Pricing on the Internet

March 2001

Arup Daripa and Sandeep Kapur*

Birkbeck College
Gresse Street
London W1P 2LL

Abstract

It is frequently claimed that the growth of e-commerce has created a more competitive environment. It is argued that lower production costs of online retailers encourages new entry in previously concentrated sectors, and a marked reduction in search costs and switching costs increase the intensity of competition. The limited evidence that exists paints a more mixed picture. Many online markets tend to be advertising-intensive, creating a tendency towards concentration; search and price comparison are not perfect; firms can create product heterogeneity and raise switching costs to dampen price competition. Where firms have some market power, as in the market for information goods, we expect discriminatory pricing to become the norm. Apart from posted prices, the internet has extended the use of auctions, even to relatively low-value goods previously traded in thin local markets. The low-cost, relatively frictionless online auction markets have increased profits as well as economic efficiency, and may emerge as the principal pricing method for a large number of goods and services.

* Department of Economics, Birkbeck College, Gresse Street, London W1P 2LL.
email: a.daripa@bbk.ac.uk, s.kapur@bbk.ac.uk

Pricing on the Internet

I INTRODUCTION

Even though e-commerce accounts for only a small fraction of total retail transactions, it has had a noticeable impact on the nature of competition and pricing behaviour in many sectors. It is frequently asserted that the growth of e-commerce has resulted in a more competitive environment in many markets. This observation is typically based on two premises. The first premise is that the internet alters the structure of costs in many industries: it is claimed that online firms have lower set-up costs and lower marginal costs of production and distribution compared to conventional firms. This promotes entry of new, virtual firms in many sectors that were previously concentrated. The second premise is that the internet has facilitated a dramatic reduction in consumers' search costs and the their cost of switching between rival sellers. Browsing a distant store's website to check prices is easier than visiting the store. 'Shopbots' -- software that simultaneously queries many stores for price information -- enables cheap and effective price comparison, especially for homogeneous goods like books, CDs, branded consumer durables and airline tickets. The ability to purchase online makes it easy to respond to discovered price differentials: switching from one online seller to another is easier than travelling from one store to another. Thus, the argument goes, by eliminating the usual frictions in the market place, the internet has increased the intensity of competition. Together these developments should take markets closer to the theoretical model of perfect competition, or at least to the intense price competition of the Bertrand variety, with prices close to marginal costs.

Information goods are particularly well suited to migrate to online markets. These include newspapers and other information services, computer software, archival databases, downloadable music, scholarly journals, etc. – indeed, anything that can be stored and transmitted in digitised form. Information goods are costly to produce but cheap to reproduce: the marginal cost of production is close to zero for many of them. There has been a noticeable trend of falling prices, and an explosion in the amount of free information available on the internet. Online editions of many newspapers are available free of charge. Financial data that only a few years ago would have been costly to acquire can be downloaded at no charge. The venerable Encyclopaedia Britannica, once marketed as a lifelong investment, is now accessible

as a free online service. While sellers have occasionally tried to charge subscription fees for such services, they have usually abandoned the idea for “competitive reasons”. In most case, competition has driven prices down to marginal cost, namely zero.

However, for many other goods and services sold online, the picture is far from clear. As is apparent to anyone watching the so-called New Economy, the premise that virtual firms have a cost advantage over their old economy rivals is not borne out by facts. While the cost of setting up a website is relatively low, to succeed in online markets, substantial investment is required to create the appropriate technological infrastructure and delivery mechanisms. There may be substantial economies of scale in these activities. Success also requires considerable investment in establishing and maintaining a brand-presence: the cost of advertising tends to rise endogenously in these markets. As Sutton (1998) has argued, such characteristics typically make for high concentration in equilibrium. The winner-take-all aspect of the internet often encourages excessive entry in early stages, but ‘penetration pricing’ – selling goods at heavily discounted prices in order to build a customer base -- has proven to be costly and short-lived. What appears extremely competitive at the moment may be less so after the dust has settled. Steady operational losses have resulted in numerous bankruptcies and consolidation. We conjecture that in most sectors only a handful of virtual firms will survive what is, in essence, a war of attrition. As Clay et al (2000) point out, the online book-retailing sector is now much more concentrated than its conventional equivalent.

Even if only a handful of firms survive, should the reduction in search and switching costs not create a frictionless environment, and lead to a competitive Bertrand-like outcome? In some sectors such price transparency has resulted in intense price competition and substantially lower prices for consumers. For instance, Brown and Goolsbee (2000) find that online price comparison in retail insurance markets has led to a substantial reduction in prices for consumers. Also, in many sectors intense price competition in online markets has put downward pressure on prices in the conventional retail elements of that sector. Goolsbee (2000) finds some evidence of this in the retail computer market.

But, more generally, the answer to this question is not straightforward. A part of the difficulty lies in that the internet increases the information available to sellers as

well as buyers. The market outcome, and whether average prices rise or fall, may depend on the relative ability of each side to manipulate and use that information to their advantage. There are reasons to believe that the outcome may not be more competitive. First, while price-comparison services help consumers to search for the lowest price, search is far from perfect. Sellers can sometimes reduce the efficiency of price search through deliberate obfuscation of the search process. Second, the internet makes it easy for online retailers to track the behaviour of their rivals: in some circumstances this makes implicit collusion more likely. Third, the internet allows sellers to collect a remarkable wealth of information about their existing and potential customers: where retailers have some market power, this enables better price discrimination, with increased revenue for retailers, and possibly *higher* prices for some buyers.

We should also expect that retailers will find devices to lock-in their customers. By raising switching costs endogenously, retailers aim to reduce price sensitivity and to counter the effect of easier search. Loyalty schemes that offer bonuses to long-term customers may be one such device. Deliberate product differentiation of seemingly homogeneous products, often achieved by changing the product-*bundle* offered as part of the transaction, may serve a similar purpose. The ability to observe consumer preferences makes it easier to create more personalised packages, by adding other goods or services to the original product, and thus linking the package to observed characteristics of a consumer's preferences.

We examine these developments for various categories of goods and services that have migrated to online markets. The first category includes homogeneous goods such as books and recorded music, but is also relevant to airline tickets and simple financial products such as insurance and mortgages. The internet has spawned a class of online intermediaries in competition with the traditional retail intermediaries in these sectors. By saving on distribution costs – say, the cost of warehousing and retail displays -- the new intermediaries claim to have a cost advantage over their bricks-and-mortar rivals. To the extent that these goods are homogeneous goods, they are well suited to automated price comparison. We examine how the internet encourages or inhibits competition in these markets and assess the evidence.

