

CHAPTER 35

INFORMATION TECHNOLOGY

One of the most radical changes in the economy in the last 15 years has been the emergence of the **information economy**. The popular press is filled with stories about advances in computer technology, the Internet, and new software. Not surprisingly, many of these stories are on the business pages of the newspaper, for this *technological* revolution is also an *economic* revolution.

Some observers have gone so far as to put the Information Revolution on a par with the Industrial Revolution. Just as the Industrial Revolution transformed the way *goods* were produced, distributed, and consumed, the Information Revolution is transforming the way *information* is produced, distributed, and consumed.

It has been claimed that these dramatically new technologies will require a fundamentally different form of economics. Bits, it is argued, are fundamentally different than atoms. Bits can be reproduced costlessly and distributed around the world at the speed of light, and they never deteriorate. Material goods, made of atoms, have none of these properties: they are costly to produce and transport, and they inevitably deteriorate.

It is true that the unusual properties of bits require new economic analysis, but I would argue that they do not require a new *kind* of economic

analysis. After all, economics is primarily about *people* not *goods*. The models we have analyzed in this book have had to do with how people make choices and interact with each other. We have rarely had occasion to refer to the specific goods that were involved in the transactions. The fundamental concerns were the tastes of the individuals, the technology of production, and the structure of the market, and *these* same factors will determine how markets for information will work . . . or not work.

In this chapter we will investigate a few economic models relevant to the information revolution. The first has to do with the economics of networks, the second with switching costs, and the third with rights management for information goods. These examples will illustrate how the fundamental tools of economic analysis can help us to understand the world of bits as well as the world of atoms.

35.1 Systems Competition

Information technology is generally used in *systems*. Such systems involve several components, often provided by different firms, that only have value if they work together. Hardware is useless without software, a DVD player is useless without DVD disks, an operating system is worthless without applications, and a web browser is useless without web servers. All of these are examples of **complements**: goods where the value of one component is significantly enhanced by the presence of another component.

In our discussion of consumer theory, we described left shoes and right shoes as complements. The cases above are equally extreme: the best computer hardware in the world can't function unless there is software written for it. But unlike shoes, the more software that is available for it, the more valuable it becomes.

Competition among the providers of these components often have to worry just as much about their "complementors" as their competitors. A key part of Apple's competitive strategy has to involve their relations with software developers. This gives competitive strategy in information technology (IT) industries a different flavor than strategy in traditional industries.¹

35.2 The Problem of Complements

To illustrate these points, let us consider the case of a Central Processing Unit (CPU) and an Operating System (OS). A CPU is an integrated

¹ See Shapiro, Carl and Hal R. Varian, *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, 1998, for a guide to competitive strategy in IT industries.

circuit that is the “brain” of a computer. Two familiar manufacturers of CPUs are Intel and Motorola. An OS is the software that allows users and applications to access the functions of the CPU. Apple and Microsoft both make operating systems. Normally, a special version of an operating system has to be created for each CPU.

From the viewpoint of the end user, the CPU can only be used if there is a compatible operating system. The CPU and the OS are complements, just as left shoes and right shoes are complements.

The most popular CPUs and OSs in the world today are made by Intel and Microsoft, respectively. These are, of course, two separate companies that set the prices of their products independently. The PowerPC, another popular CPU, was designed by a consortium consisting of IBM, Motorola, and Apple. Two commercial operating systems for the PowerPC are the Apple OS and IBM's AIX. In addition to these commercial operating systems, there are free systems like BSD and GNU-Linux that are provided by groups of programmers working on a volunteer basis.

Let us consider the pricing problem facing sellers of complementary products. The critical feature is that the demand for *either* product depends on the price of *both* products. If p_1 is the price of the CPU and p_2 is the price of the OS, the cost to the end user depends on $p_1 + p_2$. Of course, you need more than just a CPU and an OS to make a useful system, but that just adds more prices to the sum; we'll keep things simple by sticking with two components.

