

Markets for Information Goods

Hal R. Varian

University of California, Berkeley

April 1998 (revised: October 16, 1998)

Contents

| | | |
|-----------|--|-----------|
| 1 | Definition of information good | 3 |
| 2 | Information as an economic good | 3 |
| 3 | Information as experience good | 4 |
| 3.1 | Previewing and browsing | 4 |
| 3.2 | Reviews | 4 |
| 3.3 | Reputation | 5 |
| 4 | Returns to scale | 5 |
| 5 | Information as a public good | 6 |
| 5.1 | Economics of intellectual property | 7 |
| 5.2 | Software patents | 8 |
| 6 | Other ways to deal with exclusion | 9 |
| 7 | Terms and conditions | 10 |
| 8 | Piracy | 10 |
| 9 | International concerns | 11 |
| 9.1 | US as copyright pirate | 12 |
| 10 | Overload | 13 |
| 11 | Business models | 15 |
| 12 | Institutions | 16 |

Much has been written about the difficulties that “information” poses for neo-classical economics. How ironic that ICE—information, communication, and entertainment—now comprises the largest sector in the American economy. If information poses problems for economic theory, so much the worse for economic theory: real markets seem to deal with information rather well.

This paradox is the central theme of this essay: information, that slippery and strange economic good, is, in fact, handled very well by market institutions. The reason is that real markets are much more creative than those simple competitive markets studied in Econ 1. The fact that real-life markets can handle a good as problematic as is a testament to the flexibility and robustness of market institutions.

1 Definition of information good

Let us first seek a general characterization of the ICE economy. The basic unit that is transacted is what I call “information goods.” I take this to be anything that can be digitized—a book, a movie, a record, a telephone conversation. Note carefully that the definition states anything that *can* be digitized; I don’t require that the information *actually* be digitized. Analog representations, of information goods, such as video tapes, are common, though they will likely become less so in the future.

In this essay I will not be very concerned with asymmetric information. This topic has been dealt with extensively in the literature and I have little to add to the standard treatments. Instead, I want to focus on information as a good—as an object of economic transactions.

2 Information as an economic good

Information has three main properties that would seem to cause difficulties for market transactions.

Experience good. You must experience an information good before you know what it is.

Returns to scale. Information typically has a high fixed cost of production but a low marginal cost of reproduction.

Public goods. Information goods are typically non-rival and sometimes nonexcludable.

We will deal with these topics one at a time.

3 Information as experience good

You can only tell if you want to buy some information once you know what it is—but by then it is too late. How can one transact in goods that you have to give away in order to show people what they are? There are several social and economic institutions that are used to overcome this problem.

3.1 Previewing and browsing

Information producers typically offer opportunities for browsing their products: Hollywood offers previews, the music industry offers radio broadcasts, and the publishing industry offers bookstores, nowadays complete with easy chairs and cappuccinos. One of the great difficulties faced by sellers of information on the Internet is figuring out ways to browse the products. Video and previews work well, but it appears that previewing textual information would be quite difficult.

However, things are not quite as bad as they seem. The National Academy of Sciences Press found that when they posted the full text of book on the Web, the sales of those books went up by a factor of three. Posting the material on the Web allowed potential customers to preview the material, but anyone who really wanted to *read* the book would download it. MIT Press had a similar experience with monographs and online journals.

3.2 Reviews

Another way to overcome the experience good problem is for some economic agents to specialize in reviewing products and providing these evaluations to other potential consumers. This is especially common in the entertainment industry: film reviews, book reviews, and music reviewers are ubiquitous.

But reviews are also found in the purer sort of information goods. The most academic popular papers (as measured by citation) are typically surveys since the specialization required for frontier work in the sciences has created a demand for such overviews.

Peer review is the standard technique used in the sciences for evaluating the merit of papers submitted for publication, while most humanities use academic presses to provide a similar function. This institution survives because it meets an important need: evaluating information.

3.3 Reputation

The third way that producers of information goods overcome the experience good problem is via reputation. I am willing to purchase the *Wall Street Journal* today because I have read it in the past and found it worthwhile. The *Journal* invests heavily to establish and maintain its brand identity. For example, when it started an online edition, it went to great lengths to create the same “look and feel” as the print edition. The intent was to carry over the reputation from the off-line edition to the online version.