The second category of goods that we examine is information goods. Given that these products can be digitised and transmitted at low cost over the web, they

migrated to the web quite early and are likely to be sold increasingly through online channels. Due to their very nature, these products are characterised by some degree of product differentiation, providing some market power to the sellers. The question is whether the internet will enable firms to increase the degree of product differentiation and to use this for greater price discrimination. We look at the likely forms of pricing that might emerge in these markets.

Trading arrangements in these markets are based on prices posted by sellers. The internet has also made it possible to extend auction-like trading arrangements to market transactions for relatively low-value items, by enabling market participants to communicate with each other at low cost. We look at retail auctions (mostly for collectibles) and the so-called 'reverse auctions' (where buyers specify the maximum amount they are prepared to pay, say, for an airline ticket, and sellers compete for this offer). Of course, the major contribution of internet auctions has been in the 'business-to-business' sector, where they have altered the structure of firms' supply chains and, in many cases, the market structure of commodity markets. We do not consider these sectors in this paper.[1]

This paper is organised as follows. Section II examines online markets for homogeneous goods and how they are affected by comparison shopping. Section III looks at pricing behaviour in online markets for information goods, where inherent product differentiation implies that sellers have some market power. Section IV considers online auctions. Section V concludes.

II ONLINE MARKETS FOR HOMOGENEOUS GOODS

Books and recorded music are typically published by relatively large firms and are sold through retail intermediaries. The sales channels are usually non-exclusive in that most bookstores and CD retailers sell the products of multiple publishers. Books are a homogeneous product -- they are uniquely identified by their ISBN number -- and hence uncertainty about the quality of the good is not as much a deterrent to purchasing them remotely. Not surprisingly, books and recorded music were among the first products to migrate to the web. Likewise, many branded consumers durables, or generic goods (like computer hardware and memory modules) have

[1] See Lucking-Reiley and Spulber (2001) for an overview of business-to-business auctions.

migrated to online markets. Even though the market share of online retailers remains small – Goolsbee (2000) notes that the online sales of books amount to no more than 5% of the industry total – it may grow substantially over time. Online book retailers compete with each other, and increasingly with conventional bookstores. Some retailers, like Barnes and Noble, now operate in both online and offline markets. Similarly, online markets in other relatively homogeneous goods like airline tickets, and simple financial products like insurance are likely to grow over time.

At the same time, precisely because these goods are standardised, these goods are amenable to search and price-comparison services. BargainFinder.com was one of the early examples of such a service: someone who wanted to buy a particular book or CD could use BargainFinder's software to query various online stores in real time and compare prices, enabling the buyer to find the cheapest retailer.[2] Precisely because books are identified uniquely by their ISBN numbers, it was relatively easy to devise search engines for these. Now there is bewildering range of competing price-comparison services. In the US, mySimon.com and dealtime.com are among the market leaders, but there is also evenbetter.com, bottomDollar.com, addALL.com, and numerous others. Apart from price-comparison agents that work by querying prices directly, there are other services, like Pricewatch.com, in which retailers choose to post their prices on a common database, allowing buyers to search for the lowest price. Early comparison engines were simple and compared retailers only in terms of quoted prices. Over time, they have become more sophisticated, and can rank rival retailers by their quoted prices, by their final price inclusive of packing and delivery charges, by their speed of delivery, and can also report previous customers' satisfaction ratings. The evolution of XML and other common standards for organising information on the web is expected to make such multi-attribute search and price-comparison even more reliable.

In theory, the ability to search for the lowest price, combined with relative ease of switching from one online seller to another, should create intense price competition in online markets for these homogeneous products. If online retailers also have cost advantages over conventional retailers we should, *a fortiori,* expect prices in online markets to be lower. Further, any price dispersion that exists in markets for homogeneous goods – the coexistence of different prices for the same good – is

[2] See DeLong and Froomkin (2000) for more on early price comparison services.

usually attributed to information imperfections. If search costs and switching costs are low, we should expect that firms with high prices will not survive, so that price dispersion should be lower in online markets.

Existing evidence does not support these predictions. Many early studies found that average prices of these goods in online markets were higher than the prices in conventional stores (see Bailey (1998), for instance). More recent studies suggest that some online markets may be cheaper now: Brynjolfsson and Smith (2000a) find that, for books and music CDs, online stores are cheaper than conventional stores by a margin of 9 to 16%. This trend towards eventually lower price online is usually described in terms of the increasing maturity of these markets. However, on the issue of price dispersion, the results are more surprising. Almost all studies of pricing in online markets report substantial and *persistent* price dispersion. In their study of 32 online book retailers, Clay et al (2000) found that the standard deviation of online prices for books, expressed as a percentage of their average price, varies from 17 to 28%. Brynjolfsson and Smith (2000a) found that, in their data, the difference between the lowest quoted price and the tenth lowest price averaged as much as 33% for books and 25% for CDs. And, equally surprisingly from a theoretical point of view, they find that even though books and CDs are supposedly homogeneous, the firms that have the lowest price do not have the largest market shares. Amazon.com is typically 10% more expensive than the cheapest retailer, and yet dominates the online book market (by some estimates its market share is over 60%).

There are various explanations for these findings. The simplest is that these products are not quite homogeneous but differ in overall package of bundled services that accompany the transaction. These differences may include the quality of the 'online shopping experience', the speed of delivery, store policy on returning defective or unwanted items, etc. For instance, Amazon provides personalised recommendations to its registered users, based on tracking their purchase history and through 'collaborative filtering' (detecting patterns in preferences and purchases across people, and using these to suggest titles). To the extent the buyers come to value these services, they may be willing to pay for them. In other words, even though the underlying product seems to be homogeneous, the accompanying product bundle is heterogeneous, and this heterogeneity can explain some price dispersion.

(i) Frictions in Online Markets

To explain price variation in terms of heterogeneity is convenient but is, at best, a partial explanation. It is also possible that the so-called frictionless markets are not quite so, so that price dispersion can be due to market frictions.[3] For instance, buyers' purchase decisions may be distorted by the lack of trust in some retailers; switching costs may be more significant than they are claimed to be; and price comparison may not be perfect. We look at these in turn and begin with the issue of trust.

The physical separation of buyers and sellers in online markets, and the temporal separation between paying for a good and receiving it, creates a potential problem of trust. In such situations buyers may be prepared to pay a premium for the security of buying from a reputable store. Indeed, low prices are often dismissed as too good to be true, and brand-names serve as a signal for reliability in the non-contractible aspects of product bundle. That could explain why Amazon.com has a dominant share of the online book market despite the fact that its prices are not always the keenest.

