p>Industry Standard 512MB PC3200 [ADD \$15]</p> <ul style="list-style-type: none"> ● CAS 2.5 Latency ● Hand Picked 5ns ● 6 Layer Low Noise Shielded PCB Board ● 32x8 DRAM Type ● Industry Standard DRAM Chips ● 7 Days No Restocking Fee ● Return Shipping not Paid ● Improved Compatibility ● Lifetime Warranty ● Aluminum Heat Sink - Cool Down the Memory up to 35% 	<p>OEM 512MB PC3200</p> <ul style="list-style-type: none"> ● CAS 3 Latency ● 4 Layer Module Board ● 64x4 DRAM Type ● OEM DRAM Downgrade Chips ● 20% Restocking Fee According to the Market Value ● Verify Compatibility with Memory Configurator ● Return Shipping not Paid ● 9 Months Warranty

Figure 3: A website designed to induce consumers to upgrade to a higher quality memory module.

Tufshop Price: \$53.81
Price (with Selected Options): \$90.36



Super Buys

Make processor upto 30% faster or your motherboard to run with maximum efficiency. You must have this awesome value package. **(Highly Recommended)**

- Memory Upgrade - Certified intel Approved specs Memory [+\$23.11]
- Memory Upgrade - Certified AMD Approved specs Memory [+\$17.35]



Bonus Buys

Consider taking advantage of these special offers. Compare and save. Purchase everything from one location and save on shipping

- Cable Upgrades - Rounded IDE and Floppy Cables (Complete Set) [+\$11.91]
- Essential Equipment - Sony Floppy Disk Drive [+\$16.84]
- Bonus Buy - 10 pack of hand thumbscrews for Case [+\$4.95]
- Bonus Buy - 12-Pc Computer Tool Kit [+\$16.98]
- Bonus Buy - RatPadzGS Ultimate Mousepad/Gaming Surface [+\$11.97]
- Bonus Buy - CD-DVD Media Cleaning Kit [+\$4.93]
- Thermal Management - Dynatron 80mm Case Fan [+\$12.87]



Related Options

Please take advantage of these special offers.

- Memory Upgrade - CAS 2 Upgrade (Dffers Performance Increase & Helps in Overclocking) [+\$18.25]
- Memory Upgrade - CAS 2.5 Upgrade (Improves Performance over Cas3 & Helps with Applications and Games) [+\$6.35]



Pretest

Have us test your merchandise before we ship to avoid costly RMAs in the future and Maximize your time

- No Pretesting
- Pretest - Standard Pretest (Avoid costly RMAs) [+\$6.97]



Memory Performance

Options to make your hardware & applications fly

- No Memory Performance Enhancements
- Memory Upgrade - 6 Layer PCB For Stability of Memory - more layer - More = Better Design [+\$8.37]



Enhancers

Options to make your hardware & applications fly

- No System Performance Enhancers
- Memory Upgrade - Thermaltake Memory Cooling Kit (Active) [+\$19.99]
- Memory Cooling - Copper Passive Memory Cooling Kit [+\$11.15]
- Memory Cooling - Aluminium Passive Memory Cooling Kit [+\$9.91]
- Memory Upgrade - Thermaltake Memory Cooling Kit (Passive) [+\$14.99]

