Market Structure in the Network Age*

Hal R. Varian
University of California, Berkeley†

April 1999 (revised: August 30, 1999)

E-commerce will undoubtedly change the way business is done. But as we have said elsewhere, "technology changes, economic laws do not." Despite the changes introduced by e-commerce, many of the fundamental principles of competition will still be relevant.

In this paper I investigate three aspects of competition in ecommerce: marketing, interconnection, and price matching. In each case I will describe the phenomenon, illustrate its relevance for ecommerce, and describe some research issues raised.

1 Marketing

I will discuss three topics in marketing: versioning, loyalty programs, and promotions.

1.1 Versioning

I use the term "information good" to refer to a good that can be distributed in digital form. Examples are text, images, sounds, video, software, and so on. Information goods are characterized by having high fixed costs, or first-copy costs, but very low incremental costs. The challenge in pricing is to find a way to sell to a broad enough audience to cover those high first-copy costs.

*Prepared for *Understanding the Digital Economy* conference, May 25-26, 1999, Department of Commerce, Washington, DC.

†Thanks to Carl Shapiro and Marius Schwartz for helpful comments on an earlier version of this paper.

One way to accomplish this is to *version* the information good. This means offering a product line of variations on the same underlying good. The product line is designed so as to appeal to different market segments, thereby selling at a high price to those who have a high value for the product, and a low price to those who value it less.

Versioning is a common strategy for conventional information goods. Books are issued first in paperback and then in hardback; impatient, high-value users buy the hardback, while others wait for the paperback. Movies come out in theaters first, then are released 6 months later in home video.

But the flexibility of digital media offers many alternative forms of versioning. Shapiro and Varian [1998] identify several different types of versioning:

Delay Twenty-minute delayed stock quotes are given away, while a real-time feed may be costly.

User interface The professional version has an elaborate user interface; the popular version has a simple interface.

Convenience The low-price version is hard to use, but high-price version is simple to use.

Image resolution Low-resolution images sell for a low price; high-resolution images sell for a high-price.

Speed of operation The low-speed version is cheap; the high-speed version is expensive.

Flexibility of use A low-end software product may be used only for certain tasks while the high-end product is more flexible.

Capability The professional version has more capability and can do more things than the low-end version.

Features and functions The high-end version has more features and functions.

Comprehensiveness A high-end database or information service could be more comprehensive than the low-end.

Annoyance The low-end product uses "nagware," such as start-up delays or reminders, to induce the consumer to upgrade to a more expensive version.

Technical support The low-end product has no technical support; the high-end product offers this feature.

These are just a few of the dimensions on which you can version information goods. A notable feature of these dimensions is that the often involve first building the high-end product (the immediate, high-resolution, elaborate user-interface version) and then *degrading* it in some way to produce the low-end version. Often you have to go through *extra* processing or programming to create the low-end version of the product.

This, of course, raises public policy questions: should such deliberate product degradation be allowed? From the viewpoint of economic analysis, the critical issues are the size of the product degradation and whether the price differentiation increases or decreases the size of the market. The precise statement is this: if price differentiation reduces the size of the market, aggregate welfare necessarily decreases. On the other hand, if price differentiation increases the size of the market aggregate welfare may easily increase. See Varian [1985] for details.

"Aggregate welfare" counts both consumers' surplus and producers' surplus on an equal basis. A single producer does at least as well moving from a flat price to a differentiated price, since it always has the option of not differentiating. Since normally some prices go up and some go down when moving to differentiated pricing, some consumers are made worse off and some better off.[1] In balance, we expect that consumers would be worse off, but there are cases where price discrimination results in a Pareto improvement.[2]

Consider a simple example: a textbook sells in the US for $50 and a paperback, newsprint version of the same book sells for $5 in India. Does the low-quality version increase overall welfare? To answer this question we have to ask what version *would* have been produced if only one version were allowed. In this case, the likely answer is that only the high-quality, high-price version would have been produced. The ability to produce the low-quality, low-price version increases the availability of the good, and increases overall consumer surplus.

