Industrial Organization: Lessons from the Internet[1]

Glenn Ellison
MIT, NBER, and IAS

and

Sara Fisher Ellison
MIT and IAS

March 2004

[1]The paper was written while the authors were visiting the Institute for Advanced Study, the latter as the Richard B. Fisher Member. The work was also supported by NSF grant SES-0219205. We are grateful for both sources of support.

1 Introduction

Many of us have grown used to, tired of, and finally downright skeptical of claims of the transformative powers of the Internet. It was to usher in the New Economy, but we seem to mostly have the Old; it would transform retail, but Toys 'R' Us has outlasted EToys; frictionless commerce would be the norm, but plenty of friction still exists. In this paper, we will focus on an area ignored by the popular press prognosticators, but one where we think the Internet has had substantial effect.[1] We discuss how the Internet has changed the field of industrial organization (IO) and what we think some of the additional long term effects on the field might be.

We cannot cover all of the effects that Internet-inspired and -aided research has had on IO in this brief essay. We considered reducing the scope of this paper by choosing a section or two of our graduate IO syllabus and describing how it has changed in the past few years. Given that most readers of this paper have never taken a graduate IO course, we decided, instead, to structure this paper around two topics that received a lot of attention in the popular press: "frictionless commerce" and online marketplaces. These are not subfields that an IO economist would recognize, but we thought people outside the field would have an easier time recognizing them as self-contained topics. Furthermore, we thought readers would be interested to hear what has come of all the hype from 1999. In Sections 2 and 3, then, we trace through the process by which casual empiricism about the Internet, formal empirical analysis and theoretical modeling aided and inspired by the Internet have shifted the thinking of IO economists in areas related to these two topics. In each, in particular, we quickly introduce the topic; we then describe how IO economists would have thought about the topic circa 2000; we discuss what has happened both in terms of what one might read in a popular press summary and in terms of observations that have been made in research papers; and finally we discuss several branches of the IO literature that have been affected by these observations and what new lessons have been learned.

[1] After composing our rough draft, we were reminded of a 2002 *New York Times* article by Hal Varian, someone with a foot in both the academic and popular press worlds. He is, it turns out, a top-rate prognosticator, not only writing about the effect that the Internet would have on academic research, but also accurately predicting the structure of our first paragraph. See "The usual decorous waltz between prices and sales becomes a lively tango in the world of online sales," *New York Times*, December 19, 2002.

Most new developments in the world do not change the course of academic research. The *Journal of Economic Perspectives* never asked anyone to write a paper on how the invention of high-definition television or the Patriots' Superbowl victory has affected the field of industrial organization. Why would the Internet have had important effects on the topics that IO economists study and the ways in which they study them? One reason is that whenever public interest is so focused on a particular development (and that interest is backed by millions in venture capital), many good economists are drawn to studying it and questions it inspires. Such a flurry of research effort is bound to broaden our understanding both of the particular development, the Internet, as well as of more general economics topics related to the development, search theory, say. In addition, the Internet has provided some unique tools—natural experiments, the ability to carry out field experiments, new data sources—that economists have exploited to great effect. Before starting on our main task, we thought it would be useful to discuss in more detail some features of the Internet that have made it important to the field of IO.

First, the Internet has provided researchers with the opportunity to study how markets function in novel and extreme circumstances. For example, economists have long known that markets are affected by consumer search costs, and one can gain some insight into the effects by comparing markets with higher and lower search costs. It is difficult, however, to find examples of two very similar markets where just search costs vary significantly. There are also some striking theoretical insights—like Diamond's (1971) argument that there can be a great discontinuity between markets with small incremental search costs and those with no incremental search costs—that one cannot test this way. With the growth of the Internet, we suddenly did have markets with essentially no search costs. Price-search websites like Dealtime, Pricewatch, and Shopper allow customers wanting to purchase a particular book, a particular name-brand digital camera, or even a flower arrangement to get a nicely sorted and annotated list of price quotes from dozens or hundreds of retailers. They provide an interesting environment from an academic perspective and otherwise, and we would urge anyone not familiar with them to take a look. Observing what happens in an enviroment where one parameter (here the incremental search cost) takes on an extreme

value (here zero) can be very informative.

Second, as suggested by the previous example, the Internet has provided a number of interesting natural experiments for researchers to exploit. For example, in traditional retail markets "identical" products offered by different retailers were always differentiated by location, customer service, and idiosyncratic consumer preferences for a given store. With the development of e-retail, products suddenly became differentiated by just the last two of those. Comparing traditional and online retailers can thus help us think about the importance of location as a factor differentiating traditional retail stores.

Third, the Internet has made it relatively easy for researchers to conduct field experiments. The earliest example we know of is work by Lucking-Reiley (1999) who examined how differences in auction rules affected the revenue raised by auctioning of hundreds of items (in his case trading cards) over the Internet. Such experiments have the advantage of allowing researchers to examine the behavior of real and interested market participants rather than college undergraduates given artificial incentives. Furthermore, Reiley's experiments were also very inexpensive to conduct. The subsequent growth of eBay has made auction experiments even easier to carry out.[2] To our knowledge, economists using the Internet to run field experiments have so far limited themselves to the study of auctions. We imagine, however, that in the future, energetic IO economists may build their careers by running experiments in different settings. For example, one could set up and operate a number of "competing" websites selling similar products. Existing retailers already take advantage of the mutability of websites to experiment with different prices, product mixes, or look-and-feels. Such experimentation could be a powerful tool for estimating characteristics of demand and testing classical and behavioral models of consumer decision-making.

