Lessons About Markets from the Internet

Glenn Ellison and Sara Fisher Ellison

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 mostly to 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.

The Internet was also claimed to require a whole new economics with all new laws. While this, too, was very far from the truth—existing theories have mostly done quite well—the Internet has had a substantial effect on economic thought. In this paper, we discuss some ways in which the Internet has affected how economists think about markets.

That the Internet has affected economic thought is perhaps surprising. Some economic events, like the Industrial Revolution and the inflation of the mid-1970s and early 1980s, have had a major impact on the way economics is done. But other major events, like the development of the interstate highway system and the 1987 stock market crash, seem not to have made much of a difference. For this reason, we begin in the next section with a brief discussion of what about the Internet seems to be making it important to economic research.

We then turn to our main task and discuss two specific topics that received a lot of attention in the popular press and that have been viewed differently by

■ *Glenn Ellison is Professor of Economics and Sara Fisher Ellison is Senior Lecturer in Economics, both at the Massachusetts Institute of Technology, Cambridge, Massachusetts. The paper was written while the authors were visiting the Institute for Advanced Study, Princeton, New Jersey, during 2003–2004, the former as a Member and the latter as the Richard B. Fisher Member. Glenn Ellison is also a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts.*

economists since the Internet boom: how the Internet was to create online marketplaces and whether the Internet would usher in "frictionless commerce." In each case we describe how economists would have thought about these topics circa 1999, what actually happened and what lessons about the operation of markets have been learned.

The Internet and Economic Research

The Internet has affected economic research along a number of dimensions. Whenever a new phenomenon arises, especially one that receives so much attention in the popular press, it will attract study by economists. We think that this fact alone cannot account for the number of economists who have worked on the Internet and the impact this research has had. In this section, we discuss briefly why the Internet has helped economists develop new insights into the operation of markets.

First, the Internet has provided researchers with the opportunity to study how markets function in novel and extreme circumstances. A vivid example is that with the growth of the Internet, we suddenly 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 receive a sorted and annotated list of price quotes from dozens of retailers.

Why is this useful for research? Economists have long known that markets are affected by consumer search costs, and empirical studies have tried to compare markets with higher and lower search costs. It is difficult, however, to find examples of two very similar markets where only 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 by looking at markets with low, medium and high search costs. Observing what happens when a parameter—here the incremental search costs—takes on an extreme value can be very informative.

Second, 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 economists 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 the

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 market participants rather than, say, college undergraduates given artificial incentives. Furthermore, Reiley's experiments were inexpensive to conduct. The subsequent growth of eBay has made auction experiments even easier to carry out. Some interesting examples are Lucking-Reiley and Katkar (2000), Hossain and Morgan (2003) and Bajari and Hortacsu (2003). 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, 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 models of consumer decision making.

Finally, the Internet has allowed economists to exploit new data sources: it has led to the creation of new datasets and made new and old datasets easier to access. For example, we and others have collected high-frequency data on the prices offered by sets of competing firms by writing programs to carry out price searches and to record the results (for examples, see Ellison and Ellison, 2004; Baye, Morgan and Scholten, 2004; Baye, Gatti, Kattuman and Morgan, 2004). Chevalier and Goolsbee (2003) obtained both price and quantity data by downloading 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 firms now collect large databases to sell or 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 the website Autobytel.com. 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 (Ellison and Ellison, 2003). Sweeting (2003) exploits a rich dataset—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.

The immediacy of competition on the Internet has also 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.[1] Typically, however, prices do not change frequently, and empirical papers on pricing have

---

[1] 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.

therefore tended to rely on cross-sectional variation; for example, 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.