Investing in brand and reputation is standard practice in the information biz, from the MGM Lion to the *Time* magazine logo. This investment is warranted because of the experience good problem of information.

4 Returns to scale

Information is costly to produce but cheap to reproduce. It can easily cost over a hundred million dollars to produce the first CD of a Hollywood film, while the second CD can cost well under a dollar. This cost structure—high fixed costs and low marginal costs—cause great difficulties for competitive markets.

It’s even worse than that. The fixed costs for information goods are not just fixed—they are also sunk. That is, they typically must be incurred prior to production and usually are not recoverable in case of failure. If the movie bombs, there isn’t much of a market for its script, no matter how much it cost to produce.

Competitive markets tend to push price to marginal cost, which, in the case of information goods, is close to zero. But this leaves no margin to recover those huge fixed costs. How is it that information can be sold at all?

The obvious answer is that information is rarely traded on competitive markets. Instead, information goods are highly differentiated. Each pop CD is different than the others (or so the listeners think), and each movie is unique. But not *too* unique. There is still an advantage in encouraging some similarities, due to the reputation effect described earlier.

The market structure for most information goods is one of monopolistic competition. Due to product differentiation, producers have some market power, but the lack of entry restrictions tends to force profits to zero over time.

The fact that information goods generally have some degree of market power also allows producers to recover fixed costs through more creative pricing and marketing arrangements. Price discrimination for information is common: different groups of consumers pay different prices, and quality discrimination is commonplace.

Publishers first issue a book in hardback and then, a year later, in paperback. Film come out first for theaters, then then, 6 months later, on videos. Investors pay one price for real time stock prices and another much lower price for delayed prices. In each of these examples, the sellers use delay to segment in the market by willingness to pay.

There are many other dimensions along which one can “version” information goods. Shapiro and Varian [1998] describe several of these dimensions including delay, user interface, convenience, image resolution, format, capability, features, comprehensiveness, annoyance, and support.

5 Information as a public good

A pure public good is both nonrival and nonexcludable. Nonrival means that one person’s consumption doesn’t diminish the amount available to other people, while nonexcludable means that one person cannot exclude another person from consuming the good in question. Classic examples of pure public goods are goods like national defense, lighthouses, TV broadcasts, and so on.

The two properties of a public good are quite different. Nonrivalness is a property of the good itself: the same amount of defense, lighthouse services and TV broadcasts are available to everyone in the region served by the very nature of the good. Excludability is a bit different since it depends, at least in part, on the legal regime. For example, TV broadcasts in England are supported by a tax on TVs; those who don’t pay the tax are *legally* (but not technologically) excluded from watching the broadcasts. Similarly, in the US cable TV broadcasts may be encrypted and special devices are required to decode them.

For that matter, it is “merely” a legal convention that ordinarily private goods are excludable. If I want others to be prevented from consuming my car for example, I either have to use technology (such as locks) or legal authority (such as police) to prevent them.

Even such classic examples as street lights *could* be made excludable if one really wanted to do so. For example, suppose that the lights broadcast only in infrared, and special goggles were required to take advantage of their services. Or, if this seems like too much trouble, cities could offer “streetlight licenses,” the purchase of which would be required to use streetlight services. Those who don’t go out after dark, don’t need to buy.

This isn’t as farfetched as it seems. Coase [1988] describes how the English authorities collected payment for lighthouse services based on the routes followed by ocean-going vessels.

Exclusion is not an inherent property of goods, public or private, but is rather a social choice. In many cases it is *cheaper* to make a good such as streetlights universally available rather than make them excludable, either via technology or by law.

These observations have bearing on information goods. Information goods are inherently nonrival, due to the tiny cost of reproduction. However, whether they are excludable or not depends on the legal regime. Most countries recognize intellectual property laws that allow information goods to be excludable. The US Constitution explicitly grants Congress the duty “...to promote the progress of science and useful arts, by securing, for limited times, to authors and inventors, the exclusive right to their respective writings and discoveries.”