The demand for CPUs depends on the price of the total system, so we write $D(p_1 + p_2)$. If we let c_1 be the marginal cost of a CPU and F the fixed cost, the profit-maximization problem of the CPU maker can be written

$$\max_{p_1} (p_1 - c_1)D(p_1 + p_2) - F_1.$$

Similarly, the profit-maximization problem of the OS maker can be written

$$\max_{p_2} (p_2 - c_1)D(p_1 + p_2) - F_2.$$

In order to analyze this problem, let us assume that the demand function has the linear form

$$D(p) = a - bp.$$

Let us also assume, for simplicity, that the marginal costs are so small that they can be ignored. Then the CPU profit-maximization problem becomes

$$\max_{p_1} p_1[a - b(p_1 + p_2)] - F_1,$$

or

$$\max_{p_1} ap_1 - bp_1^2 - bp_1p_2 - F_1.$$

It turns out that the marginal revenue from a price increase Δp_1 is

$$(a - 2bp_1 - bp_2) \Delta p_1.$$

If profit is maximized, then the change in revenue from an increase in p_1 must be zero:

$$a - 2bp_1 - bp_2 = 0.$$

Solving this equation we have

$$p_1 = \frac{a - bp_2}{2b}.$$

In exactly the same way, we can solve for the profit-maximizing choice of the OS price:

$$p_2 = \frac{a - bp_1}{2b}.$$

Note that the optimal choice of each firm's price depends on what it expects the other firm to charge for its component. As usual, we are interested in a **Nash equilibrium**, where each firm's expectations about the other's behavior are satisfied.

Solving the system of two equations in two unknowns, we have

$$p_1 = p_2 = \frac{a}{3b}.$$

This gives us the profit-maximizing prices if each firm unilaterally and independently sets the price of its component of the system. The price of the total system is

$$p_1 + p_2 = \frac{2a}{3b}.$$

Now let us consider the following experiment. Suppose that the two firms merge to form an integrated firm. Instead of setting the prices of the components, the integrated firm sets the price of the final system, which we denote by p . Its profit-maximization problem is therefore

$$\max_p p(a - bp).$$

The marginal revenue from increasing the system price by Δp is

$$(a - 2bp) \Delta p.$$

Setting this equal to zero and solving, we find that the price that the integrated firm will set for the final system is

$$p = \frac{a}{2b}.$$

Note the following interesting fact: the profit-maximizing price set by the integrated firm is *less* than the profit-maximizing price set by the two independent firms. Since the price of the system is lower, consumers will buy more of them and be better off. Furthermore, the profits of the integrated firm are larger than the sum of the equilibrium profits of the two independent firms. Everyone has been made better off by coordinating the pricing decision!

This turns out to be true in general: a merger of two monopolies that produce complementary products results in lower prices and higher profits than if the two firms set their prices independently.²

The intuition is not hard to see. When firm 1 contemplates a price decrease for the CPU, it will increase demand for CPUs *and* OSs. But it only takes into account the impact on its own profit from cutting price, ignoring the profits that will accrue to the other firm. This leads it to cut prices less than it would if it were interested in maximizing joint profit. The same analysis applies to firm 2, leading to prices that are “too high” from the viewpoint of both profit-maximization and consumer surplus.

Relationships among Complementors

The “merger of complementors” analysis is provocative, but we shouldn’t immediately leap to the conclusion that mergers of OS and CPU manufacturers are a good idea. What the result says is that *independent* price setting will lead to prices that are too high from the viewpoint of joint profitability, but there are lots of intermediate cases between totally independent and fully integrated.

For example, one of the firms can negotiate prices for components and then sell an integrated bundle. This is, more or less, what Apple does. They buy PowerPC CPUs in bulk from Motorola, build them into computers, and then bundle the operating system and computers together for sale to the end customers.

Another model for dealing with the systems pricing problem is to use revenue sharing. Boeing builds airplane bodies and GE builds airplane engines. The end user generally wants both a body and an engine. If GE and Boeing each set their prices independently, they could decide to set their prices too high. So what they do instead is to negotiate a deal in which GE will receive a fraction of the revenue from the sale of the assembled aircraft. Then GE is happy to have Boeing negotiate to get as high a price as possible for the package, confident that it will receive its specified share.

² This rather remarkable fact was discovered by Augustin Cournot, whom we previously met in Chapter 27.

There are other mechanisms that work in different industries. Consider, for example, the DVD industry mentioned in the introduction. This has been a very successful new product, but making it work was tricky. Consumer electronics firms didn't want to produce players unless they were assured that there would be plenty of content available, and content providers didn't want to produce content unless they were sure that would be lots of DVD players out there.

On top of this, both the consumer electronics firms and the content producers would have to worry about the pricing of complements problem: if there were only a few providers of players and only a few providers of content, then they would each want to price their products "too high," reducing the total profit available in the industry and making consumers worse off.