Competing Marketplaces

Economists typically are limited to observing the slow evolution of marketplaces that already exist. But the Internet allowed economists to watch as new online marketplaces with a variety of characteristics were created, competed against each other, and evolved or disappeared. Striking things happened in several markets, which provided a great opportunity to learn about how markets work. One of the first great e-commerce success stories was eBay. By the summer of 1998, it had 80 percent of the online auction market (Lucking-Reiley, 1999). The ease with which it retained its dominance despite challenges by Amazon and Yahoo! in late 1998 and early 1999 was striking and spurred many firms to seek out similar positions in other markets. Online job search was becoming quite popular, and Monster and HotJobs were among the leading sites. The leading business-to-business (B-to-B) exchange, FreeMarkets, helped companies procure $2.7 billion worth of goods in dozens of categories in 1999. However, literally thousands of companies were planning to challenge it. Some similarly aspired to buy and sell a broad range of business products. Others, such as Covisint (established by the big three automakers), AlmondEx and ChemConnect, aimed to dominate a market niche (Lucking-Reiley and Spulber, 2001). In this section, we will discuss what 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.

What Did Economists Expect?

Economists 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 (Walford, 1883). When Daniel Defoe (1724 [1986], pp. 102–104) 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 a commodity be carried thither so far." In the modern world, great fish markets like 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—that is, they involve assumptions about technologies or utility functions that make the total surplus per participant larger in larger markets. In the trade-fair example, one source of increasing returns is that 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 commonly discussed source of increasing returns is a "preference for variety." For example, a consumer going to a shopping area with more stores may be more likely to find a product that closely matches the consumer's tastes (Gehrig, 1998). Another is risk sharing. For example, a large stock market has more liquidity and makes it cheaper for investors to execute the trades necessary to reduce the idiosyncratic risk in their portfolios (Pagano, 1989; Ellison and Fudenberg, 2003).

At the time of the Internet boom in the late 1990s, there would not have been unanimity among economists as to whether online marketplaces would become highly concentrated, nor on 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. The (nonprofit) New York Stock Exchange is able to earn a great deal of money, but New York's Fulton fish market does not earn large rents. Ownership concentration may also not be necessary for rent extraction: the U.S. real estate industry seems to extract a great deal of surplus, even though it is not very concentrated. Thus, a sophisticated prognosticator might just have said that what would happen in online marketplaces would depend on increasing returns, network issues and pricing and contracting issues in ways that had not been thoroughly explored.[2]

What Happened?

In the consumer-oriented online auction market, two remarkable things have happened. First, the market has continued its rapid growth. As Table 1 illustrates, there has been nearly a tenfold growth in the value of items successfully auctioned on eBay since 1999. eBay's revenues have grown proportionally. Second, eBay has maintained a completely dominant position. In the summer of 1999 (shortly after Yahoo! and Amazon entered) the value of goods auctioned on Yahoo! and Amazon was estimated to be about 10 percent and 1 percent, respectively, of the value of goods auctioned on eBay (Lucking-Reiley, 1999). In 2001, a Nielsen/Netratings survey put these figures at 4 percent and 3 percent. Foreign markets provide additional evidence that the consumer-oriented auction market is prone to dominance. In Japan, Yahoo! entered before eBay, and eBay abandoned the market in February 2002 after having failed to gain

---

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

Table 1
Auction Volumes and Revenues at eBay and a B-to-B Marketplace

Year	1998	1999	2000	2001	2002	2003
eBay auction dollar volume ($ million)	700	2,800	5,400	9,300	14,900	23,800
eBay revenues ($ million)	86	225	431	749	1,214	2,165
FreeMarkets auction volume ($ million)	1,000	2,700	9,900	16,700	20,300	NA
FreeMarkets revenue ($ million)	8	21	91	159	181	151

significant market share. In Europe and Australia, eBay's success led Yahoo! to abandon its auction operations.