5.1 Economics of intellectual property

The key phrase in the above quotation is “for a limited time.” Intellectual property law recognizes that no exclusion would create poor incentives for the creation of IP. But at the same time, *permanent* intellectual property rights would lead to the standard deadweight losses of monopoly.¹

Length is only one of the parameters of intellectual property protection. The others are “height”, in the sense of the standard required for novelty, and the “breadth”, in the sense of how broadly the IP rights are interpreted. Different forms of IP have different combinations of these characteristics; for example, copyright protects the expression of ideas for quite long periods (up to 75 years), with a low standard for novelty, but a narrow scope.

There has been much economic analysis of intellectual property protection for

¹Actually, this is not so obvious. If monopoly owner of information goods engage in price discrimination, as they commonly do, the deadweight losses may be much less than those generated under a single-price regime. This point definitely requires further investigation.

patents. Nordhaus [1969] examined the optimal length of a patent, finding that 20 years was not unreasonable. Scotchmer [1991] noted that invention is often cumulative and that shorter patent lives could lead to reduced incentives to invent, but more invention due to the ability to build on to earlier inventions.

Several authors, such as Dasgupta and Stiglitz [1980] and Gilbert and Newbery [1982], have recognized that the “prize” nature of patents leads to socially wasteful duplication of effort. The patent system sets up a race, which can cause firms to devote more resources to speeding up their discoveries than would be justified by a benefit/cost test. Suppose, for example, that a number of research teams were on the verge of making an important discovery, perhaps one that was the next logical step along a well-known research path. Granting the winning team long-term exclusive rights merely because they were slightly faster than others to make a discovery could well create more monopoly power than was necessary to elicit the innovative effort, and slow down future invention as well.

There has been much less investigation of the economics of copyright. The first problem is that existing copyright terms appear to be much too long from an economic point of view. At conventional interest economic transactions 30 or 40 years in the future are of negligible value so copyright terms of 50-75 years seem much too long to be based on economic calculation.

In fact as recently as the late 1960s copyrights only lasted 28 years in the US. Each subsequent reform of copyright law increased the term. The difficulty has been that each term extension grandfathered in the existing copyrights; even though no one would be willing to bargain seriously over possible cash flows 50 years down the road, the owners of about-to-expire and still valuable copyrights had significant economic incentive to extend them.

5.2 Software patents

Up until recently, the US Patent Office and the courts interpreted algorithms as “mathematical formulas” which could not be patented. However, in the mid eighties they reversed this policy and began to issue patents for software algorithms. Subsequently the patent office has issued many thousands of software patents.

There are several policy issues raised by software patents. First, until the last five years, the patent office has not had adequate expertise to evaluate the novelty of submitted patents. This has resulted in ludicrous examples such as the Compton patent on multimedia, the UCSF patent on downloading executable code, and the Software Advertising Corporation’s patent on incorporating advertising into

software programs.²

Secondly, there is the problem of “submarine patents:” patents that are not publicly available due to the fact that they are under consideration by the Patent Office. In some cases, applicants have allegedly purposely delayed their applications in order to wait for the market to “mature” so as to maximize the value of their patents, and to let them make improvements before others are apprised of their basic patent. These tactics can distort the returns to patent holders, frustrate the disclosure of patented inventions, which is a basic *quid pro quo* for patent protection under our patent system, and lead to unnecessary duplication of effort and lawsuits. The recent change in patent lifetime to twenty years after filing has gone a long way to reduce the problem of submarine patents.

Many of these problems are especially severe for software patents. Innovations that are embodied in physical goods can be bought and sold for a listed price on the open market, so there is no uncertainty about the cost of incorporating a new innovation into a product.³ However, the market for software components is still primitive, so much software is created in house. Thus, one software developer can easily infringe upon another developer’s algorithm, and, after years, find itself in a very vulnerable position if the algorithm ends up being patented.

All these reasons suggest that that patents on algorithms should be narrowly interpreted, and subject to high standards of novelty. Davis et al. [1994] also argue that software patents should have a shorter lifespan than other types of patents. Each of these policies should be carefully considered. As a practical matter, it would be far easier for the PTO to set high novelty standards and grant narrow software patents than for Congress to selectively alter patent lifetimes for software patents. Furthermore, in many cases the patent lifetime is unimportant, because the pace of progress is great enough that the patent has lost all of its value by its expiration date.