Variable	Mean	Stdev	Min	Max
128MB PC100 Memory Modules 683 website-day observations				
<i>LowestPrice</i>	62.98	33.31	21.00	120.85
<i>Range1-12</i>	6.76	2.52	1.00	13.5
<i>PLow</i>	66.88	34.51	21.00	123.49
<i>PMid</i>	90.72	40.09	35.49	149.49
<i>PHi</i>	115.27	46.36	48.50	185.50
$\log(1 + PLowRank)$	1.86	0.53	0.69	3.26
<i>QLow</i>	12.80	17.03	0	163
<i>QMid</i>	2.01	3.31	0	25
<i>QHi</i>	1.86	3.45	0	47
128MB PC133 Memory Modules 608 website-day observations				
<i>LowestPrice</i>	71.02	37.02	21.00	131.00
<i>Range1-12</i>	5.92	2.74	2.00	15.00
<i>PLow</i>	73.68	36.69	21.00	131.49
<i>PMid</i>	98.50	41.84	22.10	153.45
<i>PHi</i>	128.46	47.56	48.50	189.50
$\log(1 + PLowRank)$	1.77	0.64	0.69	3.40
<i>QLow</i>	10.08	12.13	0	99
<i>QMid</i>	2.14	4.91	0	100
<i>QHi</i>	4.79	3.93	0	35
256MB PC100 Memory Modules 575 website-day observations				
<i>LowestPrice</i>	130.30	58.17	32.46	215.00
<i>Range1-12</i>	28.95	15.78	6.39	75.80
<i>PLow</i>	143.44	67.15	47.67	283.49
<i>PMid</i>	205.29	85.36	77.49	372.89
<i>PHi</i>	250.09	93.14	98.49	417.18
$\log(1 + PLowRank)$	1.90	0.55	0.69	2.91
<i>QLow</i>	2.87	4.61	0	33
<i>QMid</i>	1.00	1.88	0	18
<i>QHi</i>	0.87	1.63	0	16
256MB PC133 Memory Modules 575 website-day observations				
<i>LowestPrice</i>	143.90	71.97	32.46	269.00
<i>Range1-12</i>	25.49	14.01	6.60	67.00
<i>PLow</i>	155.75	77.95	43.00	291.45
<i>PMid</i>	213.14	91.91	78.49	345.45
<i>PHi</i>	249.95	92.26	104.37	392.21
$\log(1 + PLowRank)$	1.97	0.49	0.69	2.73
<i>QLow</i>	5.29	10.24	0	136
<i>QMid</i>	1.10	1.93	0	12
<i>QHi</i>	3.61	3.86	0	19

Table 1: Summary statistics for memory module data

Variable	Mean	Stdev	Min	Max
Data for Website B 392 clock speed-day observations				
<i>PLow</i>	180.80	43.84	105.25	302.45
<i>PHi</i>	219.68	43.90	142.39	341.45
$\log(1 + PLowRank)$	2.08	0.42	0.88	2.64
<i>QLow</i>	0.88	1.25	0	8
<i>QHi</i>	0.57	0.92	0	6
<i>LowestPrice</i>	173.08	42.36	97.00	298.33
<i>Range1-12</i>	11.16	5.62	3.75	36.69
Data for Website C 392 clock speed-day observations				
<i>PLow</i>	180.83	44.17	103.59	302.25
<i>PMid</i>	214.71	43.96	136.49	336.25
<i>PHi</i>	230.69	44.54	150.64	365.47
$\log(1 + PLowRank)$	1.99	0.49	0.69	2.64
<i>QLow</i>	1.07	1.54	0	12
<i>QMid</i>	0.83	1.29	0	9
<i>QHi</i>	0.52	0.85	0	4
<i>LowestPrice</i>	173.08	42.36	97.00	298.33
<i>Range1-12</i>	11.16	5.62	3.75	36.69

Table 2: Summary statistics for AMD Athlon processor data

Independent Variables	Dep. var.: quantiles of each quality level		
	Low q	Mid q	High q
$\log(1 + PLowRank)$	-1.48 [†] (15.1)	-0.76* (5.8)	-0.41* (3.2)
$\log(PMid)$	0.43 (0.5)	-6.86* (5.9)	2.72* (2.0)
$\log(PHi)$	-0.06 (0.1)	2.68 (1.8)	-5.10* (3.9)
<i>SiteB</i>	-0.26* (3.5)	-0.31* (2.9)	-0.59* (5.7)
<i>Weekend</i>	-0.51* (8.6)	-0.95* (8.1)	-0.72* (5.9)
$\log(LowestPrice)$	-1.08 (1.6)	1.10 (1.3)	0.90 (1.1)
Number of Obs.	683	683	683

Absolute value of t -statistics in parentheses. * denotes significance at the 5% level. † denotes significance at the 10⁻⁴⁸% level.