Business-to-business marketplaces have evolved quite differently. Some exchanges have reported large trade volumes, but revenues and profitability have not followed. For example, the early leader, FreeMarkets, reported year after year that it was conducting online auctions for a larger dollar volume of goods than eBay. Its revenues, however, were always a much smaller fraction of its auction volume and ceased growing in 2002. FreeMarkets never achieved profitability and was recently sold to Ariba for about $350 million, whereas eBay's market capitalization has recently been about $50 billion. Other exchanges had similar experiences. Covisint reported conducting as great a dollar volume of auctions just for Daimler-Chrysler as eBay conducted in total in 2001, but subsequently sold its auction business to FreeMarkets for a mere $16 million. ChemConnect reported an $8.8 billion volume in 2002, but has since downsized to 60 employees, moved from San Francisco to Houston, and aspired to earn a profit in 2004 (Angwin, 2004). Ariba and Commerce One, who were the two leading firms selling software to business-to-business exchanges in 1999, had stock prices in early 2004 that were at about 98 and 99.8 percent below their peak values, respectively.

Online job search provides another interesting case study. These sites quickly became remarkably popular. By August 2000, one-quarter of all unemployed jobseekers and over 10 percent of those with jobs reported regularly looking for work online (Kuhn and Skuterud, 2002). The market slowed after the end of the Internet boom, but apparently is now regaining steam.[3] An interesting contrast with the eBay example is that no single site is dominant. In June 2004, comScore estimated that the number of unique users who carried out job searches on the three large sites, Monster, CareerBuilder and Yahoo! HotJobs, was 7.2 million, 6.5 million and 6.0 million, respectively. The sites have at times been profitable, but not wildly so. Monster's market capitalization is now about $2.4 billion. Yahoo! acquired HotJobs for $436 million in December 2002 (when the online job market was near its postboom nadir).

[3] Revenues at the largest job-search site, Monster, fell from $534 million in 2001 to $414 million in 2002 and $424 million in 2003. More recently, comScore reported that the number of American visitors to career services sites increased by almost 50 percent between June 2003 and June 2004.

Lessons: Sources of Increasing Returns

In the industrial organization literature, increasing returns are most often introduced by assuming that there are "network externalities." This is typically formalized with an ad hoc assumption that the payoff each agent gets from taking an action—like buying from firm X—is some action-specific constant plus some other constant times the number of other players taking the same action. This modeling device can be a good one, because it is one of the simplest ways to introduce increasing returns. However, it is often used with little or no discussion of whether the functional form is sensible. One lesson we would draw from the experiences of online markets is that it is important to think more about the source of 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. Business-to-business 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 20 potential suppliers on the telephone. When the set of potential suppliers is fairly small and easy to identify, however, this increasing return will not be very large in dollar terms and could not support a marketplace collecting anything like the share of revenues that eBay earns.

Another source of increasing returns in consumer and business-to-business 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 on each of two auction sites (and bidders could not participate in both sites), then the prices for identical containers sold on the same site would likely be identical, but we would expect to see some differences across sites. Some successful buyers at the low-priced site may have lower values than unsuccessful bidders at the high-priced site. Ellison, Fudenberg and Mobius (2004) 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 also implies that increasing returns are weak once sites are large and, hence, that multiple marketplaces can coexist if this factor is the only source of increasing returns.

Yet another source of increasing returns is preferences for product variety. If buyers get products that more closely match their tastes in larger markets, then it is possible to have an equilibrium in which everyone is happy to transact in a single market (even if the market owner is taking a cut). For this to happen, all that is necessary is for prices to be slightly higher in the large market than any smaller competing markets. Sellers can then be happy to attend the large market and receive the higher prices, while buyers are willing to pay slightly higher prices in order to get goods that better match their tastes. Another lesson from Internet markets may be that this source of increasing returns deserves more attention. For example, eBay's continued dominance suggests that product variety may provide important benefits over a wide range of market sizes. In the early days, eBay's size advantage over other auction sites may have made it more likely that someone looking for a 1943 steel penny could find one on eBay. Later, its size advantage may

have allowed one to find such a penny in as good or bad condition as one wanted and to find them in lots of 10, 20 or 50, as well as singly. Today, when eBay is conducting about 30 such auctions on a typical weekday and 70 on a typical Sunday, one is ensured that there is 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 empirical evidence of the importance of product variety: about 40 percent 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 (Brynjolfsson, Hu and Smith, 2003).[4]

The more limited success of online job sites and business-to-business 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. Indeed, 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 he 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.