6 Other ways to deal with exclusion

Assigning of property rights are not the only way to deal with intellectual property issues. A second way is to bundle the content with a good that *is* excludable. In-

²Indeed, Bruce Lehman, the Commissioner of the PTO, has conceded that there a number of software patents were granted in error.

³Also, under the first-sale doctrine of patent law, a patent holder (or applicant) who sells an item containing the patented technology loses the right to further restrict the use of that item in commerce.

deed traditional media for transmitting information goods, such as book, records, video tapes, CDs, and so on are a type of bundling. Only one person can read a book at a given time, so exclusion is not much of a problem.

This doesn't work for purely digital information goods, since the medium itself doesn't have much significance, but recent technologies like cryptographic envelopes play similar role by bundling the information good with an "excludable" authentication mechanism.

A third technique for dealing with the exclusion problem is using auditing or statistical tracking. ASCAP and BMI perform this task for the music industry while the Copyright Clearance Center deals with print media by auditing photocopying practices over a period of time and bases a yearly fee on this sample.

A fourth technique for deal with exclusion is to embrace it, and bundle the information good with information that sellers *want* to be widely disseminated such as advertising.

7 Terms and conditions


Intellectual property law assigns default property rights to users, but licenses and other forms of contract can specify other terms and conditions. This contracting choice poses an interesting tradeoff: more liberal terms and conditions will generally increase the value a particular information good to its potential users, but it will also decrease the quantity sold. That is, a license to an information good that can be shared, resold, archived, etc. will be worth more than one that cannot; however, sharing, resale, and archiving all potentially reduce the final demand for the information goods.

Roughly speaking more liberal terms and conditions increase the value of the information good, shifting the demand curve *up*. However, liberal terms and conditions also reduce the sales of the good, shifting the demand curve *in*. The profit-maximizing choice of licensing terms balances these two effects.

8 Piracy

Simply specifying terms and conditions or intellectual property laws does not ensure that they will be enforced. Illicit copying is a perennial problem.

Luckily, as with most contraband, there is a mitigating factor. In order to sell illicit copies to consumers, they must know where to find the copies. The larger



Source: /Users/hal/Books/StrategicTech/Figures/piracy.quantrix

Figure 1: Per-capita income v fraction of software that is pirated for various countries.

the scale of operation of an IP pirate the more likely it will be detected by the authorities. This means that in equilibrium, reasonable efforts to enforce the law lead to relatively small scales of operation. Varian [1998] offers a model of this phenomenon.

9 International concerns

According to estimates from the Software Publishers Association, there are many countries where software piracy is rampant. Figure 9 shows the relationship between per capita income and the fraction of illegal software in use in various countries.

Figure 9 shows that the lower the per capita income, the higher the incidence of illegal copies. This should not be surprising. Lesser developed countries have little to lose if they pirate software and have neither the resources nor the inclination to invest in enforcement.

The same effect shows up in environmental practices. In general, the lower the per capita income the less environmentally aware a country is. As per capita income grows so does the desire for a cleaner environment. Once a country passes

\$5,000 or so of per capita income they start to institute environmentally-aware policies. See Coursey [1992] and Grossman and Krueger [1991].

We expect that the same effect will occur with intellectual property piracy. As countries become richer, their desire for local content increases. But as they get more and more local content produced, the necessity of intellectual property protection becomes more and more apparent. As enforcement of intellectual property laws increase, both domestic *and* foreign producers benefit.

Taiwan is a prime example. They refused to sign the International Copyright Agreement until recently. Prior to this Taiwan was notorious for intellectual property violations. However, once the country became prosperous and developed a large publishing industry, they joined the international copyright agreement in order to assure a market for their own publishing and printing industry.

9.1 US as copyright pirate

The history of international copyright policy in the US is an instructive example of what to expect from today's underdeveloped countries.

The US Constitution gave Congress the authority to create laws regulating the treatment intellectual property. The first national copyright law, passed in 1790, provided for a 14-year copyright ...but only for authors who were citizens or residents of the US. The US extended the copyright term to 28 years in 1831, but again restricted copyright protection only to citizens and residents.