Sony and Philips, who held the basic patents on the DVD technology, helped solve this problem by licensing the technology widely at attractive prices. They also realized that there had to be a lot of competition to keep the prices down and kick start the industry. They recognized that it was much better to have a small share of a large, successful industry than to have a large share of a nonexistent industry.

Yet another model for relationships among complementors might be called "commoditize the complement." Look back at firm 1's profit maximization problem:

$$\max_{p_1} p_1 D(p_1 + p_2) - F_1.$$

At any given configuration of prices, reducing p_1 may or may not increase firm 1's revenues, depending on the demand elasticity. But lowering p_2 will *always* increase firm 1's revenue. The challenge facing firm 1 is then: how can I get firm 2 to cut its price?

One way is to try to make competition for firm 2 more intense. Various strategies are possible here, depending on the nature of the industry. In technology-intensive industries, standardization becomes an important tool. An OS producer, for example, would want to encourage standardized hardware. This not only makes its job easier, but it also ensures that the hardware industry will be highly competitive. This will ensure that competitive forces push down the price of hardware and reduce the total system price to end users, thereby increasing the demand for operating systems.³

35.3 Lock-In

Since IT components often work together as systems, switching any one component often involves switching others as well. This means that the

³ See Brandenburger, Adam and Barry Nalebuff, *Co-opetition*, Doubleday, 1997 for further analysis of strategy for complementors.

switching costs associated with one component in IT industries may be quite substantial. For example, switching from a Macintosh to a Windows-based PC involves not only the hardware costs of the computer itself, but also involves purchasing of a whole new library of software, and, even more importantly, learning how to use a brand new system.

When switching costs are very high, users may find themselves experiencing **lock-in**, a situation where the cost of changing to a different system is so high that switching is virtually inconceivable. This is bad for the consumers, but is, of course, quite attractive for the seller of the components that make up the system in question. Since the locked-in user has a very *inelastic* demand, the seller(s) can jack up the prices of their components to extract consumer surplus from the user.

Of course, wary consumers will try to avoid such lock-in, or, at the very least, bargain hard to be compensated for being locked in. Even if the consumers themselves are poor at bargaining, competition among sellers of systems will force prices down for the *initial* purchase, since the locked-in consumers can provide them with a steady revenue stream afterwards.

Consider, for example, choosing an Internet service provider (ISP). Once you have committed to such a choice, it may be inconvenient to switch due to the cost of notifying all of your correspondents about your new e-mail address, reconfiguring your Internet access programs, and so on. The monopoly power due to these switching costs means that the ISP can charge more than the marginal cost of providing service, once it has acquired you as a customer. But the flip side of this effect is that the stream of profits of the locked-in customers is a valuable asset, and ISPs will compete up front to acquire such customers by offering discounts and other inducements to sign up with them.

A Model of Competition with Switching Costs

Let's examine a model of this phenomenon. We assume that the cost of providing a customer with Internet access is c per month. We also assume a perfectly competitive market, with many identical firms, so that in the absence of any switching costs, the price of Internet service would simply be $p = c$.

But now suppose that there is a cost s of switching ISPs and that ISPs can offer a discount of size d for the first month to attract new customers. At the start of a given month, a consumer contemplates switching to a new ISP. If he does so, he only has to pay the discounted price, $p - d$, but he also has to endure the switching costs s . If he stays with his old provider, he has to pay the price p forever. After the first month, we assume that both providers continue to charge the same price p forever.

The consumer will switch if the present value of the payments to the new provider plus the switching cost is less than the present value of the

payments to the original ISP. Letting r be the (monthly) interest rate, the consumer will switch if

$$(p - d) + \frac{p}{r} + s < p + \frac{p}{r}.$$

Competition between providers ensures that the consumer is indifferent between switching or not switching, which implies

$$(p - d) + s = p.$$

It follows that $d = s$, which means the discount offered just covers the switching cost of the consumer.

On the producer side, we suppose that competition forces the present value of profits to be zero. The present value of profit associated with a single customer is initial discount, plus the present value of the profits in future months. Letting r be the (monthly) interest rate, and using the fact that $d = s$, the zero-profit condition can be written as

$$(p - s) - c + \frac{p - c}{r} = 0. \quad (35.1)$$

Rearranging this equation gives us two equivalent ways to describe the equilibrium price:

$$p - c + \frac{p - c}{r} = s, \quad (35.2)$$

or

$$p = c + \frac{r}{1 + r} s. \quad (35.3)$$

Equation (35.2) says that the present value of the future profits from the consumer must just equal the consumer's switching cost. Equation (35.3) says that the price of service is a markup on marginal cost, with the amount of the markup is proportional to the switching costs.