Lessons: Two-Sided Markets

A desire to understand the experience of online marketplaces has spurred the most rapidly developing literature in theoretical industrial organization: the analysis of two-sided markets. Some of the important papers in this literature are Spiegler (2000), Caillaud and Julien (2003), Rochet and Tirole (2003) and Armstrong (2004). This literature is suboptimally named, because the markets studied in this literature actually have three sides: buyers, sellers and an intermediary that is trying to profit from facilitating their interaction.

This literature starts with the observation that in many marketplaces, 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. Of course, 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; Microsoft can be thought of as facilitating the interaction of empirical economists with the developers of statistical software; Visa helps consumers interact with retailers. In fact, this boom in research into two-sided

---

[4] Chevalier and Goolsbee (2003) had earlier noted that the distribution of book sales appears to be well fit by a Pareto distribution.

markets was probably due as much to a confluence of current events questions, such as antitrust scrutiny of Microsoft and Visa, as it was to online marketplaces.

Benefits from cross-group interactions can be a source of increasing returns. The presence of increasing returns does not mean that there cannot be active competition between multiple marketplaces. Such active competition is easily possible if the marketplaces are differentiated and can even happen without differentiation if the increasing returns are relatively weak (Rochet and Tirole, 2003; Ellison and Fudenberg, 2003). Increasing returns can, of course, create an opportunity for a market-maker to achieve a dominant position.

One of the most basic insights of this literature is that, regardless of whether a marketplace is trying to achieve a dominant position or competing against several others, one can conceptualize the problem facing a marketplace as one of trying to attract both sides of the market simultaneously. This is different from the problem facing a firm in the standard network externality model. One way to attract buyers is with good service and low fees, but it is also possible to attract buyers despite high fees by using the revenues those fees generate to entice many sellers to the market. 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. It seems natural to presume that the particular asymmetric pricing policies these firms have come up with are neither accidental nor irrelevant.

The most obvious necessary condition for the success of any marketplace is that surplus must be generated when buyers and sellers interact. The two-sided markets literature has explored a number of additional factors that affect whether intermediaries can extract a significant portion of the surplus and whether a dominant intermediary will tend to emerge. One key characteristic is the range of pricing options available to the intermediaries: are they (like Microsoft) only able to charge simple prices, or can they (like telephone companies) charge for actual interactions, or even sign credible contingent contracts making payments dependent on subsequent participation and transaction levels? Complicated contracts obviously have the potential to extract consumer surplus more fully, but in some circumstances could also make a dominant firm much more susceptible to entry and thus greatly limit profits. For instance, a potential intermediary could attract all buyers by promising to make large payments to them if it fails also to attract all sellers away from the incumbent intermediary. Another important factor is whether participants may "multihome" or if they are constrained to interact through a single intermediary. This issue can be important, for example, because when more buyers are

multihoming, sellers are more willing to attend one market exclusively to take advantage of a reduction in fees by the intermediary.

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.

Summary

Increasing returns make it possible for a marketplace to extract rents, but the increasing returns may not be strong enough to prevent a situation of multiple active marketplaces each earning fairly low rents. Preferences for product variety may be the most important source of increasing returns in current markets. Current research, however, is not yet to the point where the question of why eBay has been so much more successful than other online marketplaces is well understood. Part of the answer may be that it is due to the nature of the product variety externalities, the costs of multihoming and other factors, but there are many other issues like the importance of specific investments, and the effects of buyer and seller concentration that call out for further exploration.

Frictionless Commerce

In the popular press of the late 1990s, the term "frictionless commerce" seemed to envision an economy where consumers were perfectly informed, retailers engaged in intense competition, pricing was at marginal cost, and the Law of One Price ruled. Many academic economists were not surprised that these predictions turned out to be greatly overstated or just plain wrong.

What Did Economists Expect?