This policy was unique among developed nations. Denmark, Prussia, England, France, and Belgium all had laws respecting the rights of foreign authors. By 1850, only the US, Russia and the Ottoman Empire refused to recognize international copyright.

The advantages of this policy to the US were quite significant: they had a public hungry for books, and a publishing industry happy to publish them. And a ready supply was available from England. Publishing in the US was virtually a no-risk enterprise: whatever sold well in England was likely to do well in the US.

American publishers paid agents in England to acquire popular works, which were then rushed to the US and set in type. Competition was intense, and the first to publish had an advantage of only days before they themselves were subject to copying. Intense competition leads to low prices. In 1843 Dickens's *Christmas Carol* sold for six cents in the US and \$2.50 in England.

Throughout the nineteenth century, proponents of international copyright protection lobbied Congress. They advanced five arguments for their position: (1) it was the moral thing to do; (2) it would help create domestic authors; (3) it

would prevent the English from pirating American authors; (4) it would eliminate ruthless domestic competition; and, (5) it would result in better quality books.

Dickens toured the US in 1842 and pleaded for international copyright on dozens of occasions. American authors supported his position, but their pleading had little impact on the public at large or on Congress.

It was not until 1891 that Congress passed an international copyright act. The arguments advanced for the act were virtually the same as those advanced in 1837. Although arguments were the same, but the outcome was different. In 1837 the US had little to lose from copyright piracy. By 1891 they had a lot to gain from international copyright—the reciprocal rights granted by the British. On top of this was the growing pride in purely American literary culture and the recognition that American literature could only thrive if it competed with English literature on an equal footing.

The only special interest group that was dead set opposed to international copyright was the typesetters union. The ingenious solution to this problem was to buy them off: the Copyright Act of 1891 extended protection only to those foreign works that were typeset in the US!⁴

There is no question that it was in the economic self-interest of the US to pirate English literature in the early days of nationhood, just as it is clearly in the economic self-interest of China and other LCDs to pirate American music and videos now. But as these countries grow and develop a longing for domestic content, they will likely follow the same path as the US and restrict foreign competition to stimulate the domestic industry.

10 Overload

Herbert Simon once said that a “wealth of information creates a poverty of attention.” This has become painfully obvious with the advent of the World Wide Web.

Despite the hype, the Web just isn’t all that impressive as an information resource. The static, publicly accessible HTML text on the Web is roughly equivalent in size to a million books. The UC Berkeley Library has 8 million volumes, and the average quality of the Berkeley library content is much, much higher! If 10% of the material on the Web is “useful,” then that means there are about

⁴This provision remained in effect until the mid-sixties! Our source for this discussion is Clark [1960].

100,000 useful book-equivalents on the Web, which is the size of good public library. The actual figure for “useful” is probably more like 1%, which is 10,000 books, or half the size of an average mall bookstore.

The value of the Web lies not in the quantity of information but rather its *accessibility*. Digital information can be indexed, organized, and hyperlinked relatively easily compared to textual information. A text is just a click away rather than a drive across town and an hour in the library.

But, of course, it isn't that simple. We've invested hundreds of millions of dollars in catalogs and cataloging for textual information, while cataloging online information is in its infancy. The information on the Web is highly accessible ... once you know where to look.

The publishing industry has developed a variety of institutions to deal with this problem: reviewers, referees, editors, bookstores, libraries, etc. There are whole set of institutions to help us find useful information. But where are the Better Bit Bureaus for the Internet?

The problem is getting worse. I would like to coin a “Malthus's law” of information. Recall that Malthus noted that number of stomachs grew geometrically but the amount of food grew linearly. Pool [1984] noted that the supply of information (in virtually every medium) grows exponentially whereas the amount that is consumed grows *at best* linearly. This is ultimately due to the fact that our mental powers and time available to process information is constrained. This has the uncomfortable consequence that the fraction of the information produced that is actually consumed is asymptoting towards zero.

Along with Malthus's law of information, I may as well coin a Gresham's Law of Information. Gresham said that bad money drives out good. Well, bad information crowds out good. Cheap, low quality information on the Internet can cause problems for providers of high-quality information.