Adding switching costs to the model raises the *monthly* price of service above cost, but competition for this profit flow forces the *initial* price down. Effectively, the producer is investing in the discount $d = s$ in order to acquire the flow of markups in the future.

In reality many ISPs have other sources of revenue than just the monthly income from their customers. America Online, for example, derives a substantial part of its operating revenue from advertising. It makes sense for them to offer large up-front discounts, in order to capture advertising revenue, even if they have to provide Internet connections at rates at or below cost.

We can easily add this effect to the model. If a is the advertising revenue generated by the consumer each month, the zero-profit condition requires

$$(p - s) + a - c + \frac{p + a - c}{r} = 0. \quad (35.4)$$

Solving for p we have

$$p = c - a + \frac{r}{1+r}s.$$

This equation shows that what is relevant is the *net* cost of servicing the customer, $c - a$, which involves both the service cost and the advertising revenues.

EXAMPLE: Online Bill Payment

Many banks offer low-cost or even free bill payment services. Some banks will even pay customers who start using their online bill payment services.

Why the big rush to pay bills online? The answer is that banks have found that once a customer goes to the trouble of setting up the bill-paying service, he or she is much less likely to switch banks. According to a Bank of America study, the frequency of switching goes down by 80 percent for such customers.⁴

It's true that once you get online bill payment up and running, it's hard to give it up. Switching to another bank to get an extra tenth of a percent of interest on your checking account doesn't seem very attractive. As in the analysis of lock-in presented above, investing in services that create switching costs can be very profitable for businesses.

EXAMPLE: Number Portability on Cell Phones

At one time, cell phone providers prevented individuals from transferring their phone numbers when they switched carriers. This prohibition increases individual switching costs significantly, since anyone who switched would have to notify all of his or her friends about the new number.

As the model presented in this chapter describes, the fact that customers could be charged more when they faced high switching costs meant that the phone providers would compete even more aggressively to sign up such highly profitable customers. This competition took the form of providing low-cost or even free phones, along with offers of "free minutes," "rollover plans," "cell-to-cell discounts," and other marketing gimmicks.

The cell phone industry was united in its efforts to block number portability and lobbied regulatory agencies and Congress to maintain the status quo.

Slowly but surely, the tide started to turn against the cell phone industry as consumers demanded number portability. The Federal Communications Commission, which regulates the telephone business, started dropping hints

⁴ Michelle Higgins, "Banks Use Online Bill Payment In Effort to Lock In Customers," *Wall Street Journal*, September 4, 2002.

that cell phone providers should consider ways in which they could implement number portability.

In June 2003, Verizon Wireless said it would drop opposition to number portability. Their decision appeared to rest on two considerations. First, it was becoming clear that they were fighting a losing battle: eventually cell-number portability would win out. Perhaps more significantly, several recent consumer surveys showed that Verizon led the industry in terms of customer satisfaction. It appeared quite possible that Verizon would gain more customers than it lost if switching costs were reduced. Indeed, it appears that ultimately Verizon benefited from number portability.

This episode provides a good lesson in business strategy: tactics to increase customer switching costs may be valuable for a while. But ultimately service quality plays a decisive role in attracting and retaining customers.

35.4 Network Externalities

We have already examined the idea of **externalities** in Chapter 34. Recall that economists use this term to describe situations in which one person's consumption directly influences another person's utility. **Network externalities** are a special kind of externalities in which one person's utility for a good depends on the *number* of other people who consume this good.⁵

Take for example a consumer's demand for a fax machine. People want fax machines so they can communicate with each other. If no one else has a fax machine, it certainly isn't worthwhile for you to buy one. Modems have a similar property: a modem is only useful if there is another modem somewhere that you can communicate with.

Another more indirect effect for network externalities arises with complementary goods. There is no reason for a video store to locate in a community where no one owns a video player; but then again, there is little reason to buy a video player unless you have access to pre-recorded video tapes to play in the machine. In this case the demand for video tapes depends on the number of VCRs, and the demand for VCRs depends on the number of video tapes available, resulting in a slightly more general form of network externalities.