Finally, the Internet has had a large effect on IO through the availability of new and rich data sources. Some of these data can be easily obtained simply by downloading webpages. For example, we and others have collected high-frequency data on the prices offered by sets of competing firms by writing programs to regularly carry out and record the results

[2]Some interesting examples are Lucking-Reiley and Katkar (2000), Hossain and Morgan (2003), and Bajari and Hortacsu (2003).

of price searches.[3] Chevalier and Goolsbee (2003) obtained price and quantity data by downloading a pages from the Barnes&Noble and Amazon websites (which include "sales ranks") at several points in time. Other data sources are private ones that exist because the growth of the Internet and other improvements in information technology have led firms to collect large databases to sell or to maintain extensive databases of their own customers. For example, Scott Morton, Silva-Risso, and Zettelmeyer (2001, 2002, 2003) have customer-level data on every car purchased from a large sample of auto dealers in a given year that includes, among many other things, the customer's name and address, the make and model of car purchased, the price paid, and whether the consumer obtained a referral from Autobytel.com in the course of his or her search. Our work on the effects of sales taxes was made possible by the fact that a dataset we obtained from a small retailer included the exact time at which each offer was received and the state in which each consumer resides.[4] An additional factor that aids data collection is that immediacy of competition on the Internet has increased the possibilities for using time-series data. Before the Internet, there were some markets where prices were posted and might change daily, such as gasoline.[5] Typically, however, prices do not change frequently, and empirical IO papers have therefore tended to rely on cross-sectional variation, *e.g.* looking at products sold at different prices in different cities. In some online markets, competitors can check prices hourly, or in real time. Hence, one may be able to estimate demand elasticities or study pricing dynamics with only a few weeks or months of time-series data, rather than several years of data.

[3]See Ellison and Ellison (2004), Baye, Morgan and Scholten (2004), and Baye, Gatti, Kattuman, and Morgan (2004) among others.

[4]See Ellison and Ellison (2003). Sweeting (2003) is also a good example of a paper exploiting a very rich dataset—in this case minute-by-minute logs of every song played by thousands of radio stations—collected so that the firm could sell the data over the Internet.

[5]Noel (2002) is an interesting study of retail gasoline prices, where the prices, posted on signs at the stations, were collected on a daily commute.

2 Competing Marketplaces

Ebay was one of the first great e-commerce successes. By the summer of 1998 it had 80 percent of the online auction market.[6] Its stock more than doubled on the day of its IPO that fall. By the spring of 1999, as it became clear that it was beating back challenges from Amazon and Yahoo, eBay's market cap went up another tenfold to over $25 billion. By the time we attended an e-commerce conference at Stanford early in 2000, much of the Internet bubble had burst. EBay, however, was still going strong, and its example contributed to making online marketplaces the Next Big Thing. Monster, HotJobs, and others were fighting to be the eBay of the online job market. FreeMarkets, which helped companies procure $2.7 billion worth of goods in dozens of categories in 1999, was the leading B-to-B exchange. Literally thousands of companies were or were planning to challenge it. Some aspired to dominate all of B-to-B e-commerce. Others like Covisint (established by the big three automakers), AlmondEx, and PaperExchange aimed just to dominate a market niche.[7] Freemarkets and the two leading firms selling software to B-to-B exchanges, Ariba and Commerce One, were themselves $10 billion dollar companies. In this section, we will discuss what IO economists might have expected, what was seen that was surprising, and how this has spurred the development of new models and reinterpretations of old ones.

2.1 What did IO economists expect?

Economists and others have long noted that trading activity can be strikingly concentrated. In Medieval Europe people would travel hundreds of miles to trade fairs to buy a year's provisions.[8] When Daniel Defoe visited the annual Sturbridge Fair in 1723 (over 500 years after it started), he was amazed to find a half mile square field full of "goldsmiths, toyshops, braziers, turners, milliners, haberdashers, hatters, mercers, drapers, pewterers, china–warehouses, and in a word all the trades that can be named in London;" as well as vast wholesale trade in woolens, hops, *etc.*, and wondered "... why this fair should be thus, of all other places in England, the centre of that trade; and so great a quantity of so bulky

[6] See Lucking-Reiley (1999).
[7] See Lucking-Reiley and Spulber (2001).
[8] See Walford (1883).

a commodity be carried thither so far:"[9] In the modern world, New York's Fulton fish market and Tokyo's Tsukiji are similar phenomena. Stock exchanges are another well-known example of the agglomeration of trade.

Explanations for concentration of trading activity usually involve some source of "increasing returns" or "network externality". In the trade-fair example, one source is transactions costs: having each trader travel once to the trade fair is much easier than having each make separate trips to visit each of the others. Another common source is a "preference for variety". For example, one might assume that the chances of finding a good match at a dating site is increasing in the number of other people using the site. Another is risk sharing. For example, a large stock market makes it easier for investors to diversify and reduce idiosyncratic risk.[10]

At the time of the boom there would not have been unanimity among IO economists as to whether online marketplaces would become highly concentrated and whether owners of online markets would be highly profitable. We see a lot of different outcomes in the physical world. Trading in stocks is typically quite concentrated, whereas bond trading does not go through a central exchange. Furthermore, concentration is not sufficient to ensure large profits for a market maker: the (nonprofit) NYSE is able to earn a great deal of money, but the Fulton fish market does not earn large rents. Concentration may also not be necessary: the U.S. real estate industry seems to be able to extract a great deal of surplus even though it is not very concentrated. What would happen in online marketplaces would depend to some extent on the strength of increasing returns, but there were also clearly relevant pricing and contracting issues that had not been thoroughly explored.[11]

2.2 What happened?

In online auction markets, the early leaders have grown and prospered. EBay's gross merchandise sales somehow seem to be able to grow by four or five percent every month and are now ten times as great as in 1999. Annual revenues of over $2 billion support a

[9] Defoe (1724). See pages 102-104 of the 1986 Penguin edition.