In industrial organization, the idea that firms might price at marginal cost is rarely seen as reasonable. Real-world firms have fixed costs that need to be covered. Firms can theoretically cover fixed costs despite pricing at marginal cost when marginal costs are increasing, but pricing above marginal cost is usually seen as much more plausible.

It seemed likely that the Internet would lead to more intense price competition because there would be less product differentiation, lower search costs and lower fixed costs. For instance, buying and selling over the Internet essentially eliminated one dimension of product differentiation, the geographical dimension, which suggests that competition should be more intense on the Internet. Likewise, Internet technologies that facilitate price search would reduce the consumer search costs that help firms avoid the most fierce competition, which in turn could lead to more intense competition online. Moreover, fixed costs appeared likely to 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.[5] Lower fixed costs implies more firms in a market, which should also lead to tougher competition.

However, judicious economists would have been hesitant to leap from the prediction of greater price competition to statements about industry structure. Perhaps online industries would look like the grocery business: an industry with over $200 billion in revenues, but many competitors and no firm worth as much as Starbucks. Or perhaps online industries 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.

Knowledgeable economists would also not have put much faith in how the Internet would lead literally to the "Law of One Price." Empirical studies have consistently reported finding price dispersion even for homogeneous products. Pratt, Wise and Zeckhauser (1979) collected prices for 39 goods and services in disparate categories randomly selected from the Boston Yellow Pages and found that the standard deviation across firms was on average 22 percent of the mean price for the product. Similar results have subsequently been found in a variety of settings. Sorensen (2000) provides a particularly clean example in his study of prices paid by cash customers for prescription drugs at different pharmacies in the same small town. He finds that the standard deviation of prices for a prescription across pharmacies is also on average 22 percent of the mean and shows that very little of the dispersion can be attributed to different pharmacies providing different levels of service. In addition, theoretical models that predict price dispersion have been around for some time, including Varian (1980), Burdett and Judd (1983) and Stahl (1989). Thus, even those who foresaw an increase in competition would not necessarily have expected identical prices to prevail across sellers.

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. Among the most noteworthy studies of price dispersion are Brynjolfsson and Smith (2000), who collected data on the prices for 20 books and 20 CDs at eight Internet retailers (and eight conventional retailers) over 15 months in 1998–1999; Clay, Krisnan and Wolff (2001), who collected data on a larger sample of books from 32 Internet bookstores over six months in 1999–2000; and Baye, Morgan and Scholten (2004a), who collected data on 1,000 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 (1998) and Bailey (1998), which reported that

---

[5] 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.

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.

However, at this level of detail, we would argue that these findings did not really challenge fundamental belief about markets. After all, few economists would have expected to find literally zero 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, because marginal costs are quite different in the two cases. For example, catalog retailers have for decades had higher costs than large discount stores because packaging and mailing items individually is less efficient than transporting them in truckload shipments and unloading them in bulk onto store shelves. One might have expected that e-retailers would have higher marginal costs for similar reasons.

At the same time, some of the details in these papers and subsequent ones are important and surprising and will have an effect on the way economists view competition. Here are four examples of such findings.

First, the magnitude of the price dispersion in markets with "branded" websites is substantial. For example, the reported standard deviation in book and CD prices across e-retailers is typically about 10 percent of the mean price, and the difference between the highest and lowest price is typically between 25 and 40 percent of the mean. These figures are only a little lower than what Sorensen (2000) 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.

Second, 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 percent above the lowest price.[6] The price dispersion in these markets is even more puzzling because most models of price dispersion are built on heterogeneously informed consumers, which is 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 neither incomplete consumer learning nor consumer heterogeneity (Baye and Morgan, 2004).

Third, 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

---

[6] Baye, Morgan and Scholten's (2004a) data from Shopper.com is comparable. However, considerable heterogeneity in pricing exists across products. The gap between the lowest and second-lowest price is less than 1 percent for almost half of the products, but has a mean of about 5 percent.

mixed equilibria, the firms would all want to change their prices as soon as they saw their rivals' prices. Although Internet retailers can and do change their prices frequently, no studies find anything like the cascade of price changes one would expect if firms were continually monitoring their rivals' prices and reoptimizing.[7] For example, our high-frequency data on memory modules show that about three of the top 24 firms change their prices in a typical hour. The identity of the top firm changes from hour-to-hour about 4 percent of the time and from day-to-day about 43 percent of the time. This amount of price turnover is tremendous, but still not nearly as much as a model of firms constantly reacting to the price changes of other firms might suggest.