The *Encyclopedia Britannica* offered an Internet edition to libraries with a site license subscription price of several thousand dollars. *Microsoft's Encarta* retails for \$49 for a CD ROM. *Encarta's* doing fine; but *Brittanica* is in serious trouble. *Brittanica* is now offering a home subscription for \$150 per year, and a home CD version for \$70, but even this may be too high.

So perhaps low-quality information really does drive out good. Maybe ... but Gresham's law really should be restated—it's not that bad money crowds out good, but that bad money sells at a discount. So bad information should sell at a discount. Good information—relevant, timely, high-quality, focussed, and useful information—like the *Britannica*—should sell at a premium. And this brings me back to the Better Bit Bureaus. The critical problem for the commercial providers

of content is to find a way to convince the user that they actually have timely, accurate, relevant, and high-quality information to sell.

When publishing was expensive, it made sense to have lots of filters to determine what was published and what wasn't: agents, editors, reviewers, bookstores, etc. Now publishing is cheap: anyone can put up an homepage on the Web. The scarce factor is *attention*. The 0-1 decision of "publish-or-not" no longer makes sense—what we need are new *institutional* and *technological* tools to determine where it is worthwhile to focus our attention.

They aren't here yet, but some interesting things are happening in this area.

One interesting approach involves *recommender systems* such as Firefly, GroupLens, etc. In FireFly you are presented with a list of old movie titles and you indicate which ones you like and dislike. The computer then finds people who have similar tastes to yours and shows you *recent* movie titles that *they* liked—with the implication that you might like them too.

In GroupLens participants rate news items that they read. When you are presented with a list of items to examine, you see a weighted average of the ratings of previous readers. The gimmick is that the weight that each person receives in this average depend on how often you have agreed with that person in the past.

Systems like FireFly and GroupLens—what we're calling "recommender systems" or "collaborative filtering systems"—allow you to "collaborate" with others who have common interests, and thus reduce your own search costs.

11 Business models

How do you pay for recommender systems? What's the economic model? There are several problems.

First, there is the issue of incentives. How do you ensure that people contribute honestly to the system? First, observe that if you can get them to contribute, it is in their interest to do it honestly. If a user of Firefly just clicks at random, then he messes up the correlations on which the system depends.

The big problem is getting people to contribute at all. Once I've seeded the system with my preferences, what is my incentive to continue to rate new movies? If I go to a movie that no one has rated, then I may see a bad movie. But everyone only goes to movies that someone else has rated, then who rate the unrated movies?

There are two solutions to this problem: you can pay people to do the ratings, or you can exclude people who refuse to do their fair share of ratings. The first

solution is the way Siskel and Ebert make their living: they specialize in recommendations and get paid by people who find their recommendations useful. The second way makes more sense in a community rating system: either you provide an appropriate share of the ratings or you are excluded from the system.

Getting people to contribute to knowledge bases—recommendations or any other sort of information—can be quite difficult. One of the major consulting firms has spent millions of dollars setting up a knowledge base. When the consultants finish a project they're supposed to file a report of useful material. I asked one of the consultants how this worked. His somewhat sheepish reply was that he was 6 months behind in filing his reports. The reason was, he said, that every time he posted something useful, he got 15 emails the next day asking him for more information! The system had *negative* incentives! The consulting firm had spent millions to set up the technology, but hadn't thought through the incentive problem. Oh well, they can always hire a consultant . . .

The production of knowledge is a tricky thing. By its nature it is easy to copy and share. And since it costs nothing to share, it is socially efficient to do so. But then how do we compensate the people that produce knowledge in the first place?

Conventional methods for protecting intellectual property don't apply: ideas can't be patented, and copyright only protects the expression of ideas, not the ideas themselves.

Let me suggest that one place that firms might look for ways to provide incentives for knowledge production is by looking to the industry whose entire economic base is knowledge—by that I mean academia. The academic system has lots of peculiar features: publish or perish, tenure, plagiarism taboos, peer review, citation, etc. When you look at these features you see that most of them are designed to provide incentives to produce good ideas.