Two, there is the possibility that switching costs are not as low as they are claimed to be. In contrast to the simplicity of buying something in a conventional store, purchasing at an online store requires an individual to fill out multiple forms, create user identities and passwords – in effect to interact with a complex database. If revisiting an online store can economise on these time-consuming tasks, switching costs are indeed positive. If familiarity with a particular electronic store-front is valuable to buyers, it may result in what has been described as 'cognitive lock-in'. As Beggs and Klemperer (1992) show, switching costs tends to result in higher prices.[4] Software innovation may reduce such switching costs in the future – a range of software and services exist that can store your essential personal data on your own computer and 'port' these on demand to expedite transactions with new retailers, but so far their take-up has been low.

---

[3] For a discussion of some of these themes, see Odylzko (1996), and Brynjolfsson and Smith (2000b).

[4] In the presence of switching costs, buyers realise that firms will be able to raise prices in the future and hence may be less responsive to price cuts today. Also each firm anticipates that aggressive pricing today may provoke aggressive pricing by their rivals in later periods. Both effects weaken the incentive to cut prices, and result in higher prices on average.

In addition, we should expect online retailers to create devices that deliberately raise the costs of switching to their rivals. If so, switching costs are, to some extent, endogenously determined. Loyalty schemes are one such time-honoured device. Just as frequent flyer programmes allowed airlines to mute the intensity of price competition, we should expect online retailers of homogeneous goods to adopt such schemes. For instance, Drugstore.com offers guidance of potential adverse interactions between prescription drugs, but only to consumers who purchase *all* their drugs from them. Amazon.com rewards its loyal customers with additional services like expedited delivery. It has developed software that enables registered users to complete the entire purchase transaction with a single click. To the extent that buyers come to value these services, they serve as a deterrent to switching. And, increasingly, online stores may resort to frequent-buyer discounts: to be really effective, these would be non-linear with proportionately-higher discounts for the most loyal customers.[5]

A third explanation is that, as yet, search and price comparison is not very comprehensive or effective.[6] Johnson et al (2000) find that many online buyers do not search at all, and even among those who search, sometimes the intensity of search *declines* with user experience. For instance, prospective buyers often use search to identify a cheap travel agent for and, once found, tend to stick with that agent rather than compare prices repeatedly. This suggest that search is costly and /or not very valuable.

Further, retailers may be able to obfuscate the search process. The first versions of shopbots compared stores on the basis of list price alone. Retailers would often quote low prices to attain high rankings on price comparisons but add unreasonably high charges for shipping and handling. Search engines are better now: they can now quote the price including all handling charges, and also report customer satisfaction ratings. A second, more brutal, obfuscation tactic used by some retailers is to prevent price-comparison engines from accessing their price data. This tactic is

---

[5] One could ask how sellers hope to profit if repeat purchases can only be induced through deep discounting. A crucial purpose may be to obtain information on personal characteristics of the consumer. That enables the seller to customise products for each individual by linking specific services that consumers value – allow them to create heterogeneity -- in order to improve the effectiveness of price discrimination.

[6] For a discussion of the effect of comparison engines on shopping behaviour, see Brynjolfsson and Smith (2000b).

somewhat self-destructive – by blocking search you lose potential customers yourself -- but nonetheless this is common practice and it does reduce the overall efficiency of search. A third tactic is to create spurious product differentiation by creating slight variations in the product bundle to confuse search engines. Lastly, Ellison and Ellison (2001) find evidence of 'bait-and-switch' tactics in the market for computer memory modules. Firms offer inefficiently low quality products at a very low price to score highly on price comparisons: this enable them to attract customers, whom they then try to convince pay extra for the better quality product they really want. The prevalence of this tactic makes search engines less effective in comparing prices of better-quality memory modules. If search is imperfect for these reasons, we should expect price dispersion may persist. Over time, search engines may improve, but then so might the obfuscation tactics of online retailers.

Given the role that price-search engines are likely to play, it is important to understand their functioning more closely. Baye and Morgan (2001) view them as the new 'information gatekeepers'. To understand their role in markets, we need to appreciate the implications of search efficiency for the information gatekeepers themselves. To survive as businesses in their own right, they must generate revenues. They can do so through advertising revenue and through direct charges levied on buyers and sellers. In practice, most search engines offer their service free to buyers while charging a fee to their sellers, either in the form of commissions on referred sales, or fees for inclusion in the listings.[7] Here there is a conundrum: a price comparison service, if effective, will intensify price competition and reduce price dispersion. In the limit, if there is no price dispersion, search is useless. But, in a world in which consumers do not search, sellers are unwilling to pay for listings or referral fees. Further, no retailer will choose to advertise on a search engine that consumers do not use. Thus, a price comparison service that was very effective in terms of intensifying price competition and reducing price dispersion would undermine all its sources of revenue. To put it differently, if price comparison engines are to survive as profitable enterprises, it is important that price dispersion is not eliminated altogether.

[7] Some shopbots, such as mySimon.com offers 'priority listings' to retailers that pay premium fees. As a result their price comparison are not always unbiased.

Baye and Morgan propose one possible resolution of this conundrum.[8] Using an approach similar to Varian (1980), they see dispersion as the outcome of deliberate price randomisation by firms. They model sellers as local monopolies of a homogeneous good, who must choose whether or not to advertise their price on a gatekeeper's site to attract buyers from other localities. Their model has an equilibrium in which the gatekeeper chooses its fee structure so that all buyers subscribe to its search services, sellers randomise over the decision to advertise on the site and also randomise over the advertised price. Buyers face some cost of travelling to the local store. They use the gatekeeper's service to buy from the store which has the lowest listed price, and if the good is not listed by any seller -- given the equilibrium randomisations, there is a non-trivial possibility of this – they buy from their local seller. At the equilibrium each seller is indifferent between listing and charging any price between the Bertrand price and the monopoly price, or not listing and charging the monopoly price. As a result of the mixed-strategy adopted by the sellers, the equilibrium outcome displays price dispersion.

(ii) Implicit Collusion

Online prices may also be higher than expected due to the greater possibility of implicit collusion in online markets. This enhanced possibility may be the result of changes in the market structure. In the travel industry, for instance, the arrival of large online travel intermediaries has made it harder for small independent travel agents to survive. In the US, the number of independent travel agents has fallen by 15% since 1997, and this tendency may be exacerbated by the airlines' decision to reduce commissions on ticket sales. The major European airlines have expressed considerable interest in setting up a commonly-owned online travel agency to replace their agents. Even the airlines realise that such steps are reminiscent of early cartel associations and have issued various 'clarifications' to divert regulators' scrutiny.