35.5 Markets with Network Externalities

Let us try to model network externalities using a simple demand and supply model. Suppose that there are 1000 people in a market for some good and we index the people by $v = 1, \dots, 1000$. Think of v as measuring the

⁵ More generally, a person's utility could depend on the *identity* of other users; it is easy to add this to the analysis.

reservation price for the good by person v . Then if the price of the good is p , the number of people who think that the good is worth at least p is $1000 - p$. For example, if the price of the good is \$200, then there are 800 people who are willing to pay at least \$200 for the good, so the total number of units sold would be 800. This structure generates a standard, downward-sloping demand curve.

But now let's add a twist to the model. Suppose that the good we are examining exhibits network externalities, like a fax machine or a telephone. For simplicity, let us suppose that the value of the good to person v is vn , where n is the number of people who consume the good—the number of people who are connected to the network. The more people there are who consume the good, the more *each* person is willing to pay to acquire it.⁶ What does the demand function look like for this model?

If the price is p , there is someone who is just indifferent between buying the good and not buying it. Let $\hat{v}$ denote the index of this marginal individual. By definition, he is just indifferent to purchasing the good, so his willingness to pay for the good equals its price:

$$p = \hat{v}n. \quad (35.5)$$

Since this “marginal person” is indifferent, everyone with a *higher* value of v than $\hat{v}$ must definitely want to buy. This means that the number of people who want to buy the good is

$$n = 1000 - \hat{v}. \quad (35.6)$$


Putting equations (35.5) and (35.6) together, we have a condition that characterizes equilibrium in this market:

$$p = n(1000 - n).$$

This equation gives us a relationship between the price of the good and the number of users. In this sense, it is a kind of demand curve; if there are n people who purchase the good, then the willingness to pay of the marginal individual is given by the height of the curve.

However, if we look at the plot of this curve in Figure 35.1, we see that it has quite a different shape than a standard demand curve! If the number of people who connect is low, then the willingness to pay of the marginal individual is low, because there aren't many other people out there that he can communicate with. If there are a large number of people connected, then the willingness to pay of the marginal individual is low, because everyone else who valued it more highly has already connected. These two forces lead to the humped shape depicted in Figure 35.1.

⁶ We should really interpret n as the number of people who are *expected* to consume the good, but this distinction won't be very important for what follows.



Network externalities. The demand is given by the curved hump, the supply by the horizontal line. Note that there are three intersections where demand equals supply.

Now that we understand the demand side of the market, let's look at the supply side. To keep things simple, let us suppose that the good can be provided by a constant returns to scale technology. As we've seen, this means that the supply curve is a flat line at price equals average cost.

Note that there are three possible intersections of the demand and supply curve. There is a low-level equilibrium where $n^* = 0$. This is where no one consumes the good (connects to the network), so no one is willing to pay anything to consume the good. This might be referred to as a "pessimistic expectations" equilibrium.

The middle equilibrium with a positive but small number of consumers is one where people don't think the network will be very big, so they aren't willing to pay that much to connect to it—and therefore the network isn't very big.

Finally the last equilibrium has a large number of people, n_H . Here the price is small because the marginal person who purchases the good doesn't value it very highly, even though the market is very large.


35.6 Market Dynamics

Which of the three equilibria will we see occur? So far the model gives us no reason to choose among them. At each of these equilibria, demand equals supply. However, we can add a dynamic adjustment process to help us decide which equilibrium is more likely to occur.

It is plausible to assume that when people are willing to pay more than the cost of the good, the size of the market expands and, when they are willing to pay less, the market contracts. Geometrically this is saying that when the demand curve is above the supply curve, the quantity goes up and, when it is beneath the supply curve, the quantity goes down. The arrows in Figure 35.1 illustrate this adjustment process.

These dynamics give us a little more information. It is now evident that the low-level equilibrium, where no one connects, and the high-level equilibrium, where many people connect, are stable whereas the middle equilibrium is unstable. Hence it is unlikely that the final resting point of the system will be the middle equilibrium.

We are now left with two possible stable equilibria; how can we tell which is likely to occur? One idea is to think about how costs might change over time. For the kinds of examples we have discussed—faxes, VCRs, computer networks, and so on—it is natural to suppose that the cost of the good starts out high and then decreases over time due to technological progress. This process is illustrated in Figure 35.2. At a high unit cost there is only one stable equilibrium—where demand equals zero. When the cost decreases sufficiently, there are two stable equilibria.




