

10

Some Lunches Are Free

“Oh, ‘tanstaaf.’ Means ‘There ain’t no such thing as a free lunch.’ And isn’t,” I added, pointing to a FREE LUNCH sign across room, “or these drinks would cost half as much. Was reminding her that anything free costs twice as much in the long run or turns out worthless.”

—Robert Heinlein¹

INSIDE THIS CHAPTER

- The basics of software platform pricing
- When to use fixed versus access fees
- Where software platforms have sought profit

When you go to a video game arcade, you have to pay every time you start a game. At home, once you have bought a game for your video game console, you can play it as often as you like. And if you play online, you can also play as often as you like, but only if you pay a monthly fee. When you play music on your computer, you can either use the media player that came with it or download others for free and use them instead. RealNetworks sells its digital content by subscription; Apple charges 99 cents per song. Computer manufacturers can pay Microsoft for the right to install Windows (with larger manufacturers receiving a volume discount) or they can install Linux for free. Software developers can either pay Sony a royalty for each copy of a PlayStation game they sell or they can get Apple’s help to write games for the Macintosh and pay Apple no royalties at all. Mobile phone users in the United States

1. Robert A. Heinlein, *The Moon Is a Harsh Mistress* (New York: Tom Doherty Associates, 1997).

are generally charged according to the number of minutes they are connected to the network, with various volume discounts available, while i-mode users in Japan pay according to the amount of data they send or receive, regardless of how long they are connected.

This chapter tries to explain why different industries have chosen different pricing methods to get both sides on board and maximize profits from the participants in their ecosystems. In so doing it provides insights into the pricing policies that new businesses based on software platforms should consider. We begin by explaining the basic economics of pricing in multisided platform businesses. We then consider three important dimensions of pricing for software platforms. First, what is priced? Some businesses based on software platforms charge for access to their platform, some charge for use of the platform, and a few charge for both. Second, how do prices vary across the customers on a given side of a platform (game developers, for example)? One price does not fit all in either theory or practice. Third, who gets the (nearly) free lunch? Contrary to Milton Friedman's observation that there's no such thing as a free lunch, in the industries we have examined, the customers on one side get services for free, or at least for a price that at best covers out-of-pocket cost.

Pricing in Multisided Businesses

In single-sided businesses, pricing analysis mainly focuses on the level of price, both at introduction for new products and at maturity, and on price discrimination—differences in price paid by different customers. The profit-maximizing price for a single product depends on the cost of supplying an additional unit—the incremental or marginal cost—and on the responsiveness of demand to price. Demand will be more responsive to the price of a particular product—economists would say the price elasticity of demand will be higher—the easier it is for buyers to reduce purchases of the product in question when its price rises, either by switching to competitive products or by doing without, and the less likely competitors are to match price increases or ignore price cuts. The higher a product's price elasticity of demand, the lower is the optimal markup of that product's price over marginal cost.

Pricing is more complex for new products and for firms that produce multiple products. Firms sometimes engage in “penetration pricing”: they charge low introductory prices to get the attention of buyers and penetrate a new market, then raise these prices over time as the market matures. Under other conditions firms practice “cream skimming,” charging high prices to early, eager buyers and then lowering prices over time to capture a larger market. Similarly, firms that produce several products must adjust prices to take into account that some may be substitutes for each other (pricing high reduces cannibalization of sales from substitutes) and some may be complements for each other (pricing low boosts the sales of complementary products).

Even in these and other more complex cases, pricing in single-sided markets always begins with product-specific marginal cost. This tight connection between the incremental cost of a product and its price weakens considerably in multisided businesses in ways that have important implications for pricing strategy.

The fact that there is a tight connection between prices and costs in single-sided businesses doesn’t mean that all customers get charged the same price. In fact, many firms charge higher markups over cost to some customers (or groups of customers) than others, depending on their intensity of demand. Economists gave this practice the name “price discrimination” before “discrimination” acquired such negative connotations.² Not only is price discrimination common in market economies, but often it enhances economic welfare by, for example, better enabling firms to recover the costs of research and development and thus increasing incentives to perform R&D.