In markets where economies of scale allow only a handful of online firms to survive, the nature of information flows on the internet could potentially increase the likelihood collusive pricing. In particular, the internet alters the speed with which firms can monitor and react to their rival's prices. It could be argued that in conventional markets retailers do not respond immediately to their rival's price cuts because it takes time to learn about the rival's price cut and the menu costs of changing prices

[8] For other approaches, see Ellison and Ellison (2001), and Kephart and Greenwald (1998).

prohibit frequent changes in response. In online markets, the story may be different. Menu costs are low because prices can be changed quite easily in a central database. Firms can observe and react to each other's prices quite readily. Indeed the process can be automated by using software that tracks rivals' prices and uses simple algorithms to respond to it. If so, firms can react to their rivals' price changes in minutes. They may use this ability to match their rival's prices, dropping their prices and raising their prices in tandem with their rivals. Price matching may well become the norm in some markets. While this seems like a pro-competitive development, its real effect on prices could be perverse. The essence of price competition lies in the fact that if a retailer drops its price, it expects to gain market share: it makes sense to reduce prices as long as the gain of a larger market share outweighs the loss due to lowered price. If your rivals match your price cuts instantly, the incentive to lower prices is dampened: a price reduction does not increase the market share but results in lower profits on existing, infra-marginal sales. On the whole, the widespread prevalence of price-matching could result in higher prices in the aggregate.

We must not overstate these possibilities as they depend quite crucially on the precise form of price-matching behaviour. In particular, if firms can carry out price discrimination – charge different prices to different customers for the same good, price-matching guarantees may be implemented only for buyers who are informed about rivals' lower prices. Corts (1996) points that in some cases this kind discriminatory price matching may actually lower prices in the aggregate. In the next section we look at how the internet affects the possibility of price discrimination.

III ONLINE MARKETS FOR INFORMATION GOODS

The internet is eminently suited to online distribution of many information goods – the list includes computer software, recorded music, electronic newspapers and scholarly journals. All these can be stored in digitized form and distributed at relatively low cost over the internet. Improvements in the transmission capacity – the bandwidth – should allow feature-length movies and other entertainment products to be readily downloadable in the near future.

Information goods differ from conventional goods in the structure of their costs. Producing an information good is costly, but reproduction is relatively cheap. The

cost of producing a Hollywood feature film runs into millions of dollars, but it is possible to make near-perfect copies of the first print at negligible cost – the cost need not be much more than the cost of the physical storage medium. Or, to put it another way, information goods have relatively high fixed costs of production but their marginal cost is close to zero. The problem is that if the marginal cost is indeed zero, or close to zero, pricing based on marginal cost is not in the producers' interest. A firm that sets its price at marginal cost will not be able to recover its fixed costs. So how are these goods priced, and how is their migration to the internet likely to affect pricing behaviour?

(i) Price discrimination

Consider a monopoly seller of an information good who faces a large market with heterogeneous buyers. The buyers differ in their valuation of this information good. A single posted price would not maximize the seller's revenue in such cases. A better outcome, from the seller's perspective, would be to set each buyer a price equal to his or her maximum willingness to pay for the information good.[9] This would quite naturally involve price discrimination: charging different prices to customers for the same good. Of course, price discrimination is not peculiar to e-commerce but the nature of online transactions may alter the ease with which it is carried out.

For price discrimination to be feasible, the market should satisfy some conditions. One, the firm must have some market power. Two, it should know about each consumer's willingness to pay, or at the very least, should be able to sort customers who are willing to pay more from others who are not. Three, it should be able to segregate the markets, to prevent people who can buy the good cheaply to resell the good at a higher price to others.

One time-honoured way to sort customers is to discriminate on the basis of some observable consumer attribute that is correlated to their willingness to pay (this is called third-degree price discrimination). Airlines offer discounted airfares to students on production of student IDs: here, student status is correlated with lower willingness

[9] Of course, it may be hard for the seller to ascertain the buyer's willingness to pay. Often the buyer herself cannot judge the value of information prior to receiving it. Of course, once you have acquired the information, there is no incentive to pay or to acknowledge its value. Smooth transactions in information goods require a mix of many devices to circumvent these problems. Free samples and previews aim to reassure buyers about the quality of what they are about to buy; transactions may occur within long-term relationships between the buyer and seller, or be restricted amongst reputable traders.

to pay. Publishers of college textbooks often set lower prices in developing countries: they use geography to sort and segregate their markets. How does the internet affect the ability of firms to price discriminate along these lines?

Of course, it is precisely these forms of price discrimination that begins to fray with e-commerce. It is harder to use geography as the basis for discriminatory prices. Geographic location does not matter when people order online, provided that delivery costs are not too sensitive to distance and if there are no legal and tariff restrictions on the free mobility of goods. If people can buy cars online, and if there are no restrictions to transporting them within the EU, the geographic price discrimination practised by European car manufacturers will decline.

However, there are other ways in which the internet encourages price discrimination. Online sellers may be remarkably well-informed about their potential customers and their preferences. An online retailer can identify a customer who returns to their site by using 'cookies' – these are small bits of information lodged on the user's computer that allow the online seller's computer to 'recognise' the customer. It is then possible to match the customer to his previous history of browsing and purchases at the site. And it does not end there: once you provide the seller a delivery address, they may be able to use existing third-party databases to get a better idea of the market value of the house you live in, the average income in your neighbourhood, and so on.[10] The seller can use this wealth of information to customise prices for each potential buyer. The nature of the online transactions permits this: it is easy to quote a different price to each customer who logs on, and prices can be personalised just as pageviews are. This is much harder to do in a conventional store where prices are posted publicly.

How would an online seller want to use this ability? For one, it could charge a low price to attract new customers, while extracting a higher price from loyal customers (i.e., those who are locked-in). Indeed, it has been reported that Amazon.com has attempted such price discrimination, quoting higher prices to existing customers (see BBC (2000)). The helpful advice offered by the ever-alert hacking community was to conceal your identity by blocking Amazon's cookies from your browser. Amazon's

[10] The seller can also determine the kind of browser and software you use to access their site, and the quality of your connection: a fast (expensive) connection or a slow (cheap) one. To the extent these are correlated with income, they could provide additional bases for price discrimination.

spokespersons dismissed this charge and argued that they vary prices randomly as a part of their normal experimentation, but this is not very credible. In the long run it is unreasonable to expect that online retailers will not use the vast amount of information they collect about their customers to improve their revenue.[11]

We should also expect more temporal price discrimination: price discrimination based on the timeliness of a good or service. Such price discrimination is not peculiar to online markets. A person who is eager to enjoy a new movie or novel soon after its release can be made to pay more than those who are willing to wait. The early hardcover version of a novel costs significantly more than the paperback version released months later: the price-differential does not reflect the cost of binding but is a form of price discrimination. Movies can be seen more cheaply through video-rentals if you are prepared to wait a few months. Such temporal price discrimination is not easily undermined by e-commerce (though illegitimate e-commerce, say violations of copyright laws, restrict its depth), and may even grow. For instance, online services that report financial data are available in different versions -- stock quotes in real time are expensive services bought by market traders, while time-delayed quotes are available freely. Indeed, we can view these as a special case of 'versioning', which we discuss next.