[10] See Pagano (1989) and Ellison and Fudenberg (2003).

[11] Spiegler (2000) had provided an intriguing analysis of how an intermediary, even without market power, might use be able to extract a great deal of surplus obtained from a match.

market capitalization of $43 billion. Yahoo! Auctions, on the other hand, has less than one-tenth as many listings and few of the listed products sell. Amazon is much further behind. The battle for Japan provides a striking contrast that suggests a large first mover advantage: Yahoo! entered before eBay, and eBay was never able to garner a significant market share before abandoning the market in February of 2002.

B-to-B marketplaces, however, have evolved quite differently. Even though some exchanges have reportedly achieved large trade volumes, revenues and profitability have not followed. For instance, in 2001 Covisint reported conducting as great a dollar volume of auctions just for Daimler-Chrysler as eBay conducted in total, but subsequently sold its auction business to FreeMarkets for a mere $16 million.[12] FreeMarkets itself recently agreed to be acquired by Ariba for what will likely turn out to be about $350 million. ChemConnect reported an $8.8 billion volume in 2002, but has since downsized to 60 employees, moved from San Francisco to Houston, and aspires to earn a profit in 2004.[13] Ariba and Commerce One were recently trading at about 98% and 99.8% below their peak values.

Online job search provides yet another case study. These sites quickly became quite popular. By August of 2000 one-quarter of all unemployed jobseekers (and over 10% of those with jobs) reported regularly looking for work on online.[14] No single site is dominant: Monster remains the largest site, but CareerBuilder is gaining on it and HotJobs is also very active.[15] The sites have at times been profitable, but not wildly so. Monster's market capitalization is now about $2.6 billion. Yahoo acquired Hot Jobs for $436 million in December of 2002.

2.3 Lessons: sources of increasing returns

Over the last decade the IO literature has seen a surge in papers on network externalities. Many of these papers start by assuming that the payoff to taking an action (*e.g.,* buying from firm X) is assumed to be some action-specific constant plus a constant times the

[12]The remainder of Covisint was sold to Compuware for an undisclosed amount in March of 2004.

[13]See "Top Online Chemical Exchange is an Unlikely Success Story," *Wall Street Journal* January 8, 2004.

[14]See Kuhn and Skuterud (2002)

[15]Both Monster and CareerBuilder (which also operates Headhunder) were reported by comScore Media Metrix to be among the top 25 Internet companies in January 2004 with 17.9 and 15.6 million unique visitors, respectively. HotJobs is a division of Yahoo! and its visitors are not separately reported.

number of other players taking the same action, with little or no discussion of whether this functional form matters or is sensible. One lesson that IO will take away from the experiences of online markets is that it is important to think more about the sources of network externalities and increasing returns and what functional forms are appropriate.

An important source of increasing returns in the Sturbridge fair example is the savings on travel costs. B-to-B exchanges should also have such an increasing return: it is easier to post an ad stating one's desire to buy polypropylene on ChemConnect than to call twenty potential suppliers on the telephone. These economies are not so large in dollar terms, however, and could not support a marketplace collecting anything like the share of revenues that eBay earns.

Another source of increasing returns in consumer and B-to-B auctions is that having a single market more efficiently allocates goods to high value users. For example, if multiple containers of polypropylene were auctioned off simultaneously in each of two auction sites (and bidders could not attend both sites), then the price for containers should be the same at each site, but will likely differ across sites. Some successful buyers at the low-priced site may have lower values than unsuccessful bidders at the high price site. Ellison, Fudenberg, and Mobius (2003) derive formulas for the strength of this increasing return. They note that it can lead markets to tip toward a dominant firm, but that it implies that increasing returns are weak once sites are large, and hence that multiple marketplaces can coexist if this is the only source of increasing returns.

Another source of increasing returns is preferences for product variety. All traders will prefer a larger marketplace if sellers benefit from receiving higher prices and buyers benefit from obtaining goods that more perfectly match their preferences. A key lesson from observing the Internet may be that this source of increasing returns deserves even more attention. EBay's continued dominance suggests that benefits from product variety may provide important benefits over a wide range of market sizes. For example, in the early days, eBay's size advantage may have made it more likely that someone looking for a 1943 steel penny could find one. Later it may have allowed one to find such a penny in as good or bad condition as one wanted and to have them available in lots of 10, 20, or 50

as well as singly. Later, still, it may have ensured that there was an auction for a coin in the desired condition ending very soon rather than in a few days. Studies of online book sales have provided striking evidence of the importance of product variety there: about 40% of Amazon's sales appear to be books too obscure to find even in a Barnes & Noble superstore, and the ability to buy these books from Amazon may have contributed about $1 billion in consumer surplus in 2000.[16] The more limited success of online job sites and B-to-B exchanges suggests that product variety may be less important in these markets. Allowing a firm to look over 100 programmers' resumes instead of ten may not make it much more likely that the firm will actually hire a better candidate.[17]