Fourth, 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 receive reasonable margins. Amazon's accounting statements, for example, report that average percentage markups (over the "cost of sales" plus "fulfillment costs") are about 15 percent. In our study of one firm 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 percent.

These observations challenge several elements of the standard economic theory on search costs and product differentiation. One can react either by developing new theories or by applying existing theories differently. In the sections that follow, we will describe developments of both of these types.

Lessons: Product Differentiation

What are the most important elements of product differentiation, and in particular, how important is geography? If economists had been asked how they thought Barnes & Noble and Borders avoided marginal cost pricing in the offline world, many of them would have put geographical differentiation at the top of the list.

Observations of the online book industry have changed this view, leading to a belief that geographical differentiation is less important than previously believed and idiosyncratic preferences for the service, atmosphere, or other attributes that make one retailer different from another are more important. Chevalier and Goolsbee's (2003) study of Amazon and Barnes&Noble online provides the most convincing evidence. Although the two sites have similar price levels and often have identical prices, they find two useful sources of price variation: a "pricing experiment" Amazon conducted in June 2001 that increased prices in many categories; and idiosyncratic variation that arises, for example, when the best-seller lists (which

---

[7] 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.

are discounted) differ across the two sites.[8] They cleverly use the "Sales Ranks" that the sites post to infer demand. They find that the demand for books at Amazon is highly inelastic: a 1 percent increase in the price of a book on Amazon is estimated to reduce demand for that book at Amazon by only 0.5 percent. Most of the lost customers go to Barnes&Noble. Barnes&Noble's own price elasticity is much larger, about -4, and most of the customers it loses 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 the extremes 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 a substantial amount of product differentiation. The standard single-good markup formula implies that Barnes&Noble.com might support a markup of 25 percent. This markup is comparable to the gross margin that Barnes&Noble reports for its traditional bookstores in its accounting statements. These findings bolster 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 a well-known Internet retailer rather 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.

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. But while it seems intuitive that limited consumer information should contribute to positive markups and price dispersion, search models have always been a bit marginalized; for example, the standard graduate industrial organization text, Tirole (1988), does not even include them. When search costs are discussed in industrial organization, it tends to be as part of a discussion of price dispersion, rather than as a central part of understanding why firms receive positive markups. In fairness, search models 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

---

[8] On a title-by-title basis prices differences averaged about 10 percent, whereas Amazon's average price was about 3 percent higher during the experiment and 4 percent lower in the other weeks.

of obtaining a price quote decreases and when the fraction of consumers who can search costlessly increases.

Because the Internet seemed to promise considerably lower search costs, search models became much more prominent in the literature. But on the other side, evidence from the Internet also challenged the existing search models, because we did not see the tremendous decrease in prices and price dispersion that many had predicted.

One way to respond is to 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 (for example, with different options for shipping), to make price offers that search engines will misinterpret (like products bundled together), 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. Obfuscation could be modeled in a number of ways: with a 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 (2004a) model obfuscation as an action that increases the differentiation parameter in the random utility model of product differentiation. Whether the Internet will prove to aid search or obfuscation more is not clear a priori.

A second potential response is to develop new theories of search. 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 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 achieve the same goals. The possibility that we find most intriguing is Baye and Morgan's (2004) observation (made in a simpler Bertrand context) that results substantially different from the traditional ones may be possible if one assumes that firms are approximate profit maximizers rather than exact profit maximizers.