Cost adjustment and network externalities. When the cost is high, the only equilibrium implies a market of size zero. As the cost goes down, other equilibria become possible.

Now add some noise to the system. Think of perturbing the number of people connected to the network around the equilibrium point of $n^* = 0$. These perturbations could be random, or they could be part of business strategies such as initial discounts or other promotions. As the cost gets

smaller and smaller, it becomes increasingly likely that one of these perturbations will kick the system up *past* the unstable equilibrium. When this happens, the dynamic adjustment will push the system up to the high-level equilibrium.

A possible path for the number of consumers of the good is depicted in Figure 35.3.


It starts out at essentially zero, with a few small perturbations over time. The cost decreases, and at some point we reach a critical mass that kicks us up past the low-level equilibrium and the system then zooms up to the high-level equilibrium.



A real-life example of this kind of adjustment is the market for fax machines. Figure 35.4 illustrates the price and number of fax machines shipped over a period of 12 years.⁷

⁷ This diagram is taken from "Critical Mass and Network Size with Applications to the US Fax Market," by Nicholas Economides and Charles Himmelberg (Discussion Paper no. EC-95-11, Stern School of Business, N.Y.U., 1995). See also Michael L. Katz and Carl Shapiro, "Systems Competition and Network Effects," *Journal of Economic Perspectives*, 8 (1994), 93-116, for a nice overview of network externalities and their implications.

7 This diagram is taken from "Critical Mass and Network Size with Applications to the US Fax Market," by Nicholas Economides and Charles Himmelberg (Discussion Paper 1005, MIT, 1995). See also Michael



Fax market. The demand for fax machines was small for a long time since so few people used them. During the mid-eighties the price fell significantly and the demand suddenly exploded.

EXAMPLE: Network Externalities in Computer Software

Network externalities arise naturally in the provision of computer software. It is very convenient to be able to exchange data files and tips with other users of the same software. This gives a significant advantage to the largest seller in a given market and leads software producers to invest heavily in acquiring market share.

Examples of this abound. Adobe Systems for example, invested heavily in developing a "page description language" called PostScript for desktop publishing. Adobe realized clearly that no one would invest the time and

resources necessary to learn PostScript unless it was the clear “industry standard.” So the firm deliberately allowed competitors to “clone” its language in order to create a competitive market in PostScript interpreters. Adobe’s strategy paid off: several competitors emerged (including one that gave its product away) and PostScript became a widely used standard for desktop publishing. Adobe kept a few things proprietary—for instance, techniques for displaying fonts at low resolution—and managed to dominate the high end of the market. Ironically, Adobe’s market success was due to its ability to encourage entry by its competitors!

In recent years, many software producers have followed this model. Adobe itself gives away several software products, such as the Adobe Acrobat reader. One of the hot new stock issues of 1995, Netscape Communications Corporation, acquired the lion’s share of the Web browser market by giving away its main product, making it a prime example of a company that “lost money on every sale, but made up for it in volume.”

35.7 Implications of Network Externalities

The model described above, simple though it is, still yields a number of insights. For example, the critical mass issue is very important: if one user’s demand depends on how many other users there are, it is very important to try to stimulate growth early in the life cycle of a product. Nowadays it is quite common to see producers offering very cheap access to a piece of software or a communications service in order to “create a market” where none existed before.

Of course, the critical question is how big does the market have to be before it can take off on its own? Theory can provide little guidance here; everything depends on the nature of the good and the costs and benefits the users face in adopting it.

Another important implication of network externalities is the role played by governmental policy. The Internet is a prime example. The Internet was originally used only by a few small research labs to exchange data files. In the mid-eighties the National Science Foundation used the Internet technology to connect several large universities to 12 supercomputers deployed at various locations. The original vision was that researchers at the universities would send data back and forth to the supercomputers. But a fundamental property of communications networks is that if you are all connected to the same thing, you are all connected to each other. This allowed researchers to send email to each other that had nothing to do with the supercomputers. Once a critical mass of users had been connected to the Internet, its value to new users increased dramatically. Most of these new users had no interest in the supercomputer centers, even though this was the original motivation for providing the network.

EXAMPLE: The Yellow Pages

The familiar local yellow pages phone directories are a \$14 billion business. Ten years ago, it was dominated by telephone companies, who had about 95 percent of the market. Nowadays, they have only 85 percent.