2.4 Lessons: two-sided markets

A desire to understand the experience of online marketplaces has also spurred the most rapidly developing literature in theoretical IO: the analysis of two-sided markets.[18] This literature starts with the observation that many marketplaces are also unlike the standard network externality model in that the population is naturally divided into two groups that benefit from cross-group interactions: job sites match workers with firms; dating sites match men with women; Covisint matched auto parts manufacturers with auto manufacturers. Such network effects are not unique to online marketplaces or even to marketplaces at all. The makers of video game consoles can be thought of as matchmakers who match game players with game developers; newspapers match readers and advertisers; and so forth. The boom in research into two-sided markets appears, in fact, to be due to a confluence of current events questions with antitrust scrutiny of Microsoft and Visa playing as important a role as online marketplaces.

At this point, some readers may be a bit puzzled over what we would claim is simply

[16] See Brynjolfsson, Hu, and Smith (2003). Chevalier and Goolsbee (2003) had earlier noted that the distribution of book sales appears to be well fit by a Pareto distribution with parameter 1.2.

[17] Hadass's (2004) study of one large firm reports that employees hired via online sites were significantly more likely to leave the firm in the first few years than workers hired via employee referrals and interprets this as evidence that online job sites do not provide better matches and are used mostly for economies of the "travel cost" variety. An alternate explanation, of course, is that the online employees might have been better matches and performed better in their time on the job, but left more quickly because they were more adept at searching for even better matches.

[18] Some of the important papers in this literature are Spiegler (2000), Caillaud and Julien (2003), Rochet and Tirole (2003), and Armstrong (2004)

a suboptimal name for this literature. All markets have two sides, after all, buyers and sellers. The markets studied in this literature actually have three sides: buyers, sellers, and an intermediary that is trying to profit from facilitating their interaction. In Rochet and Tirole's (2004) view, the defining feature of two-sided markets is that there are different ways of breaking up prices across buyers and sellers, and how prices are set is not neutral. For example, dating sites may make membership free for women and costly for men, newspapers sell papers to consumers at less than the marginal production cost and make money on advertisers, and eBay devotes a part of its revenue stream to providing services to large sellers.

The two-sided markets literature has discussed a number of market characteristics that play an important role in determining whether intermediaries can extract a significant amount of the surplus created when buyers and sellers interact. In doing so, it has also addressed the question of whether a dominant intermediary will tend to emerge. One characteristic is the range of pricing options available to the intermediaries: are they only able to charge simple prices, or can they credibly commit to subsidizing one side of the transaction or use contracts with payments contingent on subsequent participation and transaction levels? Another characteristic is whether participants may "multihome," or if they are constrained to interact through a single intermediary. Another is the extent to which the intermediary can monitor or regulate transactions between buyers and sellers. Basic pricing principles developed in this literature have added to our understanding of which side of the market intermediaries are likely to want to subsidize and what welfare implications various policies might have.

3 Frictionless commerce

At the time of the Internet boom, the popular press was full of stories about how the Internet would usher in a new era of frictionless commerce. Although the precise meaning of "frictionless commerce" was never entirely clear to us (as is often the case with phrases culled from the popular press), we believe journalists and others were envisioning a world where consumers were perfectly informed, retailers were engaged in intense competition,

pricing at marginal cost, and the Law of One Price ruled. Most of these stories were ridiculous, and IO economists were not surprised that they turned out to be wrong. In particular, the ideal of perfect competition with prices at marginal cost makes no sense in the real world: firms have fixed costs and go bankrupt if they price at marginal cost. In this section we will again discuss what IO economists might have expected, what was seen that was surprising, and how this has spurred the development of new models and reinterpretations of old ones.

3.1 What did IO economists expect?

Any models of a particular industry need to account for why prices stay above marginal cost. IO economists have a number of explanations for why firms are able to avoid this intense competition, including the existence of switching costs or search costs. By far the most pervasive explanation, though, is product differentiation. Firms avoid intense competition if their products appeal more to some niche of consumers than their rivals' products do, for idiosyncratic taste reasons, or simply reasons of geography. IO economists, therefore, might have thought that by eliminating one dimension of differentiation, the geographical dimension, that competition among a *fixed* number of firms should be more intense on the Internet, leading to lower margins. Likewise, to the extent that consumer search costs help firms avoid the most fierce competition, IO economists might have predicted that Internet technologies that facilitate price search could also lead to more intense competition online. IO economists would also have recognized, though, that markups are determined not only by these factors which soften competition but also by the number of firms in the market, a variable which is also determined by fixed costs. Here, they would likely have assumed that fixed costs would be lower for Internet firms. Wal-mart, for example, built 276 stores before it reached $1 billion in sales, whereas Amazon needed just six warehouses to service over $3 billion in North American sales in 2003.[19]

Little differentiation. Low search costs. Low fixed costs. In such industries one might expect to see several firms engaged in intense price competition. Perhaps online industries

[19]Wal-mart claims to have been the first company to have reached $1 billion in sales in such a short time when it did so in its 17th full year of operation. Amazon reached the $1 billion milestone in its fourth.

would look like the grocery business: an industry with over $200 billion in revenues but no firm worth as much as Starbucks. Or perhaps they might look like discount retail, where Wal-mart's lower costs have made it hugely profitable while Woolworth, K-mart, and others struggle to survive.