(ii) Versioning

Timeliness is just one of the qualitative attributes that information goods may differentiated across. Digital images can differ in the fineness of their resolution, while software can differ in the number of features available. With information goods, producing multiple versions is not very costly. Typically, the inferior version is produced by downgrading the best available version, and sold at a lower price. For instance, software can often be bought in 'student' and 'professional' versions; the student version is similar to the professional version with some features disabled, and is sold at a significantly lower price. This kind of pricing strategy is useful when the seller cannot distinguish between customers' willingness to pay on the basis of any observable characteristic. It works as long as consumers with a high willingness to pay are also the ones who value quality more highly. By offering multiple versions, and allowing buyers to select which one they want, the seller can sell at a high price

[11] At the same time, there is the rise of sites like Safedoor.com that help customers to anonymize their transactions.

to the quality-conscious buyers and, at the same time, capture some revenue from the rest, by selling them a cheaper, inferior version. A single version, if of the high-quality, expensive variety would lose potential revenue from the lower end of the market, while a single low-quality-version would unnecessarily compromise revenue from the upper end of the market. Interestingly, as Deneckere and McAfee (1996) pointed out, it pays to deliberately reduce the quality of the inferior version, if it persuades some customers who were on the margin of choice between the two versions, to upgrade to the expensive one. Less legroom on the economy sections in aircrafts makes business customers more inclined to upgrade.

Versioning and welfare: an example

Suppose the quality-sensitive, high-willingness-to-pay consumer values a high-quality version at £10 and the low-quality version at £4. A second consumer, with low willingness to pay, does not care for quality, and values both versions at £3. Suppose the cost of production is £1 for either version. Indeed, the inferior version can be thought of as a 'damaged' version of the superior one. If only the superior version is produced, the price can be set at £10. Only the high willingness-to-pay customer would buy, yield £9 in profit. In this case, the social surplus (i.e. profits plus consumers' surplus) is £9 too. If only the inferior version is produced, profit is maximised by setting price at £3 and selling 2 units: this generates a profit of £4, and a social surplus of £5. Suppose both versions are produced, and the inferior version is priced at £3. The low willingness-to-pay customer would buy this version at this price, but the quality-sensitive customer may be tempted to buy this too. Buying the inferior version gives him a net utility of 4-3 = £1. To persuade him to buy the superior version, that must cost no more than £9. The total profit is now (9-1) +(3-1) = £10, and the social surplus is £11. In this case, producing two versions increases profit, and also enhances social welfare.

Versions produced, price	**Quantity**	**Profit**	**Social Surplus**
Superior only, £10	1 unit	£ 9	£ 9
Inferior only, £3	2 units	£4	£ 5
Both versions with price discrimination superior at £9 inferior at £3	 1 unit 1 unit	 £ 10	 £ 11

Table 1: Profit and social surplus with versioning

In this simple example, versioning is welfare-enhancing. Varian (1985) points out that, in general, if price discrimination enlarges the market, welfare may increase. If, on the other hand it reduces the size of the market, aggregate welfare necessarily decreases. For information goods characterised by low marginal costs, versioning would allow sellers to also serve consumers with lower valuations and thus for some categories the price would be closer to the marginal cost. It is then reasonable to conjecture that versioning will typically lead to market expansion. Note that marginal cost pricing is not feasible for information goods, as it does not allow firms to cover their fixed costs. Price discrimination may be one way to ensure that low value customers are served at all – thus price discrimination might be the optimal outcome.[12]

(iii) Aggregation of Goods

Many of the things we buy through conventional channels are bundled commodities: a daily newspaper or a magazine is a bundle of many articles and reports on various subjects, such as news, sports and financial information. You may care a lot for some of these pages and not at all for others. Similarly, this issue of OXREP is a bundle of many articles on the internet. These are bundled together because in the editors' opinion the articles are linked by a common theme. But in addition, the technology of production, printing and distribution makes it economical to sell these articles as a bundle rather than individual articles. The aggregation of goods and services extends further. Instead of buying individual issues of the journal, you may choose to subscribe to the journal – thus bundling the purchase of issues over time. Institutional purchase of the journal bundles the purchase across people.

The ability to distribute information goods online creates the possibility of novel forms of pricing. The structure of costs is quite different for the online version of this journal. If articles can be downloaded directly, the production economies that make it worthwhile to bundle articles together are less important. Might it be sensible to allow

[12] Of course, successful price discrimination requires some market power. To what extent would the growth of price-comparison services erode this market power? The possibilities here are quite interesting. The technology allows retailers to distinguish between customers referred by price comparison sites and those that come directly. To the extent the former are more likely to be price sensitive, it could set a lower price for them, while charging a higher price for those that come directly.

readers to buy individual articles, each for a small fee? Abstracts might be available free, as they would serve to advertise the contents of the article (but we may need greater adherence to truth in advertising, a feature that abstracts often seem to lack). The development of micropayment systems to enable small financial transactions should, in principle, allow unbundling. Indeed, why buy even an entire article? Get it page by page, paying initially for only the first page, and then for the second page only if the first page seems worth it.

While the internet creates the technological possibility of unbundling, it may not always be in the seller's interest to do so. Bundling may allow the seller to extract more revenue than if they would get if they sold the goods separately. The simplest way to understand this considers an electronic journal that has only two articles to sell and a potential readership of two individuals. Assume that the potential readers place different valuation on the articles, and the marginal cost of delivering each article is zero.

	Article 1	Article 2
Individual A	£2	£3
Individual B	£3	£2

Table 2: Individuals' willingness to pay for articles

If the publisher was aware of these valuations, they could personalise prices, making each individual pay exactly their valuation for each article. In the absence of this information, price discrimination is not possible. Suppose the publisher does not have precise information on valuations, but knows, for every article, the distribution of the valuations in the population. If so, selling the two articles as a bundle may be more profitable than selling them separately. To see this, let us look at the publisher's choice when setting the optimal price for Article 1. If they set the price at £3, only one individual would buy it. At the lower price of £2, both individuals would buy it, increasing the revenue to £4. Likewise, Article 2, sold individually is optimally priced at £2, yielding £4 as revenue. Total revenue from sell both articles separately is £8. However, selling the two articles as a bundled product can increase profits. Each individual is prepared to pay upto £5 for the two articles together, so that selling the bundle for £5 would increase revenue to £10. Bundling enhances revenue in this case, because while the readers have heterogeneous preferences for individual

articles, their valuation for the bundle is the same. The argument here is pretty general. By the law of large numbers, as long as individual's valuations are not correlated, their valuation of the bundle is likely to be less dispersed than their valuations for individual components of the bundle.