Lessons: Multiproduct Competition and Switching Costs

Multiproduct competition is probably underemphasized as a factor that keeps prices away from marginal cost. In a sense, almost every firm is involved in multiproduct competition: it hopes to sell products in the future as well as today. For example, we mentioned above that Amazon's sales of a particular

book would go down by one-half of 1 percent if the price of that book were increased by 1 percent. If one thought about Amazon's decision on each book in isolation, then one would say from this that Amazon's price was too low: a firm should always price on the elastic part of the demand curve. We think, however, that the fact that the demand for a single book at Amazon is inelastic is not at all mysterious. 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 those customers might make fewer of their book purchases from Amazon in the future. In this way, switching costs can be an important determinant of markups.

Competition with switching costs and the hope of repeat sales in the future is a particular type of multiproduct competition. More traditional multiproduct competition might be important as well in understanding 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 elasticities are extremely high. In the standard single good pricing model, markups are the inverse of elasticities, so elasticities this high imply that markups can be only 2.5 to 4 percent.[9] These markups over marginal cost are so low that a firm could not be expected to cover its fixed costs and survive. To account for how firms do exist, we think that it is necessary to recognize that retailers selling through Pricewatch.com are multiproduct firms, in the traditional sense of selling multiple products at the same time, 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 are reminiscent of ones common in many other industries: hotels with low room rates but high prices for phone calls, minibar items and restaurant meals; rental car companies with a low rate on rentals but high prices for insurance and refueling; appliance stores with low prices on refrigerators or washing machines but expensive extended service warranties; and so forth.

Any discussion of how add-ons affect pricing needs to recognize the standard 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? Lal and Matutes's (1994) model of loss leaders provides a formalization of this critique: it exhibits a model in which any regulation of add-on prices would be completely irrelevant in the sense that it would affect neither the bundle of goods each consumer purchases nor the total amount each

---

[9] Our use of percentages refer to markups as a percentage of the price. The inverse elasticity pricing rule, $(P - mc)/P = -1/\varepsilon$, where $\varepsilon = \partial Q/\wedge \partial P(P/Q)$, follows immediately from a rearrangement of the standard marginal cost equals marginal revenue condition, $Q\partial P/\partial Q + P = mc$.

consumer pays. However, Ellison (2004) developed a model to illustrate that under other assumptions we find more reasonable, the practice of selling add-ons can prevent firms selling minimally differentiated products from driving profits close 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 buy few or no add-ons. In such a population, the marginal customers attracted by a price cut will be different from the representative sample of customers for each firm: the marginal customer pool will be full of cheapskates. 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.[10]

Our empirical work supports both the assumptions and the conclusions of this argument (Ellison and Ellison, 2004). For example, on the assumption side we find that when a website has a lower price, it is able to convince a smaller fraction of its customers to purchase upgrades: on days when the firm we studied offered the lowest- or second lowest-price for 256MB memory modules, it was only able to convince 37 percent of its customers to upgrade to a higher quality product, whereas it was able to convince 65 percent of its customers to upgrade on days when it had the tenth-lowest price. On the conclusions side, we find that markups are not being driven down as low as the standard single-good elasticity formula and intuition about profits on add-ons being competed away would imply. Cost data indicate that the retailer we study is able to maintain markups of about 10 percent.

Selling add-ons is a strategy that firms in many industries could adopt. It appears to provide another reason why "frictionless commerce" will not eliminate price-cost margins.

Summary

Popular press visions about frictionless commerce seem unlikely ever to be realized. Geography may be a less important source of differentiation than had been realized, with preference-based differentiation more important. Search is not just a one-sided game and firms may take an active role in ensuring that perfect consumer learning never occurs. The multiproduct nature of retail may also play a role in keeping prices above cost.

---

[10] 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 (2004b) 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.

Conclusion

In this paper, we have discussed some lessons about markets and market competition that have been learned from the Internet. In the future, data resources on these issues should continue to improve. Economists will gain additional experience with looking at natural experiments created by the Internet and by setting up experiments of their own. A number of the intriguing questions still do not have good answers and should provide ample opportunities for future theoretical work. How should economists 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 to account for why eBay has been so much more successful than other marketmakers? Will eBay continue to be so successful? Its product variety advantage could diminish as the market grows, and trust is also an important issue. The story we have told about how the Internet teaches us about markets is far from over.