IO economists would not have put much faith in the "Law of One Price". Empirical studies have consistently reported finding price dispersion, even for such homogeneous products as prescription drugs sold at different pharmacies in the same small town, and models that predict price dispersion have been around for some time.[20]

3.2 What happened?

The earliest academic literature on the question of whether the Internet would provide a frictionless form of commerce produced two widely heralded findings: there is price dispersion online and online prices are not much lower (if at all) than offline prices.[21] At this level of detail, we would argue that these findings are not very important to IO. IO economists would have expected to find some price dispersion. Also, comparisons of online and offline prices do not tell us much about price-cost margins and the relative intensity of online and offline competition becuase marginal costs are quite different in the two cases.[22]

At the same time, we feel that some of the details in these papers and subsequent ones are important and surprising and will therefore have a greater effect on the way IO economists view competition. Among these findings are:

[20] Sorensen (2000) reports that the range in prices for a prescription drug across the eight to ten pharmacies in a town is about 50% of the drug price and the standard deviation is on average 22% of the mean. For theoretical models see Varian (1980), Burdett and Judd (1983), and Stahl (1989) among others.

[21] Among the most noteworthy studies of price dispersion are Brynjolfsson and Smith (2000), which collected data on the prices for 20 books and 20 CDs at eight Internet retailers (and eight conventional retailers) over fifteen months in 1998–1999, Clay, Krisnan, and Wolff (2001), which collected data on a larger sample of books from 32 Internet bookstores over six months in 1999–2000, and Baye, Morgan, and Scholten (2004a), which collected data on 1000 computer and electronic products for eight months in 2000–2001 by repeatedly asking Shopper.com to carry out price searches. The earliest studies of online versus offline prices we know of are Lee (1997) and Bailey (1998), which reported that prices for used cars, books, CDs, and software, were on average higher online than offline. Brynjolfsson and Smith (2000) looked at a more representative sample of traditional retailers and found that book and CD prices were 9% to 16% percent less online.

[22] Catalog retailers have for decades had higher costs than large discount stores for a simple reason: packaging and mailing items individually is less efficient than transporting them in truckload shipments and unloading them in bulk onto store shelves. One should have expected that e-retailers would have higher marginal costs for similar reasons.

• The magnitude of the price dispersion in the markets with "branded" websites is substantial. The reported dispersion in book and CD prices across e-retailers, for example, is typically about 10%, and the range is typically between 25% and 40%. This is only a little lower than what was reported for prescription drugs at traditional pharmacies, which does not fit well with a view of the Internet as a place where frictions have been greatly reduced.

• Price dispersion has been found even in environments that one would expect to be intensely competitive, such as markets where consumers find retailers via price search engines. For example, in our data on small computer-parts retailers selling through Pricewatch.com, the twelfth lowest price is typically about 10% above the lowest price.[23] In addition, these markets seem to lack the features around which most current models of equilibrium price dispersion are built.[24]

• The intertemporal properties of the price dispersion on the Internet seem incompatible with the traditional search-based explanations. In the standard models, dispersion usually takes the form of firms using mixed strategies. In these mixed equilibria, the firms would all want to change their prices as soon as they saw their rivals' prices. While a few studies have noted that Internet retailers can and do change their prices frequently, none find anything like the cascade of price changes one would expect if firms were continually monitoring their rivals' prices and reoptimizing.[25]

• Price-cost margins on the Internet are not extremely low. Although some well-known firms like Pets.com set prices below cost and went bankrupt, many other firms are getting

[23] Baye, Morgan and Scholten's (2004a) data from Shopper.com is comparable. One noteworthy aspect is the heterogeneity across products. The gap between the lowest and second-lowest price is less than one percent for almost half of the products, but has a mean of about 5%.

[24] Most models of equilibrium price dispersion rely on heterogeneously informed consumers, clearly not a salient feature of price search engine environments. Incidentally, price dispersion has also been found in experimental Bertrand markets, which by design have no consumer heterogeneity. See Baye and Morgan (2004).

[25] Although Baye, Morgan, and Scholten (2004b) emphasize the turnover of the low-priced firms in their data, the data are at a monthly frequency, and even at this frequency there appears to be significant inertia. Our high-frequency data on memory modules shows that there is a great deal of movement even from hour to hour: about three of the top twenty-four firms change their prices in a typical hour. Part of this is due to common reactions to wholesale price changes. The identity of the top firm changes from hour-to-hour about 4% of the time and from day-to-day about 43% of the time. This is a tremendous amount of turnover, but not nearly as much as a mixed strategy model might suggest.

reasonable margins. Amazon's accounting statements, for example, report that average percentage markups (over the "cost of sales" plus "fulfillment costs") are about 15%. In our study of one of the firms operating in the intensely competitive Pricewatch.com environment, we found that the firm was able to sustain average markups over marginal cost of about 10%.

These observations challenge the standard applications of IO theory on search costs and product differentiation. One can react either by developing new theories or by applying existing theories differently. In the section that follows we will describe developments of both of these types.

3.3 Lessons: product differentiation

The first explanation IO economists typically offer for why prices stay above marginal cost is product differentiation. Hotelling's competition-on-a-line model is the canonical example, used both for teaching students and in many current theoretical papers. Although one always emphasizes that the "transportation costs" in the model could represent either the physical cost of travelling to a store or distances in some abstract taste space, the geographic metaphor is compelling. It naturally makes one think that spatial differences are important. We assume if most IO economists had been asked how they thought Barnes & Noble and Borders avoided marginal cost pricing in the offline world, they would have put geographical differentiation at the top of the list.