An even more dramatic example can be constructed where the cost of production is such that the product could not be produced at all without access to both

[1] Somewhat surprisingly, there are cases where all prices move in the same direction. See Nahat et al. [1990] for examples.

[2] See Hausman and MacKie-Mason [1988] for some theorems about when this can happen in the case of monopolistic price discrimination. In the case of competitive price discrimination, Armstrong and Vickers [1999] show that as long as the markets are competitive enough price discrimination normally results in increase in overall social welfare.

the US and Indian revenue streams. Here the ability to version, and price discriminate, is critical to the economic viability of the product in question. This case is rather common with information goods, due to the high first-copy costs.

One can imagine other cases with the opposite result—cases where versioning ends up reducing the total amount of the product sold. However, these cases do not seem to be very robust, and I believe that, in general, versioning tends to increase overall welfare.

1.2 Loyalty programs

On the Internet, the competition is just a click away. This has the potential to lead to intense price competition for commodity products. ACSES[3] is a search engine for books; it queries roughly 45 sites and reports back price, availability and shipping charges for all books. On a recent query the total price for a particular book varied from $24.07 to $40.94!

Given these dramatically reduced search costs, it is natural for firms to try to build customer loyalty. Obviously the best way to do this is to have low prices, high quality, and good service. But there are other ways: Amazon.com gives out free coffee mugs, t-shirts, upgraded shipping, and other bonuses. They offer customized services based on shopping history of their customers that would be difficult for competitors to imitate.

Another strategy that we expect to become more and more widespread are frequent-purchaser programs. Frequent-flyer programs have been around for 25 years and have dramatically changed the pricing and marketing programs of what is essentially a commodity business. Frequent-purchaser programs on the Internet have the same potential.

If all interaction with a customer is via a Web browser, it is trivial to capture that information and offer rewards to customers based on their purchase history. Intermediaries such as Cardpoints.com[4], SmartFrog.com[5], Data Marketing Group[6] and many others allow merchants to monitor and archive consumer purchase information, rewarding them with dollars, points, or prizes.

What should the reward system look like for such loyalty programs? Suppose that Amazon.com and Barnesandnoble.com both use a linear reward system—you get one point credit for each dollar spent that can be turned in for cash. This system

[3] http://www.acses.com
[4] http://cardpoints.com
[5] http://www.smartfrog.com
[6] http://www.datamarketinggroup.com

will be attractive to consumers, but won't build encourage exclusive use, since a point on BarnesandNoble.com is worth just as much as a point on Amazon.com. What the sellers should use is a *nonlinear* rewards scheme—a system that gives the customer a big prize after he has hit a certain level. This way, the consumer is encouraged to stay with one vendor rather than diversify across both vendors. Even better, the sellers could use a sequence of milestones, each more difficult to achieve, that are rewarded by prizes and/or through special services. This strategy is taken right from the airlines, who offer their frequent flyers free flights at certain milestones, and first-class upgrades at others.

If such loyalty programs become widespread, as we expect they will, price competition will tend to be reduced—just as it has with the airlines. The competition will occur upfront, in an attempt to encourage new users to sign up with one program or another. Consumers will still benefit from competition, due to the payments and prizes used to award loyal customers, but the heavy purchasers will benefit disproportionately.

1.3 Promotions and shopbots

Kephart and Greenwald [1998] have investigated the "economics of shopbots." In their model, some consumers have access to shopbots that can search out the lowest price for a generic product, while other consumers do not. As Kephart and Greenwald [1998] note, their model is similar to that of Varian [1980] in which some consumers read the newspaper for sales, while other consumers shop at random.

It is not surprising that the models generate similar equilibria: firms randomize their prices—hold "sales"— in order to price discriminate between the searchers and the non-searchers. Those who invest in using shopbots end up with a lower price, but at the cost of a more elaborate search. In addition, these price-sensitive customers have to give up the benefits conferred by the loyalty programs described above.