The difference is due to competition. Several small upstarts entered the market in recent years, taking business away from the local phone companies. This is no easy task, as the local business directories exhibit a classic form of network effects: it used to be that consumers all used the yellow page directory provided by their local phone companies, so local merchants were forced to advertise in them.

One upstart, Yellow Book, managed to overcome the network effects by using clever business strategies, such as dramatically undercutting the phone companies' ad rates and distributing its directory just before the local phone company's directory came out. The incumbent providers, thinking that their market was secure, dismissed the threat of aggressive newcomers until it was nearly too late. In the last few years, competition has heated up in this industry. This example goes to show that even industries with strong network effects aren't immune to competitive forces, particularly when the incumbents become overconfident.

35.8 Rights Management

There is much interest these days in new business models for intellectual property (IP). IP transactions take a variety of forms: books are sold outright and also borrowed from libraries. Videos can either be sold or rented. Some software is licensed for particular uses; other software is sold outright. Shareware is a form of software in which payment is voluntary.

Choosing the terms and conditions under which a piece of intellectual property is offered is a critical business decision. Should you use copy protection? Should you encourage users to share a news item with a friend? Should you sell to individuals or use site license?

Some simple economics helps to understand the relevant issues. Let's consider a purely digital good, such as an online newspaper, so we don't have to worry about marginal cost of production. First let us consider behavior under some default set of terms and conditions. The owner of the digital good will choose a price and, implicitly, a quantity to sell so as to maximize profit:

$$\max_y p(y)y \quad (35.7)$$

This yields some optimal (p^*, y^*) .

Now the seller of the good contemplates liberalizing terms and conditions: let's say extending a trial period of free use from 1 week to 1 month. This

has two effects on the demand curve. First, it increases the value of the product to each of the potential users, shifting the demand curve up. But it also may easily result in less of the item being sold, since some users will find the longer trial period enough to meet their needs.

Let us model this by defining the new amount consumed by $Y = by$, where $b > 1$, and the new demand curve by $P(Y) = ap(y)$, where $a > 1$. The new profit-maximization problem now becomes

$$\max_Y P(Y)y.$$

Note that we multiply price times the amount sold, y , not the amount consumed, Y .

Applying the definitions $Y = by$ and $P(Y) = ap(y)$, we can write this as

$$\max_Y ap(y) \frac{Y}{b} = \max_Y \frac{a}{b} p(Y)Y.$$

This maximization problem looks like problem (35.7) except for the constant a/b in front of the max. This will not affect the optimal choice, so we can conclude that $Y^* = y^*$.

This simple analysis allows us to make several conclusions:

- The amount of the good consumed, Y^* , is independent of the terms and conditions.
- The amount of the good produced is y^*/b which is less than y^* .
- The profits could go up or down depending on whether a/b is greater or less than 1. Profits go up if the increase in value to the consumers who buy the product compensates for the reduced number of buyers.

EXAMPLE: Video Rental

Video stores can choose the terms and conditions under which they rent videos. The longer you can keep the video, the more valuable it is to you, since you have a longer period of time during which you can watch it. But the longer you keep the video, the less profit the store makes from it, since it is unable to rent it to someone else. The optimal choice for the rental period involves trading off these two effects.

In practice, this has tended to lead to a form of product differentiation. New releases are rented for short periods, since the the profits from other renters being excluded are very substantial. Older videos are rented for longer periods, since there is less cost to the store from the video being unavailable.

35.9 Sharing Intellectual Property

Intellectual property is often shared. Libraries, for example, facilitate the sharing of books. Video stores help people to “share” videos—and charge a price for doing so. Interlibrary loan helps libraries share books among themselves. Even textbooks—such as the one you are holding—are shared among students from one term to the next via the resale market.

There is considerable debate in the publishing and library communities about the proper role of sharing. Librarians have established an informal “rule of five” for interlibrary loan: an item may be loaned out up to five times before additional royalty payments should be made to the publisher. Publishers and authors have traditionally been unenthusiastic about the resale market for books.

The advent of digital information has made this situation even more acute. Digital information can be perfectly reproduced, and “sharing” can be taken to new extremes. Recently, a well-known country music singer engaged in a vociferous public relations campaign against stores selling used CDs. The problem was that CDs do not deteriorate with replay and it is possible to buy a CD, tape it, and then sell the CD to the used-CD store.