Observations of the online book industry have changed this view. We assume most IO economists would continue to believe in the product differentiation model as a useful paradigm, but shift their interpretation to regard geographical differentiation as less important and idiosyncratic taste differences as more important. Chevalier and Goolsbee's (2003) study of Amazon and Barnes&Noble provides the most convincing evidence. It cleverly uses "Sales Ranks" the sites post to infer demand. They report that the demand for books at Amazon is highly inelastic: a one percent increase in the price of a book on Amazon is estimated to reduce demand at Amazon by only one-half of one percent. Most of the lost customers would go to Barnes&Noble. Barnes&Noble's own price elasticity is

much larger, about -4, and most of the customers it would lose from a price increase are estimated not to go to Amazon. These observations are consistent with a model where online bookstores are differentiated in taste space, with Amazon at one end, a group of other retailers at the opposite end, and Barnes&Noble being between but closer to the other retailers. Although demand at Barnes&Noble is much more elastic than demand at Amazon, the demand for books from Barnes&Noble still suggests that there is a substantial amount of product differentiation. The standard single-good markup formula implies that Barnes&Noble.com might support a markup of 25%. This is comparable to the gross margin that Barnes & Noble reports for its traditional bookstores in its accounting statements and bolsters the view that geographic differentiation may not have been a very important part of what differentiated traditional retail stores.

Additional evidence of the importance of brand preferences is provided by Brynjolffson and Smith (2001). In a study covering a wider set of retailers, they find that even people who use a price search engine to compare book offers seem to be willing to pay between $1.50 and $2 more to buy from one of the well-known Internet retailers than from an unknown retailer. Within the well-known retailer group, consumers will pay $1 more to buy from Amazon.com than from Borders.com or Barnes&Noble.com.

3.4 Lessons: search costs

Gathering price and product information is a costly activity, and real-world consumers clearly do only a limited amount of information gathering. It seems intuitive that limited consumer information should contribute to positive markups and price dispersion. Search models have, nonetheless, always been a bit marginalized in IO. Tirole's (1988) textbook does not even include them. When they are discussed, it tends to be as part of an isolated discussion of price dispersion, rather than as a central part of understanding why firms receive positive markups.

One reason for the inattention to search models is that people tend to use models that are simple and have predictions they like. The Cournot model, for example, seems wholly inadequate as a description of price-setting in most industries, and yet for years it was the

most popular oligopoly model. This popularity can be traced to its tractability as well as its appealing prediction that prices fall smoothly from the monopoly price to marginal cost as the number of firms increases. Search models, in contrast, tend to be complicated to work out, and many early papers had awkward solution concepts, comparative statics that seemed wrong, or both.

More recent search models have more compelling predictions. In Stahl's (1989) model, for example, equilibrium prices decrease toward marginal cost both when the cost of obtaining a price quote decreases and when the fraction of consumers who can search costlessly increases. One way in which the Internet immediately affected IO is that search models became more prominent in graduate IO classes, ours included. Evidence from the Internet also challenged these models, however, because we did not see the tremendous decrease in prices and price dispersion that many predicted.

One potential way to respond is to reaffirm one's belief in the models and question whether the Internet has reduced search costs much. Our study of firms selling computer parts via Pricewatch.com suggested this possibility (Ellison and Ellison, 2004). Although the Internet is a powerful information-gathering tool for consumers—witness the popularity of sites like Google, AskJeeves, Yahoo, and Pricewatch—we noticed that retailers are also harnessing the power of the Internet to carry out what we call obfuscation strategies. For example, the Internet makes it easy for e-retailers to offer complicated menus of prices with different options for shipping, etc., to make price offers that search engines will misinterpret (like products bundled with unattractive terms of trade), to personalize prices, and to make the process of examining an offer sufficiently time-consuming so that customers will not want to do it many times.[26] Whether the Internet will prove to aid search or obfuscation more is not *a priori* clear.

A second potential way to respond is to develop new search theories. If consumers can only find retailers via search engines (and search engines have market power), it is easy to explain why search engines might not reduce prices: search engines will extract profits

[26]The variety of these examples suggests that obfuscation could be modeled in a number of ways: with some cognitive model; as an increase in per-item search costs; as changing the game to one where firms offer menus of prices; and so on. Gabaix and Laibson (2003) treats obfuscation simply as an action that increases product differentiation in the standard model.

by charging click-through fees that lead retailers to raise their prices. Baye and Morgan (2001) developed the first complete model of a search engine capable of charging consumers and firms, and note that one can account for both positive margins and price dispersion by assuming that retailers have some captive consumers and cannot price discriminate. Subsequent papers have explored alternate ways to achive the same goals. The possiblity that we find most intriguing is Baye and Morgan's (2004) observation (made in a simpler Bertrand context) that current models may have ϵ-equilibria (equilibria where players are satisfied as long as they achieve with ϵ of the best possible payoff) that look quite different from exact Nash equilibria for small values of ϵ.

3.5 Lessons: multiproduct competition and switching costs

Like search costs, multiproduct competition is probably underemphasized as a factor which keeps prices away from marginal cost. In some sense, almost every firm is involved in multiproduct competition: it hopes to sell products in the future as well as today. It is not paradoxical that Amazon sets prices such that the elasticity of demand is only -0.5. If Amazon were to set higher prices today, more of its customers might incur the real or psychic costs involved in trying out another bookstore, and Amazon might, in addition, lose a large number of sales in the future. Another lesson of the Internet might be that we should keep in mind that switching costs can be an important determinant of markups.