Sometimes it is even more profitable to sell goods both as a bundle and separately as individual components: the bundle is typically priced at less than the sum of the prices of individual components. This is called mixed bundling, and is more profitable than pure bundling if the marginal cost of producing and distributing the components are not trivial and the heterogeneity in consumers' preferences takes particular forms. Consider standard software such as Microsoft Office which bundles together software for word-processing, spreadsheet, presentation, etc. Most individuals may value these components differently and yet be prepared to pay a similar amount for the bundle as a whole. However, if there are individuals who do not care for the bundle at all, and are willing to pay a lot for, say, the spreadsheet, allowing them to buy that individually can be profitable. The OXREP practices mixed bundling too. You can subscribe to the journal but you can also buy individual issues at a higher unit price.[13]

Such aggregation extends in other dimensions. Site licensing of computer software is an alternative to selling it directly to individual users, and as such saves on administration costs and provides greater interoperability to all users at the site. But, as Bakos and Brynjolfsson (2000) argue, this is also a form of aggregation – it aggregates the preferences across different individuals who wish to use the software at a site – this typically increases the seller's revenue and enhances welfare. For instance, the value of a site license to a university is the sum of all individual valuations, so that within the university the software becomes a public good.

In sum, what are the implications of bundling for the manner in which information goods will be sold in online markets? Even though the technology of online transactions make it cheap to unbundle goods, we should expect bundling to survive. Indeed new forms of mixed bundling may emerge and concomitantly, a wide variety of pricing schemes to develop. For example, the seller can sell several bundles with different two-part prices. Intensive users might prefer a larger "entry" fee coupled with

[13] For an interesting discussion of alternative forms of pricing for online journals, see Chung-I Chuang and Sirbu (2000).

low prices for products in the bundle. Occasional users might prefer a lower entry fee and then pay somewhat higher prices for the parts of the (smaller) bundle they use. Creating a wide range of bundles is relatively cheap for online sellers.

(iv) Price Discrimination under Oligopolistic Competition

The above discussion of price discrimination abstracts from the effects of oligopolistic competition. What effect could competition have on profits, consumer welfare and aggregate efficiency? In general, modelling markets where firms compete by offering different versions and/or bundles is difficult, but Armstrong and Vickers (2001) use a novel approach. Since, ultimately, consumers rank bundles by the total amount of utility that each bundle generates, firms can be modelled as competing by directly offering 'utility' to consumers. The welfare effects can then be analysed by looking at the maximum utility that firms can provide subject to breaking even. The idea is that whenever it is true that under sufficiently strong competition, a firm can only attract consumers by delivering close to this maximum utility, restricting the ways in which a firm can deliver utility reduces this maximum. Thus, under strong enough competition, price discrimination tends to improve welfare[14]. This may prove to be a useful insight in the welfare analysis of price discrimination on the internet.

(v) Pricing and copyright enforcement

The discussion so far implicitly assumes that the goods and services in question are excludable — only those consumers who pay the posted price have access to the good or service. For information goods available in a digital form, this assumption is not innocuous. Napster.com distributes software that enables a subscriber to search through the hard disks of other subscribers to locate a piece of music (in MP3 format), and access it. As long as a single subscriber has access to a piece of music, others can download and listen to it without paying for it. Recently, the US Supreme Court has asked Napster to remove all copyrighted material from its database. However, there are numerous other 'swap' programmes that make virtual tracing of a

[14] Armstrong (1996) and Rochet and Choné (1998) analyze multi-product non-linear pricing under monopoly where consumers have private information about their tastes, and show that some buyers are excluded from the market and among those included, some buyers with different tastes buy the same bundle. Armstrong and Vickers (2001) and Rochet and Stole (2001) show that under competition the optimal pricing policy is much simpler – the need to compete dominates the need to screen and it is an equilibrium for all firms to offer efficient two-part prices.

file's usage nearly impossible. For certain information goods like textbooks, `experiencing' a free copy on the internet might even enhance demand for a printed version you have to pay for (which is perhaps more comfortable to use), but for goods like MP3 music files, this is less relevant. The music industry is understandably apprehensive of the possibility of file swapping. Varian (2000) examines the pricing of shared information goods under the assumption that groups such as libraries pay for access. However, if swapping cannot be prevented, libraries are unlikely to be able to solve the free-rider problem, and the market could break down. Such complex issues have no clear resolution at present. Perhaps one way to combat the problem is by developing better and better encryption technology.

IV RETAIL AUCTIONS

So far we have looked at pricing methods in which the sellers post prices. For many goods, a better alternative is to arrive at a price through an auction. Auctions are often used to sell goods for which buyers have heterogeneous, privately known valuations (e.g. antiques). In such cases, an auction allocates objects efficiently – the good is allocated to the person who values it the most – and raises greater revenue than a posted price would. In conventional markets setting up an auction has significant transaction costs: congregating in one place at a particular time was sensible only for transactions in high value items. The internet enables geographically-dispersed market participants to communicate with each other asynchronously and quite cheaply, so it quite feasible to extend auctions for buying and selling items of relatively low value too.

There is a large number of rival online auction sites. On any single day, eBay, the largest of among these, has a few million objects for sale, arranged in thousands of categories, and the total volume of transactions through internet auctions is over a billion dollars a year. Objects being auctioned include comic books, concert tickets, home office and laboratory equipment, exotic holidays, real estate, collectibles (such as antiques, toys, stamps, coins, magic cards and pokémon cards), and even (as reported by Lucking-Reiley (2000)) a date with 'an attractive woman trying to pay off her credit card debts'! Such auction markets have very significantly expanded the markets for items such as low-value collectibles, which could otherwise be sold only in small local markets. Sellers who could expect to sell only from stalls in the local flea market (or the local pub, in case of the indebted woman) suddenly find

themselves facing a world-wide customer base. Even expensive items traditionally auctioned through established auction houses are gradually moving to the internet.

As Lucking-Reiley (2000) points out, sellers with goods in limited supply and of unknown demand have most to gain from using auctions compared to other pricing methods. This is certainly relevant for collectibles, which has emerged as the single largest category of goods auctioned on the internet. Concert tickets, hotel rooms and household services such as plumbers, electricians, have similar features, and are likely to be sold increasingly through auctions.

Most of the auctions on the internet sell precisely one unit of some object. There are three well-known formats for single-unit auctions. In a first-price sealed-bid auction, the highest bidder wins and pays his own bid. In an 'English auction', bidders raise their bids sequentially until all but a single bidder drops out. And finally, in a `Dutch auction,' the price descends until someone is willing to buy at the current price. The Dutch auction is strategically equivalent to a first-price sealed-bid auction, and as William Vickrey showed, the English auction is equivalent to a second-price sealed-bid auction (Vickrey auction), in which the highest bidder wins and pays the second highest bid. These standard auctions are efficient mechanisms in that they allocate objects to the bidders with the highest valuations. Thus, pricing through auctions not only enhances revenue, but also improves economic efficiency.