Table 3: Demand for 128MB PC100 Memory Modules

128MB PC100 Modules				128MB PC133 Modules			
	Low	Mid	Hi		Low	Mid	Hi
<i>PLow</i>	-25.0*	-12.9*	-7.0*	<i>PLow</i>	-33.0*	-11.5*	-4.1*
<i>PMid</i>	0.4	-6.9*	2.7*	<i>PMid</i>	0.2	-1.5	-0.1
<i>PHi</i>	-0.1	2.7	-5.1*	<i>PHi</i>	-0.1	-5.5*	-4.5*

256MB PC100 Modules				256MB PC133 Modules			
	Low	Mid	Hi		Low	Mid	Hi
<i>PLow</i>	-36.8*	-16.3*	-5.5*	<i>PLow</i>	-41.4*	-28.4*	-11.5
<i>PMid</i>	7.9	-3.9	-4.1*	<i>PMid</i>	-2.7	4.5	3.6
<i>PHi</i>	-3.4	2.4	-2.6	<i>PHi</i>	-1.4	-5.5	-1.0

Table 4: Price elasticities for memory modules: three qualities in each of four product classes

Independent Variables	Dep. var.: quantities of each quality level		
	Low q	Mid q	High q
$\log(1 + PLowRank)$	-1.14* (6.1)	-0.86* (4.7)	0.10 (0.3)
$\log(PMid)$	-0.34 (0.1)	-4.81 (0.8)	4.14 (0.4)
$\log(PHi)$	-2.70 (0.7)	-2.98 (0.5)	-21.48 (1.5)
<i>Weekend</i>	-0.47* (2.4)	-0.26 (1.2)	-0.81* (3.0)
$\log(LowestPrice)$	-1.47 (0.5)	2.86 (1.0)	6.77 (1.5)
<i>TimeinDays</i>	-0.02* (2.3)	-0.02* (2.3)	-0.01 (0.7)
<i>LowHasFan</i>	0.46 (1.8)	-0.25 (0.9)	-1.14* (2.3)
<i>Speed650</i>	-2.88* (2.5)	-2.97* (2.9)	-5.74* (3.0)
<i>Speed700</i>	-1.66* (1.9)	-1.67* (2.2)	-3.07* (2.1)
<i>Speed750</i>	-0.82 (1.5)	-1.19* (2.5)	-2.34* (2.5)
Number of Obs.	392	392	392

Absolute value of t -statistics in parentheses. * denotes significance at the 1% level.

Table 5: Demand for Athlon CPU's at site C

Dependent Variable	Product Category					
	128MB PC100 Memory	128MB PC133 Memory	256MB PC100 Memory	256MB PC133 Memory	Athlon CPUs Site B	Athlon CPUs Site C
Low Quality Sales	-1.47 (0.10)	-1.41 (0.09)	-1.79 (0.36)	-2.29 (0.36)	-1.28 (0.27)	-1.14 (0.19)
Medium Quality Sales	-0.76 (0.13)	-0.49 (0.12)	-0.79 (0.29)	-1.57 (0.79)	-0.97 (0.31)	-0.86 (0.18)
High Quality Sales	-0.41 (0.13)	-0.18 (0.06)	-0.66 (0.25)	-0.64 (0.35)		0.10 (0.30)

Standard errors in parentheses.

Table 6: Effect of low quality Pricewatch rank on sales of each quality level: estimates from six datasets.

Independent Variables	Dep. variable: Dummy for choice of site A		
	Medium and high	Medium quality	High quality
$\log(1 + P_{LowRank})$	-0.38* (5.12)	-0.68* (4.57)	-0.22* (3.03)
$\log(P)$	-3.01* (3.18)	-0.87 (0.67)	-5.03* (3.53)
Number of Obs.	10845	4078	6757

The table presents estimates of logit models. The dependent variable is a dummy for whether a consumer chose to buy from site A (as opposed to site B). The first column pools the observations from all eight specific products: medium- and high-quality; 128MB and 256MB; PC100 and PC133. The second and third are separate regressions for the medium- and high-quality subsamples. Absolute value of z-statistics in parentheses. * denotes significance at the 5% level. The regressions also included category dummies and a linear time trend.