The sellers price discriminate between searchers and non-searchers by randomizing their prices. That way they compete for the searchers, when they happen to have the lowest price, but still manage to charge a relatively high price on average.

The question is, why would anyone *not* use a shopbot? Presumably the answer is the loyalty programs mentioned above: if the customer stays with one merchant, he can receive benefits (in the form of lower prices or coffee mugs) that cannot be offered by the low-price merchant.

This suggests that we will see a division between cheap, low-service merchants and high-cost, high-service merchants in cyberspace, just as we have seen in the ordinary marketing. Indeed, it may easily happen that the high-service and low-service merchants rely on the same underlying infrastructure. Clemons et al. [1998] describes a case study of seven on-line travel agencies and show exactly this outcome: the highest price and the lowest price merchants are both owned by the same company!

1.4 Research topics in marketing

The phenomena described above, versioning, loyalty programs, and promotions all raise interesting research questions.

- Online merchants will be collecting many megabytes of data about their customers' buying habits. How can these data be analyzed effectively? Many existing techniques rely heavily on parametric forms, but I expect that the massive amounts of data becoming available will readily allow for effective non-parametric analysis. Nevo and Wolfram [1999] has been applying some of these techniques to supermarket scanner data with considerable success.

- It is also worth learning about the people who don't buy. The "clickstream" captures the search process of online users and, to some extent, reflects the underlying cognition surrounding the purchase decision. It would be very helpful to understand this in more detail.

- Since we expect to see much more price discrimination for information goods, it would be helpful to have better tools for welfare analysis, especially for the case of quality discrimination. Armstrong and Vickers [1999] describe a very promising method of analysis that may help with this issue.

- Loyalty programs push the competition for consumers up front. However, after the consumer has chosen a merchant, they tend to be locked-in, allowing the merchant to exploit some monopoly power. In a world of forward-looking, rational consumers, the ex post monopoly would have small social cost. However, in a world with myopic consumers, the subsequent monopoly may be less benign. A better understanding of the welfare economics of loyalty programs would be helpful.

- A weakness of standard economic theory of production is that we tend to focus too much on the one-product firm. In the online world there will be

dramatic economies of scope. For example, firms can use a single underlying transactions technology that can then be branded in different ways, depending on the market they wish to attract. One would expect that competitive forces would be less strong in such environments, but it would be interesting to work out the details.

2 Interconnection

Economists say that there is a *network externality* when the value of a good depends on the number of other people who use it. Examples are goods like the telephone network, the fax machines network, the email network, or the Internet itself. Generally consumers would like to be connected to as large a network as possible. This implies that if there are several different providers of networks, then it is very advantageous to consumers if they interconnect.

In the examples above, telephones, faxes, and email are valuable precisely because they all work according to a common standards and anyone can call, fax or email anyone else who is connected to the network. The Internet is valuable because it is built on a common platform of open standards that allows many different networks to interconnect.

Despite the fact that interconnection is typically in the social interest, it may or may not be in the private interest. There may be cases where a large incumbent may find it attractive to avoid interconnection with new entrants in order to preserve its market power. Shapiro and Varian [1998] discuss several historical examples.

It is important to understand that if the value of the network increases through interconnection, then there should be a way to divide that increase in value so to make all participants better off. If the pie gets bigger, everyone can get a larger price. However, the increased size of the pie also means that *threats* not to interconnect become more significant. And, of course, a larger pie is a more tempting target for someone to try to snatch than a smaller one.

Let us see how these effects play out in a simple algebraic example. Suppose that the value of a network to a user is proportional to the total number of people on the network, n. For simplicity choose the constant of proportionality to be 1. The value of the entire network is then n^2, in accordance with "Metcalfe's Law."