Competition with switching costs is a particular type of multiproduct competition, but more traditional multiproduct competition might be important as well in understanding what factors are important in how firms set prices, both on the Internet and more generally. In our study of small unknown retailers selling through Pricewatch.com, we noted that own-price elasticities for listed products ranged from -25 to -40. Apparently, in this environment the Internet has reduced differentiation and search frictions to the point where single-product models would make us think firms could not survive: the $(p-mc)/p = 1/\epsilon$ formula indicates that markups should be 2.5% to 4%, which is not enough to cover fixed costs. To account for how firms do exist, we think that it is necessary to recognize that retailers selling through Pricewatch.com are also multiproduct firms (in the traditional sense) and

to develop new multiproduct models.

It is common for firms to list low prices for low quality items on Pricewatch and to design their websites to try to induce consumers who visit to purchase additional items or higher quality items at higher prices. The tactics remind us of ones we have seen in many other industries: hotels having expensive restaurants and charging high prices for phone calls and minibar items; rental car companies with expensive insurance and high refueling charges; appliance stores selling extended warranties, and so forth. Our more detailed empirical analysis of these practices has led us to feel that the existence of these add-ons is an important determinant of markups, both online and offline.

Immediately, a discussion of add-ons would prompt the Chicago-school critique of many supposedly anticompetitive practices: even if hotels are able to earn high profits on their restaurants, will they not compete the profits away in the form of lower room rates? In fact, Lal and Matutes (1994) provided a formal demonstration that add-ons can be completely irrelevant in some circumstances. However, Ellison (2004) developed a model to illustrate that under other assumptions we find more reasonable, the practice of selling add-ons is not neutral and can prevent firms selling minimally differentiated products from driving profits close to to zero. The key assumption of the model is that some customers are cheapskates: they will be willing to switch between firms to exploit a small price difference and will also be hard to talk into buying add-ons. In such a population, the marginal customers attracted by a price cut will be different from the representative sample of customers each firm gets in symmetric equilibrium: the marginal customer pool will be full of cheapskates. In a single good model, this would not matter. A cheapskate's money is as good as anyone else's. In a multigood model with expensive add-ons, firms may lose money on every cheapskate they serve. Hence, price competition in a population with cheapskates is similar to competition in insurance coverage in a population with sick people. The adverse selection effect makes price-cutting less attractive and raises equilibrium prices.[27]

[27]Whether firms advertise prices for add-ons is not an endogenous choice in the base model of Ellison (2004). In the simplest extension of the model with costless advertising, firms would advertise prices for add-ons and eliminate the excess profits. Ellison (2004) discusses several mechanisms by which this result could be avoided. Gabaix and Laibson (2004) explore one of them—bounded rationality by some consumers—and show how it can make high add-on prices immune to advertising under various demand structures.

Our empirical work supports both the assumptions and the conclusions of this argument.[28] We have found that when a website has a very low price, it is able to convince a much smaller fraction of its customers to purchase upgrades. For example, looking at predicted values from our model on a typical day, we report that a website would sell 60 units but only convince 14% of customers to upgrade if it charged the lowest price on Pricewatch, whereas it would sell 6 units and convince 45% of them to upgrade if it set the twelfth lowest price. We also find that markups are not being driven down as low as the Chicago-school critique would imply. Cost data indicate that the retailer we study is able to maintain markups of about 10%.

4 Conclusion

In this paper, we have discussed how a few parts of industrial organization have been affected by lessons learned from the Internet. We think the story we have told is far from over. Data resources should continue to improve. A number of the intriguing questions that have been raised still do not have good answers and should provide ample opportunities for future theoretical work: How should we think about price dispersion in environments where firms can move in continuous time and static mixed-strategy equilibria seem unreasonable? Is Amazon's current success reflective of consumer preferences or just of switching costs? Are preferences for variety really sufficiently important so as to account for why eBay has been so much more successful than other marketmakers?

In the introduction we discussed a number of features of the Internet that have made it a rich source for obtaining economic insights as well as economic data. These features are not advantageous just for these the two topics we have covered, nor just for studying industrial organization. Hence, we imagine that many other papers like this one could be written in the years to come.

[28] See Ellison and Ellison (2004).

References

Armstrong, Mark. 2004. "Competition in Two-Sided Markets." mimeo, University College London.

Bajari, Patrick and Ali Hortacsu. 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *Rand Journal of Economics*. 34:2, 329-355.

Bailey, Joseph. 1998. "Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce." Ph.D. thesis, Massachusetts Institute of Technology.

Baye, Michael R., Rupert Gatti, Paul Kattuman and John Morgan. 2004. "Online Pricing and the Euro Changeover: Cross-Country Comparisons." mimeo, Indiana University.

Baye, Michael R. and John Morgan. 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *American Economic Review*. 91, pp. 454-474.

Baye, Michael R. and John Morgan. 2003. "Information Gatekeepers and Price Discrimination on the Internet." *Economics Letters*. 76, pp. 47-51.

Baye, Michael R. and John Morgan. 2004. "Price Dispersion in the Lab and on the Internet: Theory and Evidence." *Rand Journal of Economics*. forthcoming.

Baye, Michael R., John Morgan and Patrick Scholten. 2004a. "Price Dispersion in the Small and the Large: Evidence from an Internet Price Comparison Site." *Journal of Industrial Economics*. forthcoming.