Price discrimination turns out to be important in multisided businesses as well, as we discuss later in this chapter. As background for these discussions, it is helpful to summarize the three major types of price discrimination that economists have identified.

Firms would like to be able to charge every individual buyer the absolute most they would be willing to pay. That is what a very good used car salesperson tries to do: he tries to figure out the most each buyer

2. Mankiw, “Principles of Economics,” p. 334; Carlton and Perloff, “Modern Industrial Organization,” pp. 297–299.

would pay and then quotes just a little bit less. (Economists would say he is practicing first-degree price discrimination.) In practice, sellers rarely have enough information to pull this off, so they try a couple of cruder methods.

One common method involves various bulk or volume discount schemes that give customers a lower average cost per unit the more units they buy. (Economists call this second-degree price discrimination, or nonlinear pricing.) The simplest bulk discount is called a “two-part tariff,” because customers have to pay both a fixed charge to buy anything (an access charge) and a per-unit charge for each unit purchased (a usage charge). The average cost per unit falls as more units are purchased.³ In principle, firms can generally increase profits by using both access and usage charges, though collecting both charges sometimes entails higher transactions costs that swamp any potential profit gain. If there are costs associated both with providing access and with providing usage, it seems more common for firms to make the bulk of their profits from usage, which measures strength of demand, rather than from access. That is, to use the most familiar example, giving away the razor (or selling it at cost) and making money on the blades seems to be more common than selling the blades at cost and making profits on the razor.

Another method of price discrimination reflects the fact that it is often much easier for firms to guess how much, on average, a particular group would be willing to pay than to guess how much a particular individual might pay. An airline may not know how much a particular individual will pay for a seat, but it does know that people who travel at the last minute are typically business travelers who are willing to pay much more than people who book weeks in advance. The airline can therefore charge the group of last-minute travelers more than advance-booking passengers. Of course, to engage in this group method of pricing, the seller must find groups of buyers who have different sensitivity to the product’s price, who can be identified and therefore charged separate prices, and

3. The classic exposition is Walter Oi’s “A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly,” *Quarterly Journal of Economics* 85 (1971): 77–96. For a more technical treatment, see Richard Schmalensee, “Monopolistic Two-Part Pricing Arrangements,” *Bell Journal of Economics* 11 (Autumn 1981): 445–466.

who can't defeat the scheme by having the low-price group sell to the high-price group. Although there are challenges here, many businesses in practice figure out ways to engage in this sort of discrimination. (This is known as third-degree price discrimination.)

Pricing for Balance . . . ?

We have emphasized that in theory and in practice, multisided firms have to balance the demand on the two sides by pricing, and that they often do this by pricing low to one side and high to another. Any student of basic economics, though, should question this statement for any platform in which the two sides interact directly with each other. After all, under textbook competition it doesn't matter whether the government imposes a product-specific tax (2 cents per bushel of wheat, for instance) on buyers or on sellers, since sellers will pass the full amount of the tax through to the buyers. Now consider payment cards. Card systems typically charge cardholders a zero price for using their cards and charge merchants about 2 percent of the transaction amount when cardholders use their cards to pay. This effort to tilt pricing in favor of cardholders would be defeated if merchants imposed a surcharge of 2 percent on all transactions using a card. Then the cardholder would end up paying, and the merchant wouldn't.

Could a similar result be true for video game consoles? Does it matter whether the console maker collects royalties from game developers or from game users? If game developers pass the royalty cost on to the users, the users end up paying in the end anyway.

In practice, it generally does matter which side pays, because two key assumptions made in the textbook discussion don't apply. First, there are often significant transactions costs that prevent the customers on the two sides of most markets from just "sorting it out" themselves. Take the payment card example. Although most card systems prohibit merchant surcharging because it degrades the value of their product to cardholders, several countries have barred card systems from imposing such a no-surcharge rule. In those countries, however, most merchants don't surcharge. One reason is that it is costly to impose small charges on customers. Those merchants that do surcharge often charge more than they are charged by the card system—an indication that they are using the

fact that a customer wants to use her card as a basis for groupwise price discrimination.