If two networks of size n_1 and n_2 interconnect, what increase in value accrues

to each one? A simple calculation shows

$$\begin{array}{lll} \Delta v_1 & = n_1(n_1+n_2) - n_1^2 & = n_1 n_2 \\ \Delta v_2 & = n_2(n_1+n_2) - n_2^2 & = n_1 n_2 \end{array} \tag{1}$$

Note the surprising result that each network gets *equal* value from interconnecting. Each person in the large network gets a little bit of extra value from connecting to the small network, but there are a lot of people in the large network. Conversely, each person in the small network gets a lot of extra value from connecting to the large network, but there are only a few people in the small network.

This calculation, simple though it is, gives some insight into why "peering," or settlement-free interconnection, is common among large backbone providers. The gains from interconnection are split more-or-less equally, even among somewhat different size networks.

But, in a way, it proves too much, since not all networks are willing to interconnect on a payment-free basis. The answer to this seeming paradox is to look what happens if one network *acquires* the other. Suppose, for example, that network 1 pays network 2 its standalone value and then merges it with its own network. The increase in value is then

$$\Delta v_1 = (n_1+n_2)^2 - n_1^2 - n_2^2 = 2n_1 n_2.$$

In this case, network 1 captures twice as much value by buying out network 2 rather than interconnecting with it. This is why I said above that the threat of not interconnecting can be very valuable, since it can be used to induce another network to merge or be bought out. Essentially the threat of non-connection increases the larger network's bargaining power.

Of course, the linear-in-value assumption underlying Metcalfe's Law may be wrong, as suggested by Kling [1999]. Still, we might expect that the value function is *locally* linear, suggesting that network providers don't have to be perfectly symmetric to gain more-or-less equally from settlement-free interconnection.

2.1 Research topics in networks

- The theory of network effects is more evolved than the empirics, and it would be very helpful to have some detailed empirical analyses. This is very hard to do at the "macro" level, due to data limitations, but more feasible at the micro level. See Goolsbee and Klenow [1999] for a nice example of micro analysis of network externalities.

- Interconnection is likely to be a very contentious issues as the Internet evolves, and it is important to try to work out some sensible institutions to facilitate this. Varian [1998], for example, argues that an industry arbitration board might make sense. The problem with such a board would be the temptation to use it as a device for collusion.

- The strategic analysis of interconnection is in its infancy. Much of the analysis based on telecommunications deals with the asymmetric case of local telephony providers interconnecting with long-distance carriers. This is hard enough, but the symmetric case of Internet providers interconnecting is even more difficult. Crémer et al. [1999] is a nice start on this set of issues.

3 Price competition

I indicated above that the intense price competition would induce online merchants to look for ways to increase customer loyalty. Presumably they will also try to adopt pricing strategies that will reduce the intensity of the competition.

Web-based price comparison agents, sometimes known as "shopbots," have been generally viewed as being beneficial to consumers. But this is not so obvious since shopbots not only allow consumers easy access to other firms' prices, they also allow the firms themselves to monitor each others price movements.

For example, suppose that there are two dominant firms in a given industry, A and B. A adopts the following price strategy: whenever B cuts its price, A will immediately cut its price by the same amount. Whenever B raises its price, A will immediately raises its price by the same amount.

What will be the impact of this policy on price competition? The answer depends on how fast "immediately" is. If the consumers move more rapidly than the firms, then cutting price may be advantageous to the first price cutter, since the flood of extra consumers makes up for the (small) reduction in price necessary to attract them. But if the firms move faster than the consumers, then this may not be the case. Suppose, for example, that firm B matches firm A's price change *before* consumers can respond. In this case, there is no flood of consumers from price cutting, and the incentives to cut price are dramatically reduced. If both firms pursue the price-matching strategy, the equilibrium price is the same as if there were a single monopolist. The only check on this upward drift in prices comes from competitive suppliers such as local merchants who may find it more difficult to change prices so rapidly.