Baye, Michael R., John Morgan and Patrick Scholten. 2004b. "Temporal Price Dispersion: Evidence from an Online Consumer Electronics Market." *Journal of Interactive Marketing*. forthcoming.

Brynjolfsson, Erik and Michael Smith. 2000. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science*. 46:4, 563-585.

Brynjolfsson, Erik and Michael Smith. 2001. "Consumer Decision-making at an Internet Shopbot: Brand Still Matters," *Journal of Industrial Economics*. 49, pp. 541-558.

Burdett, Kenneth and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion," *Econometrica*. 51, pp. 955-969.

Caillaud, Bernard and Bruno Jullien. 2003. "Chicken & Egg: Competition Among Intermediation Service Providers." *Rand Journal of Economics*. 34:2, pp. 309-328.

Chevalier, Judith and Austan Goolsbee. 2003. "Measuring Prices and Price Competition Online: Amazon vs. Barnes and Noble." *Quantitative Marketing and Economics*. 1, pp.

203-222.

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, pp. 521-540.

Defoe, Daniel. 1724. *A Tour Through the Whole Island of Great Britain*. Reprinted by Penguin: London, 1986.

Deneckere, Raymond and R. Preston McAfee. 1996. "Damaged Goods." *Journal of Economics and Management Strategy*. 5, pp. 149-174.

Diamond, Peter. 1971. "A Model of Price Adjustment." *Journal of Economic Theory*. 3, pp. 156-168.

Ellison, Glenn. 2003. "A Model of Add-on Pricing." mimeo, MIT.

Ellison, Glenn and Sara Fisher Ellison. 2003. "Tax Sensitivity and Home State Preferences in Internet Purchasing." mimeo, MIT.

Ellison, Glenn and Sara Fisher Ellison. "Search, Obfuscation, and Price Elasticities on the Internet." mimeo, MIT.

Ellison, Glenn and Drew Fudenberg. 2003. "Knife-Edge or Plateau: When do Market Models Tip?." *Quarterly Journal of Economics* 118:4, pp. 1249-1278.

Ellison, Glenn, Drew Fudenberg and Markus Möbius. 2004. "Competing Auctions." *Journal of the European Economic Association*. 2:1, 30-66.

Gabaix, Xavier and David Laibson. 2003. "Some Industrial Organization with Boundedly Rational Consumers." mimeo, Massachusetts Institute of Technology.

Gabaix, Xavier and David Laibson. 2004. "Shrouded Attributes and Information Suppression in Competitive Markets." mimeo, Massachusetts Institute of Technology.

Hadass, Yael. 2004. "The Effect of Internet Recruiting on the Matching of Workers and Employers." mimeo, Harvard University.

Hossain, Tanjim and John Morgan. 2003. "A Test of the Revennue Equivalence Theorem using Field Experiments on eBay." mimeo, University of California-Berkeley.

Kuhn, Peter and Mikal Skuterud. 2002. "Internet Job Search and Unemployment Durations." IZA Discussion Paper No. 613.

Lal, Rajiv and Carmen Matutes. 1994. "Retail Pricing and Advertising Strategies." *Journal of Business*. 67, pp. 345-370.

Lee, Ho Geun. 1998. "Do Electronic Marketplaces Lower the Price of Goods?." *Communications of the ACM*. 41:1, pp. 73-80.

Lucking-Reiley, David. 1999. "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet." *American Economic Review.* 89:5, pp.1063-1080.

Lucking-Reiley, David. 2000. "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics.* 48:3, pp. 227-252.

Lucking Reiley, David and Rama Katkar. 2000. "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokmon Field Experiment." mimeo, University of Arizona.

Lucking-Reiley, David and Daniel F. Spulber. 2001. "Business-to-Business Electronic Commerce." *Journal of Economic Perspectives.*15:1, pp. 55-68.

Noel, Michael. 2002. "Edgeworth Price Cycles: Firm Interaction in the Toronto Retail Gasoline Market." mimeo, University of California-San Diego.

Pagano, Marco. 1989. "Trading Volume and Asset Liquidity." *Quarterly Journal of Economics* 104, pp. 255-274.

Rochet, Jean-Charles and Jean Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association.* 1:4, pp. 990-1029.

Rochet, Jean-Charles and Jean Tirole. 2004. "Defining Two-Sided Markets." mimeo, Institut d'Economie Industrielle.

Scott-Morton, Fiona, Florian Zettelmeyer, and Jorge Silva Risso. 2001. "Internet Car Retailing" *Journal of Industrial Economics.* 49:4, pp. 501-519.

Scott-Morton, Fiona, Florian Zettelmeyer, and Jorge Silva Risso. 2002. "Cowboys or Cowards: Why are Internet Car Prices Lower?" mimeo, University of California-Berkeley.

Scott-Morton, Fiona, Florian Zettelmeyer, and Jorge Silva Risso. 2003. "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?" *Quantitative Marketing and Economics.* 1:1, pp.65-92.

Spiegler, Ran. 2000. "Extracting Interaction-Created Surplus." *Games and Economic Behavior.* 30, pp. 142-162.

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

Sweeting, Andrew. 2003. "Product Differentiation, Station Listenership, and Station Ownership in the Music Radio Industry." mimeo, Northwestern University.

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

Walford, Cornelius. 1883. *Fairs: Past and Present: A Chapter in the History of Commerce.* Reprinted by A. M. Kelly: New York, 1968.