The effects of transactions costs are visible in many two-sided markets. We discuss their likely role in video games below.

Second, competition in most real markets is less intense than in textbook markets. Competition among suppliers on at least one side of two-sided businesses is often imperfect, either because there are only a few major sellers or because products are differentiated. In this case, per-unit charges on that side generally are not passed on dollar for dollar. The exercise of market power (which leads to output restrictions) on that side may complicate the problem of balancing the two sides, and the presence of excess profits (deriving either from market power or from a few highly successful differentiated products) makes it more attractive to charge that side.⁴

When balance matters in a mature two-sided business, the pricing problem is much more complex than in a single-sided business. Marginal cost and price responsiveness on both sides matter for both prices, and so does the pattern of indirect network effects. In general, if side *A* cares more about side *B* than *B* cares about *A*, then, all else equal, *A* will contribute more total revenue. Thus, newspapers make their money from selling advertising, not from selling papers.

The textbook pricing formula for a single-sided market gives the optimal markup over marginal cost as 1 over a measure of price responsiveness (the price elasticity of demand), so low price responsiveness implies high markups. The corresponding formula for a two-sided business involves marginal costs on both sides, price responsiveness on both sides, and measures of the strength of indirect network effects in both directions. In particular, balance may require charging a price below marginal cost to a group with low price responsiveness, something a single-sided business would never do, if it is critical to attract members of that group in order to get members of the other group on board.

4. For a general discussion of the issues discussed in this paragraph and the remainder of this section, see Jean-Charles Rochet and Jean Tirole, "Two-Sided Markets: A Progress Report," mimeo, IDEI and GREMAQ, Toulouse, France. On the effects of imperfect competition, see Andrei Hagiu, "Pricing and Commitment by Two-Sided Platforms," *Rand Journal of Economics*, 37 (2006): forthcoming.

As we mentioned in Chapter 3, most two-sided businesses earn all or almost all of their profits from only one of the customer groups they serve. The standard economic theory of pricing in these businesses indicates that such pricing structures may be optimal, but it does not imply that they should be the norm. One explanation for the observed pattern is that sensitivity to price typically differs substantially between the two sides, so that it is optimal to price low to the price-sensitive side in order to attract the price-insensitive side, which can then serve as the main source of profit.⁵ Another explanation is that the standard theory neglects the transactions costs of collecting revenue from two sides rather than one. If, for instance, standard theory says that an 80/20 revenue split between two sides is optimal, but the costs of monitoring usage and excluding nonpayers required to collect the 20 would be significant in practice, the true optimum, taking those costs into account, may be 100/0.

Entry and Platform Competition

Two other strands of the economic literature would seem to be relevant to software platforms. The first deals with entry strategies. Recent work has argued that it may be necessary for new two-sided businesses to use a “divide and conquer” pricing strategy to deal with the chicken-and-egg problem, or the reluctance of either side to come on board without the other. The idea is initially to subsidize one side (or, more generally, to do whatever it takes) in order to get it on board even though the other side is not yet on board, and to use the presence of the subsidized side to attract the other side.⁶ This differs from the single-sided penetration pricing strategy discussed above because the key here is to generate indirect network effects, to use the subsidized side as a magnet to attract the other side. After entry has been successfully effected and both sides are on board, of course, the rationale for the initial subsidy vanishes, and

5. The technical argument that “corner solutions” of the 100/0 sort are the rule under the most plausible assumptions is given by Wilko Bolt and Alexander F. Tieman, “Skewed Pricing in Two-Sided Markets: An IO Approach” (working paper 13, De Nederlandsche Bank, Amsterdam, October 2004).