References

Angwin, Julia. 2004. "Top Online Chemical Exchange is an Unlikely Success Story." *Wall Street Journal.* January 8; Available at ⟨http://webreprints.djreprints.com/907660072246.html⟩.

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, pp. 329–55.

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

Baye, Michael R. and John Morgan. 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Mar-

kets." *American Economic Review.* 91:3, pp. 454–74.

Baye, Michael R. and John Morgan. 2003. "Information Gatekeepers and Price Discrimination on the Internet." *Economics Letters.* 76:1, 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.* 35:3, forthcoming.

Baye, Michael R., John Morgan and Patrick Scholten. 2004a. "Price Dispersion in the Small and in 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.

Baye, Michael R., Rupert Gatti, Paul Kat-

tuman and John Morgan. 2004. "Online Pricing and the Euro Changeover: Cross-Country Comparisons." Mimeo, Indiana University.

Brynjolfsson, Erik and Michael Smith. 2000. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science.* 46:4, pp. 563–85.

Brynjolfsson, Erik and Michael Smith. 2001. "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters." *Journal of Industrial Economics.* 49:4, pp. 541–58.

Burdett, Kenneth and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica.* 51:4, pp. 955–69.

Caillaud, Bernard and Bruno Jullien. 2003. "Chicken & Egg: Competition Among Intermediation Service Providers." *Rand Journal of Economics.* 34:2, pp. 309–28.

Chevalier, Judith and Austan Goolsbee. 2003. "Measuring Prices and Price Competition Online: Amazon vs. Barnes and Noble." *Quantitative Marketing and Economics.* June, 1, pp. 203–22.

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.* 49:4, pp. 521–40.

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

Deneckere, Raymond and R. Preston McAfee. 1996. "Damaged Goods." *Journal of Economics and Management Strategy.* 5:2, pp. 149–74.

Diamond, Peter. 1971. "A Model of Price Adjustment." *Journal of Economic Theory.* June, 3, pp. 156–68.

Ellison, Glenn. 2004. "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. 2004. "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. 1249–278.

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

Gabaix, Xavier and David Laibson. 2004a. "Competition and Consumer Confusion."

Mimeo, Massachusetts Institute of Technology and Harvard University.

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

Gehrig, Thomas. 1998. "Competing Markets." *European Economic Review.* 42:2, pp. 277–310.

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 Revenue Equivalence Theorem using Field Experiments on eBay." Mimeo, University of California at 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.* 67:3, pp. 345–70.

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. 1063–080.

Lucking-Reiley, David. 2000. "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics.* 48:3, pp. 227–52.

Lucking Reiley, David and Rama Katkar. 2000. "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon 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 at San Diego.

Pagano, Marco. 1989. "Trading Volume and Asset Liquidity." *Quarterly Journal of Economics.* 104:2, pp. 255–74.

Pratt, John W., David A. Wise and Richard Zeckhauser. 1979. "Price Differences in Almost Competitive Markets." *Quarterly Journal of Economics.* 92:2, pp. 189–211.

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, Jorge Silva-Risso and Florian Zettelmeyer. 2001. "Internet Car Retailing." *Journal of Industrial Economics.* 49:4, pp. 501–19.

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

Scott-Morton, Fiona, Jorge Silva-Risso and Florian Zettelmeyer. 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.

Sorensen, Alan. 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *Journal of Political Economy.* 108:4, pp. 833–50.

Spiegler, Ran. 2000. "Extracting Interaction-Created Surplus." *Games and Economic Behavior.* 30:1, pp. 142–62.

Stahl, Dale O. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review.* 79:4, pp. 700–12.

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

Tirole, Jean. 1988. *The Theory of Industrial Organization.* Cambridge: MIT Press.

Varian, Hal. 1980. "A Model of Sales." *American Economic Review.* September, 70, pp. 651–59.

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