Let us try to construct a model of this sort of sharing phenomenon. We begin with the baseline case in which there is no sharing. In this case a video maker chooses to produce y copies of a video to maximize profit:

$$\max_y p(y)y - cy - F. \quad (35.8)$$

As usual, $p(y)$ is the inverse demand function, c is the (constant) marginal cost, and F is the fixed cost. Let the profit maximizing output be denoted by y_n , where the n stands for “no sharing.”

Now suppose that a video rental market is allowed. In this case the number of videos *viewed* will be distinct from the number of copies produced. If y is the number of videos produced and each video is shared among k viewers, then the number of viewings will be $x = ky$. (For simplicity we are assuming that *all* copies of the video are rented in this case.)

We need to specify how the consumers sort themselves into the “clubs” that share the videos. The simplest assumption is that the consumers with high values associate with each other, and the consumers with low-values associate with each other. That is, one club consists of consumers with the k highest values, another club consists of the consumers with the next k highest values, and so on. (Other assumptions could be used, but this one gives a very simple analysis.)

If y copies of the video are produced, $x = ky$ copies will be viewed, so the willingness to pay of the marginal individual will be $p(x) = p(ky)$. However, it is clearly the case that there is some inconvenience cost to renting a video

rather than owning it yourself. Let us denote this “transactions cost” by t , so that the willingness to pay of the marginal individual becomes $p(x) - t$.

Recall that we have assumed that all copies of the video are shared among k users. Therefore the willingness to pay of a *video store* will just be k times the willingness to pay of the marginal individual. That is, if y copies are produced, the willingness to pay of the video store will be

$$P(y) = k[p(ky) - t]. \quad (35.9)$$

Equation (35.9) contains the two key effects that arise from sharing: the willingness to pay goes *down* since more videos are viewed than are produced; but the willingness to pay also goes *up* since the cost of a single video is shared among several individuals.

The profit maximization problem of the producer now becomes

$$\max_y P(y)y - cy - F,$$

which can be written as

$$\max_y k[p(ky) - t]y - cy - F,$$

or

$$\max_y p(ky)ky - \left(\frac{c}{k} + t\right)ky - F.$$

Recalling that the number of viewings, x , is related to the number produced, y , via $x = ky$, we can also write the maximization problem as

$$\max_x p(x)x - \left(\frac{c}{k} + t\right)x - F.$$

Note that this problem is identical to problem (35.8), with the exception that the marginal cost is now $(c/k + t)$ rather than c .

The close relationship between the two problems is very useful since it allows us to make the following observation: *profits will be larger when rental is possible than when it is not if and only if*

$$\frac{c}{k} + t < c.$$

Rearranging this condition, we have

$$\left(\frac{k}{k+1}\right)t < c.$$

For large k , the fraction on the left is about 1. Hence the critical issue is the relationship between the marginal cost of production, c , and the transactions cost of renting, t .

If the cost of production is large and the cost of renting is small, then the most profitable thing for a producer to do is to produce a few copies, sell them at a high price, and let the consumers rent. On the other hand, if the transactions cost of renting is larger than the cost of production, it is more profitable for a producer to have renting prohibited: since renting is so inconvenient for the consumers, video stores aren't willing to pay much more for the “shared” videos, and so the producer is better off selling.

Summary

1. Because information technology works together in systems, it is costly to consumers to switch any one component.
2. If two monopoly providers of complementary products coordinate their price setting, then they will both set their prices lower than they would than if they set them independently.
3. This will increase profit for the two monopolists *and* make consumers better off.
4. There are many ways to achieve this coordination, including merger, negotiation, revenue sharing, and commoditization.
5. In a lock-in equilibrium the discount offered first period is paid for by increased prices in future periods.
6. Network externalities arise when one person's willingness to pay for a good depends on the number of other users of that good.
7. Models with network externalities typically exhibit multiple equilibria. The ultimate outcome often depends on the history of the industry.
8. Rights management involves a tradeoff between increased value and prices versus reduced sales.
9. Information goods like books and videos are often rented or shared as well as purchased. Rental or purchase can be more profitable depending on how transactions costs compare with production costs.

REVIEW QUESTIONS

1. If the cost to a customer from switching long-distance carriers is on the order of \$50, how much should a long-distance carrier be willing to pay to acquire a new customer?
2. Describe how the demand for a word processing package might exhibit network externalities.
3. Suppose that the marginal cost of producing an extra video is zero and the transactions cost of renting a video is zero. Does a producer make more money by selling the video or by renting it?