The above discussion assumes that buyers have privately known valuations for the good being sold. However, for some objects, instead of such 'private values,' all buyers have a 'common value.' Imagine bidding for the right to drill a particular plot of land for oil. The actual amount of oil under the plot of land is the (unknown) common value for bidders, each of whom make their own estimates of this common value. A central result in auction theory states that if each bidder has a private value that is not correlated with the private values of other bidders, all the above auctions raise the same revenue. However, if bidders have common values, the English auction raises more revenue than the Vickrey auction, which in turn generates more revenue than the first-price sealed-bid (and Dutch) auction.

Among these, the English auction is probably the most popular and commonly-understood auction format. The optimal bidding strategy is particularly simple when bidders have private values - the optimal bid is simply the private value. Even under common values, as noted above, it raises the most revenue among all standard

auctions. Thus, unsurprisingly, most of the auction sites on the internet adopt variations of this format, although a few adopt the Vickrey format directly.

However, online versions of English auctions allow bidders to wait till a few seconds before the auction closes, and then just beat the current high bid. Such 'sniping' distorts the outcome. To solve this problem, some auction sites (such as Amazon) extend the end time to a few minutes after the last bid. Another solution is to have a 'proxy-bidding' system. Under such a system if a bidder places a bid above the current high bid, the price only jumps by the minimum increment over the current high bid. If another bidder now submits an even higher bid, price increases smoothly to just above the penultimate bid. Thus when the auction closes, the highest bidder wins, but only has to pay a small increment over the second highest bid. Thus proxy bidding implements a Vickrey auction, and it is incentive compatible for buyers to simply submit their true value to the seller. Note, however, that the proxy bidding solution holds only under private values. Under common values, there is no known true value, and buyers actually learn from observing the dynamic bidding behaviour of others. As expected, for auctions with a dominant common value (such as coin auctions analysed by Bajari and Hortacsu (2001) in which coins are investment rather than mere collectibles, and there is uncertainty about the exact quality of the coins), bidding tends to be concentrated towards the end. However, bidders are likely to have known private values for the bulk of goods and services sold through the internet – thus proxy bidding probably fares well most of the time.

As Lucking-Reiley (2000) shows, pricing behaviour in internet auctions also depends on the fee charged by the auction site to the seller. Sites that charge a higher listing fee to sellers make them more intent on selling – thus reserve prices are set lower. This also results in a greater proportion of auctions resulting in a sale. eBay charges significant listing fees and 54% of its auctions result in a sale. In contrast Yahoo! does not charge an entry fee to sellers and only 16% of its auctions result in a sale. This raises the question of optimal reserve price in internet auctions. This question is studied in a field experiment by Lucking-Reiley (1999b). He finds that an important distinguishing feature of internet auctions is that bidders consider entry costly, and thus entry is endogenous. The field experiment suggests that the optimal policy is to set a zero reserve price even if the seller has a positive value for the good himself. Thus considerations of inducing entry seem to be the dominant factor.

One particular auction format, somewhat unique to the internet, combines the notions of price discrimination and auctions. Priceline.com runs what is sometimes called a 'reverse auction' and has the stated objective of collecting demand information. Priceline asks bidders to bid for objects such as hotel bookings, airlines tickets and car rental offered by different sellers, and once bids are placed, sellers decide whether to accept the offers. Bidders can specify several aspects such as which brands they like, whether they are willing to be flexible about dates (in case of airline tickets). Clearly, such an auction helps price discrimination. Some buyers pay the regular market price, others - who are willing to wait and see if their offer of a lower price is accepted by any seller - pay a lower price. To the extent that this allows allocation of some of the excess capacity of sellers, such price discrimination improves social efficiency. Further, since demand from lower-value customers is likely to be fairly uncertain, auctioning the excess capacity allows efficient allocation and raises more revenue compared to creating a low-end version of the ticket and posting a single price for it. There are two alternative ways to think about such an auction. If all bids are submitted before the sellers decide which to accept, the highest bidders win and this is no different from a standard multi-unit sealed-bid first-price auction. However suppose buyers arrive randomly (say, by a Poisson process), and the sellers must accept or reject offers sequentially, the auction takes on features of a financial option. An airline seat (say) is very much like a put option that an airline holds, and, for each bid, must decide whether to exercise the option. This kind of trading format merits future research.

V CONCLUSIONS

It is widely believed that the internet promotes greater competition in markets. The ability to buy and sell online extends the reach of sellers and weakens location-specific market power. Online comparison of prices makes the process of price discovery easier. If online firms have cost advantages relative to their conventional rivals, we should expect new virtual entrants to capture substantial market shares and eventually exert downward pressure on prices in the offline world.

The existing evidence, though somewhat limited given the relative infancy of this new medium of commerce, paints a mixed picture. It is hard to deny that in some sectors the effect of electronic commerce has been to increase the intensity of

competition.[15] But in many sectors the emerging online markets display high concentration and prices substantially above marginal costs.

To some extent this is not hard to understand. While setup costs are relatively low in online markets, brand names and advertising come to play an important role. As Sutton (1991, 1998) has argued persuasively for advertising-intensive and technology-intensive markets in general, a fragmented market structure may not be quite sustainable in such circumstances.[16] Starting from a low degree of market concentration, each firm would have an incentive to escalate advertising to capture market share, and others must either respond, or be forced to exit as consumers are attracted away by advertising rivals. In such an environment the level of sunk costs rises endogenously and the equilibrium market structure is relatively concentrated, with a handful of firms operating with high fixed costs. The welfare implications of such market outcomes are not transparent. To the extent that the escalation of advertising costs is competitive, it serves no social purpose and reduces welfare. And if the equilibrium market structure involves firms operating with high fixed costs, competitive pricing will be unable to cover such costs and hence, be unsustainable. Price discrimination and other strategies that increase revenue may be consistent with the second-best outcome in such markets.

Online markets for homogeneous goods present an interesting test case for these hypotheses. In theory, the relative ease of price comparison and the low costs of switching between firms should increase the intensity of competition. Search and price comparisons are not perfect, in part because they involve psychic costs that all buyers are not prepared to pay, and because sellers can deliberately obfuscate search procedures. Switching costs are not that low, and online retailers have found various devices to increase these. For instance, we expect online loyalty schemes to become more prevalent in the future. Even when the underlying product is homogeneous, we expect online firms will create heterogeneity (by bundling the product with various services, for instance): all these will serve to reduce the intensity of price competition.

---

[15] See Borenstein and Saloner (2001) and Bakos (2001) for a discussion of some sectors where electronic commerce has enhanced competition.