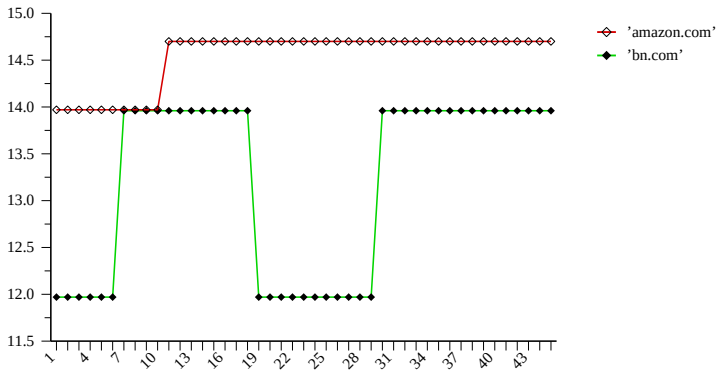

Figure 1: Price matching between Amazon.com and Barnesandnoble.com?

Does this sort of price matching occur online? Dillard [1999] offers some suggestive evidence. Figure 1 shows the timeline of the price of a bestseller on Amazon.com and Barnesandnoble.com. Amazon.com was selling the book at a higher price initially, and when Barnesandnoble.com raised its price, Amazon.com raised its price even higher. Amazon.com did not response to Barnesandnoble.com's price cut, and Barnesandnoble.com then returned to its higher price.

Another interesting example is depicted in Figure 2. Here we see Amazon.com leading by raising the price, with Wordsworth and BooksAMillion following shortly after.

Clearly these examples can only be taken as suggestive; it would certainly be worthwhile to look at more cases. One great advantage of the online world is that it is easy to monitor sellers' behavior.

In addition to the theoretical argument, we can also look at experience with pricing on more mature electronic markets. NASDAQ recently paid a $1.01 billion settlement to drop charges of price fixing on their exchange. The charges were prompted by an academic study described in Christie and Schultz [1995]. Apparently the traders were able to avoid selling securities at odd-eighth quotes by identifying deviators and refusing to trade with them. This increased the average spread, resulting in higher profits for traders.

Another illustrative example is airline reservation systems, such as SABRE. It is common to see competitive carriers engaging in "price signaling" and other techniques to maintain stable, oligopolistic prices. In this industry airlines clearly respond to each others' prices more rapidly than do consumers, and modern in-

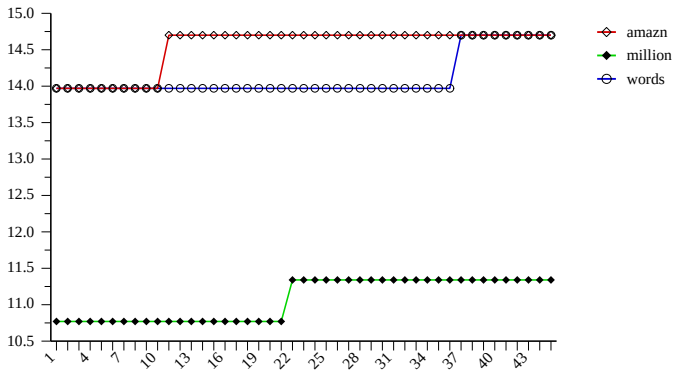

Figure 2: Amazon.com, Wordsworth.com and BooksAMillion.com

formation systems apparently help enforce higher prices. (See Nomani [1990] for a popular account of this practice.)

However, the airline system is different from online merchants in several ways. In particular, there are generally only a couple of major airlines in each city-pair market. As we saw above, there can be dozens of competitors in a market for a commodity product such as books. The price-matching strategy described above may work with two firms, but not with dozens.

Another difference is that entry costs are probably much lower for Internet vendors than for airlines. However, although anyone can set up a Web site, marketing costs can be prohibitive, especially if one has to compete with entrenched incumbents. Furthermore, marketing costs are entirely sunk costs in the sense that they aren't recoverable if the entrant fails to survive.

Customer loyalty is another issue. A new entrant may spend enough money to get noticed, but will it get patronized? Customers who are building up points through loyalty programs may be loathe to desert their current vendor for a slightly cheaper price.