Table 7: Evidence of incomplete consumer learning: effects of product prices and low-quality price rankings on where consumers buy medium- and high-quality memory modules.

Independent Variables	Dependent variables: Quantity of 128MB PC100 memory modules of each quality level					
	Other speed instruments			Cost instruments		
	Low	Mid	Hi	Low	Mid	Hi
$\log(1 + P_{LowRank})$	-1.48* (7.8)	-0.45 (1.5)	-0.22 (0.6)	-1.89* (5.0)	-0.34 (0.2)	0.79 (1.7)
$\log(P_{Mid})$	-1.06 (0.7)	-7.78* (3.8)	-0.47 (0.1)	11.93 (1.9)	-9.07 (1.1)	-2.49 (0.5)
$\log(P_{Hi})$	3.58 (1.9)	2.84 (1.0)	-3.35 (0.9)	5.54 (1.0)	4.26 (0.5)	-5.30* (1.7)
<i>SiteB</i>	-0.28* (2.9)	-0.49* (3.2)	-0.61* (3.3)	-0.40 (1.9)	-0.48 (1.2)	-1.10* (4.3)
<i>Weekend</i>	-0.48* (7.9)	-0.97* (8.0)	-0.67* (5.1)	-0.57* (2.0)	-0.98* (6.9)	-0.75* (5.2)
$\log(LowestPrice)$	-1.85* (2.6)	0.71 (0.7)	0.80 (0.8)	-5.01* (2.5)	1.02 (0.3)	1.21 (0.7)
Number of Obs.	683	683	683	683	683	683

Absolute value of t -statistics in parentheses. * denotes significance at the 5% level.

Table 8: Instrumental variables estimates of PC100 128MB Memory Demand Model

Quality	Product Category					
	128 MB Memory		256 MB Memory		CPU/fan combos	
	PC100	PC133	PC100	PC133	Site B	Site C
Low	-0.7%	-2.4%	4.2%	2.9%	-3.4%	-4.4%
Mid	17.0%	15.6%	16.3%	19.9%		9.3%
Hi	27.2%	26.9%	24.3%	24.9%	9.5%	14.7%
Overall	7.6%	11.5%	12.6%	15.8%	2.8%	6.3%
Average in \$	\$5.06	\$10.25	\$17.63	\$22.41	\$5.62	\$12.40

The table presents revenue-weighted mean percentage markups for products sold by websites A, B, and C in each of six product categories.

Table 9: Mean percentage markup in six product classes

Rank	Price	Predicted site B quantities			Predicted profits			
		<i>QLow</i>	<i>QMid</i>	<i>QHi</i>	Base costs		<i>c</i> = 78	<i>c</i> = 68
					Site B	Site A	Site B	Site B
1	68	54.4	5.3	2.8	-5	147	-343	275
2	69	29.9	3.9	2.3	92	229	-95	259
3	70	19.5	3.1	2.1	122	246	-2	239
4	72	14.1	2.6	1.9	145	264	55	235
5	74	10.7	2.3	1.8	154	267	85	227
6	74	8.5	2.0	1.6	141	244	85	202
7	74	7.0	1.8	1.6	130	225	84	184
8	75	5.9	1.7	1.5	127	218	88	175
9	77	5.0	1.5	1.4	129	217	96	172
10	77	4.4	1.4	1.4	121	204	92	160
11	78	3.9	1.3	1.3	118	198	92	154
12	78	3.4	1.3	1.3	112	188	89	145

The table presents predicted sales and profits that website B would have earned if it had occupied each position on the Pricewatch list pictured in Figure 1 throughout the day on October 12, 2000.

Table 10: Predicted profits as a function of rank: October 12, 2000

Date Due

Date Due		

Lib-26-67

MIT LIBRARIES



3 9080 02618 0114