[16] For an interesting application of Sutton's ideas to online markets, see Latcovich and Smith (2001).

Internet transactions generate a wealth of information, and the relative abilities of sellers and buyers to use this information may affect the market outcome. Sellers may be able to gather much more information on individual preferences and will use this to customise products and prices for individual consumers. In general where individual firms come to acquire market power, we should expect more price discrimination. With exogenous set-up costs, and with some degree of competition, price discrimination is likely to lead to welfare improvement. If sunk costs rise endogenously (as in advertising-intensive industries), the welfare analysis is more complex and merits more research.

REFERENCES

Armstrong, M. (1996), 'Multiproduct Nonlinear Pricing', *Econometrica*, **64**, 51-75.

Armstrong, M., and Vickers, J. (2001), 'Competitive Price Discrimination',' mimeo, Nuffield College, Oxford.

Bakos, Y. (2001), 'The Emerging Landscape for Retail E-Commerce', *Journal of Economic Perspectives,* **15,** 69-80.

Bakos, Y. and Brynjolfsson, E (2000), 'Aggregation and Disaggregation of Information Goods:Implications for Bundling, Site Licensing and Micropayment systems', in Brian Kahin and Hal Varian (eds) in *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*, The MIT Press.

Bailey, Joseph (1998), 'Electronic Commerce: Prices and Consumer Issues for three production – Books, Compact Discs and Software', technical report, OECD.

Bajari, Patrick and Ali Hortaçsu (2001), 'Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay auctions,' mimeo, Stanford University, available at http://www-econ.stanford.edu/faculty/workp/swp00004.pdf

Baye, M. and Morgan, J (2000), 'Information Gatekeepers on the Internet and Homogeneous Product Price Competition,' *American Economic Review.*

BBC (2000), 'Amazon's old customers pay more', online report available at http://news.bbc.co.uk/hi/english/business/newsid_914000/914691.stm

Beggs, A. and Klemperer, P (1992), `Multi-Period Competition with Switching Costs,' *Econometrica,* **60**, 3, 651-666.

Borenstein, S. and Saloner, G. (2001), 'Economics and Electronic Commerce', *Journal of Economic Perspectives,* **15**, 3-12.

Brown, J. and Goolsbee, A. (2000), 'Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry', NBER Working paper 7996.

Brynjolfsson, E. and Smith, M. (2000a), 'Frictionless Commerce? A comparison of Internet and Conventional Retailers', *Management Science.*

Brynjolfsson, E. and Smith, M. (2000b), 'The Great Equaliser? Consumer Choice Behaviour at Internet Shopbots', mimeo, MIT Sloan School.

Chung-I-Chuang John and Marvin Sirbu (2000), 'Network Delivery of Information Goods: Optimal Pricing of Articles and Subscriptions', in Kahin and Varian, *ibid*.

Clay, K., Krishnan, R., and Wolff, E. (2000), 'Pricing strategies on the Web: Evidence from Online Book Industry', mimeo, Heinz School, Carnegie Mellon University.

Corts, Kenneth (1996), 'On the competitive effects of price-matching policies', *International Journal of Industrial Organization*, **15**, 283-299.

DeLong, J. B. and Froomkin, M. (2000), 'Speculative Microeconomics for Tomorrow's Economy', in Brian Kahin and Hal Varian (eds) *ibid*.

Deneckere, R. J., and McAfee, R.P. (1996), 'Damaged goods', *Journal of Economics and Management Strategy,* **5**, 149-174.

Ellison, Glenn and Sara Fisher Ellison (2001), 'Search, Obfuscation, and Price Elasticities on the Internet', mimeo, MIT.

Goolsbee, A. (2000), 'Competition in the Computer Industry: Online versus Retail', mimeo, Graduate School of Business, University of Chicago.

Johnson, E., Moe, W., Fader, P., Bellman, S and Lohse, J. (2000), 'On the Depth and Dynamics of Online Search Behavior', mimeo, Columbia University.

Kephart, J.O., and Greenwald, A. (1998), 'Shobot Economics', technical report, IBM Institute for Advanced Commerce.

Klemperer, Paul (1999) 'Auction theory: a guide to the literature', http://www.nuff.ox.ac.uk/users/klemperer/papers.html

Latcovich, S. and Smith, H. (2000), 'Sinking the Internet: Pricing, Sunk Costs and Market Structure Online', mimeo, Oxford University

Lucking-Reiley, D. (1999a), 'Using Field Experiments to Test Equivalence between Auction Formats: Magic on the Internet', *American Economic Review*, 1063-1080.

Lucking-Reiley, D. (1999b), 'Experimental evidence on the endogenous entry of bidders in internet auctions', mimeo, Vanderbilt University, available at http://www.vanderbilt.edu/Econ/reiley/cv.html

Lucking-Reiley, D. (2000), 'Auctions on the Internet:: What's being auctioned, and how?', *Journal of Industrial Economics*, **48**, 227-252.

Lucking-Reiley, D., and Spulber, D.F. (2001), 'Business-to-Business Electronic Commerce,' *Journal of Economic Perspectives*, **15**, 55-68.

Odylzko, A. (1996), 'The Bumpy Road of Electronic Commerce', mimeo, ATT Research.

Rochet, J-C. and Choné, P. (1998), 'Ironing, sweeping and multi-dimensional screening', *Econometrica*, **66**, 783-826.

Rochet, J-C and Stole, L. (2001), 'Nonlinear pricing with random participation', *Review of Economic Studies*, forthcoming.

Shapiro, C. and Varian, H. (1999), `*Information Rules: A Strategic Guide to the Network Economy*,' Harvard Business School Press.

Smith, M. D., Bailey, J., and Brynjolfsson, E. (2000), 'Understanding Digital Markets: Review and Assessment' in Erik Brynjolfsson and Brian Kahin (eds.), *Understanding the Digital Economy: Data, Tools and Research*, MIT Press.

Sutton, J. (1991), *Sunk Costs and Market Structure*, MIT Press.

Sutton, J. (1998), *Technology and Market Structure*, MIT Press.

Varian, H. (1980), 'A Model of Sales', *American Economic Review*, **70**, 651-659.

Varian, H. (1985), 'Price Discrimination and Social Welfare', *American Economic Review*, **75**, 870-875.

Varian, H. (1989),`Price discrimination', in Richard Schmalensee and Robert Willig (eds), *Handbook of Industrial Organization*, North-Holland Press.

Varian, H. (2000a), 'Buying, Sharing and Renting Information Goods', mimeo, University of California at Berkeley.

Varian, H. (2000b), 'Versioning Information Goods' in Brian Kahin and Hal Varian (eds), *ibid*.

Wilson, R.B. (1993), `*Nonlinear Pricing'*, Oxford University Press.