One scenario is that we will see intense price competition until the small players are weeded out, leaving only a handful of online merchants who can engage in the kinds of pricing strategies described above. However, the threat of entry and competition from offline vendors will prevent significant abuse of market power. If this scenario is right, online markets may not be as cutthroat as is commonly expected.

3.1 Price matching: research questions

- There is an extensive literature on price matching that examines under what conditions it is pro- or anti-competitive; see Corts [1997] for a recent overview of this literature. However, the models are not very robust, and more analysis is needed. Shaffer [1999] describes some empirical work, which is very helpful. Such work should be easier to do online than off.

- How can we distinguish "competitive" price matching from "oligopolist" price matching? After all, we *want* firms to cut prices in response to their rivals' cuts. This has been a vexing problem with the airlines, and I don't see much hope of its being resolved easily. It is probably a better idea to focus on ease-of-entry or alternative providers as a means of price discipline.

- Upfront sunk costs are likely to play a large role in online industries. If competition were so intense as to really compete price down to marginal cost, there would be no way to recover those initial upfront investments. Perhaps strategies like loyalty programs and price matching allow for somewhat lessened competition is a healthy one. Clearly more research is needed to understand the nature of competition in such environments.

References

Mark Armstrong and John Vickers. Competitive price discrimination. Technical report, Nuffield College, Oxford, 1999.

William Christie and Paul H. Schultz. Are nasdaq spreads determined by competition or implicit collusion? *Journal of Economic Perspectives*, 9(3):199–208, 1995.

Eric K. Clemons, Lorin M. Hitt, and Il-Horn Hann. The nature of competition in electronic markets: A comparison of on-line travel agent offerings,. Technical report, Wharton School, University of Pennsylvania, 1998.

Kenneth S. Corts. On the competitive effects of price-matching policies. *International Journal of Industrial Organization*, 15(3):283–299, 1997.

Jacques Crémer, Patric Rey, and Jean Tirole. Connectivity in the commercial Internet. Technical report, Institut D'Économie Industrielle, Toulouse, 1999.

Martin Dillard. The economics of electronic commerce: A study of online and physical bookstores. Bachelor's honors thesis, UC Berkeley, 1999.

Austan Goolsbee and Peter J. Klenow. Evidence on learning and network externalities in the diffusion of home computers. Technical report, University of Chicago, GSB, 1999. `http://gsbadg.uchicago.edu/vitae.htm`.

Jerry A. Hausman and Jeffrey K. MacKie-Mason. Price discrimination and patent policy. *RAND Journal of Economics*, 19(2):253–265, 1988.

Jeffrey O. Kephart and Amy Greenwald. Shopbot economics. Technical report, IBM Institute for Advanced Commerce, 1998. `http://www.research.ibm.com/infoecon/researchpapers.html`.

Arnold Kling. The last inch and metcalfes law. Technical report, AIMST Essays, 1999. `http://home.us.net/~arnoldsk/aimst2/aimst212.html`.

Babu Nahat, Krzysztof Ostaszewsky, and P. K. Sahoo. Direction of price changes in third-degree price discrimination. *American Economic Review*, 80:1254–1258, 1990.

Aviv Nevo and Catherine Wolfram. Prices and coupons for breakfast cereals. Technical report, UC Berkeley, 1999. http://emlab.Berkeley.EDU/users/nevo/.

A. Nomani. Fare warning: How airlines trade price plans. *Wall Street Journal*, page B1, October 9 1990.

Greg Shaffer. On the incidence and variety of low-price guarantees. Technical report, University of Rochester, 1999.

Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide the Network Economy*. Harvard Business School Press, Boston, 1998. http://www.inforules.com.

Hal R. Varian. A model of sales. *American Economic Review*, 70:651–659, 1980.

Hal R. Varian. Price discrimination and social welfare. *American Economic Review*, 75(4):870–875, 1985.

Hal R. Varian. How to strengthen the Internet backbone. *Wall Street Journal*, page A22, June 8 1998.