Take tenure for example. As Carmichael [1988] points out, one role of tenure is to encourage experts to truthfully evaluate people who are close substitutes for themselves. It's hard to get people to hire their own replacements—unless you offer them a tenure guarantee that says they won't be replaced.

12 Institutions

Another approach to the filtering problem is the institutional approach: creating the equivalents of the editors, publishers, and reviewers for online content. This is the strategy of AOL, CompuServe, and Microsoft. They hope to become the intermediaries that filter and organize online information for the masses.

I have my doubts about this strategy. I think that the “mass market” is going to be less significant in the future than it has in the past.

One of the most striking features of the print media in the last 20 years has been the demise of the newspapers and the rise of the magazine. Most major cities have only one newspaper; and in those few cities with two newspapers, it’s pretty clear that one is going to go.

But you can now get magazines for just about every possible interest group, from butterfly collectors to body builders—and there is probably one for those who do both!

The same thing has happened with TV. In the last 10 years the big 3 TV networks have seen their market share drop while dozens of new channels have sprung up to serve niche markets. The Science Fiction Channel, the Discovery Channel, the History Channel are all offering content targeted to those with very specific interests.

I think that the Internet will accelerate this trend. People will be able to coalesce around their particular interest, be it butterfly collecting or bodybuilding. Everybody who wants to will be a publisher. Editors will filter with respect to topic and quality—but there will be lots and lots of different editors to choose from, so the search problem for individual users will be just as severe, if not more so, than it is now.

There’s no getting away from the fact that information management is going to be a bigger and bigger part of our lives. We’ll need to have better tools to do this task ourselves, and we’ll need to utilize information management specialists when necessary. Whether we are producers or consumers of information we will need additional expertise to help us locate, organize, filter, retrieve and use the information we need.

This expertise is what we have set out to produce at Berkeley. We’ve created a School of Information Management and Systems, whose mission is twofold: our research mission is to produce more powerful tools to manage information and our teaching mission is to train the information management specialists of the future. We’re giving our students a core curriculum with computer science, library science, law, and management. After these core courses, the students will take electives in areas of specialization such as electronic documents, archiving, data bases, information retrieval, human-computer interface, and so on.

Our students will be skilled in building and using information management tools. We think this expertise will be attractive to anybody who needs to manage information—which means just about everybody, these days. Whether you are a producer or a consumer, a professional or a dilettante, you’ve got some informa-

tion to manage—and our students will be there to help you do it.
So take heart—help is on the way!

References

- Lorne Carmichael. Incentives in academics: Why is there tenure? *The Journal of Political Economy*, 96(3):453–472, 1988.
- Aubert J. Clark. *The Movement for International Copyright in Nineteenth Century America*. Catholic University of American Press, Washington, DC, 1960.
- Ronald Coase. *The firm, the market, and the law*. University of Chicago Press, Chicago, 1988.
- Don Coursey. The demand for environmental quality. Technical report, Public Policy School, University of Chicago, 1992.
- P. Dasgupta and J. Stiglitz. Uncertainty, market structure and the speed of r&d. *Bell Journal of Economics*, pages 1–28, 1980.
- Randall Davis, Mitchell Kapor, J. H. Reichman, and Pamela Samuelson. A manifesto concerning the legal protection of computer program. *Columbia Law Review*, 94, 1994.
- Richard J. Gilbert and David M. Newbery. Preemptive patenting and the persistence of monopoly. *American Economic Review*, 72:514–526, 1982.
- Gene M. Grossman and Alan B. Krueger. Environmental impacts of a north american free trade agreement. Technical report, Department of Economics, Princeton University, 1991.
- William Nordhaus. *Invention, Growth, and Welfare*. MIT Press, Cambridge, MA, 1969.
- Ithiel De Sola Pool. *Communications Flows: A Census in the United States and Japan*. Elsevier Science, New York, 1984.
- Suzanne Scotchmer. Standing on shoulders of giants: Cumulative innovation and patent law. *Journal of Economic Perspectives*, pages 29–42, 1991.
- Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide for the Network Economy*. Harvard Business School Press, Cambridge, MA, 1998.
- Hal R. Varian. Intellectual property piracy. Technical report, SIMS, UC Berkeley, 1998.