6. Bernard Caillaud and Bruno Jullien, “Chicken and Egg: Competition Among Intermediation Service Providers,” *Rand Journal of Economics* 34, no. 2 (Summer 2003): 521–552.

one would expect to see a corresponding shift in pricing policy. One of the regularities we discuss below, however, is that pricing structures—the relative amounts paid by the various sides—appear fairly robust over time; there are not many examples of pricing low to one side at first and then raising prices significantly later.

A slightly different entry problem arises when members of one side must be attracted before members of the other side.⁷ In order to attract buyers for a new video game console, for instance, an array of attractive games must be available at the console's launch, but this won't happen unless developers have been persuaded earlier to invest in developing those games. Developers, of course, won't make those investments unless they expect the console to be popular. This requires at least that they expect the console to be sold for a low price. In order to create such expectations, console makers often commit publicly to low prices months before their products are launched, in announcements directed at both game developers and end users. Steve Race, then president of Sony Computer Entertainment, describes such an announcement he made at a large trade show six months before the launch of Sony's Playstation⁸:

Olaf [Olafsson, President of Sony Electronic Publishing] was about two-thirds of the way through his speech when he said, 'I would like to call up Steve Race to tell you a little bit more about the Sony Playstation.' So I walked up. I had a whole bunch of sheets of paper in my hands, and I walked up, put them down on the podium, and I just said '\$299,' and I walked off to this thunderous applause.

The other relevant strand of the economic literature considers competition among multisided platform businesses. At one level, the standard pricing formula mentioned above deals with this: as in single-sided markets, the presence and behavior of competitors are important determinants of price responsiveness of demand. A new element here is the distinction between "single-homing" and "multihoming." When faced with two or more competing platforms, a business or household is said to single-home if (because of switching costs or for other reasons) it can deal with at most one of them; it is said to multihome if it is able to deal

7. Andrei Hagiu, "Pricing and Commitment by Two-Sided Platforms," *Rand Journal of Economics* 37 (2006): forthcoming.

8. Steven L. Kent, *The Ultimate History of Video Games*, 2001 p. 516.

with two or more of them. In the PC world, most households single-home, while many software developers multihome.

This is a fairly general pattern: most members of one side single-home and most members of the other multihome. While it seems plausible that this difference should affect pricing, it is less clear which side should benefit in general. The standard theory says that pricing on one side will tend to be lower, all else equal, when the number of single-homing members increases on that side and higher when the number of multihoming members increases.⁹ The argument is that if members on one side become more inclined to single-homing—which happens, for example, if their switching costs become higher—then competition will be more intense on that side, since it becomes competition for all of a member’s business, not just for some of it. On the other hand, casual observation of the video game industry suggests that as multihoming has become more common among video game developers over time, their royalty rates have come down substantially. This is consistent with the opposed argument that the easier it is for an important player on one side to multihome, the lower its switching costs and thus the greater its ability to shift its business between competing platforms. This in turn enhances its bargaining power vis-à-vis platforms and thus its ability to command lower prices. Perhaps in part for this reason, as we discuss below for most software platforms, end users, who generally single-home, contribute much more to the net revenue of the platform business than application or content developers, who commonly multihome.

What Is Priced?

Most of the preceding section implicitly assumed a platform business that charged only usage fees. In fact, an important choice in the industries studied here is between access fees and usage fees, which can be exemplified as the difference between buying a video game for home use, and thereby getting the right to play it as often as you like, versus being

9. Jean-Charles Rochet and Jean Tirole, “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association* 1, no. 4 (2003): 990–1029.

charged each time you play a game in an arcade. Even though, as we discussed above, it is theoretically preferable to employ both kinds of fees, we know of only one case in which this is done—most plausibly because of transactions costs. That case is massive, multiplayer online role-playing games (MMPRPGs to gamers), where players do face a two-part pricing regime—and more.

Price Discrimination in MMPRPGs

Online gaming has spawned several new and original pricing business models, which have first appeared on the PC platform and are now increasingly emulated by console manufacturers and console game publishers.¹⁰

Most of these novel pricing models have been created by the developers of online massive multi-player role-playing games (MMPRPG), which are hosted on the publishers' servers and played online simultaneously by thousands of users who enter and exit the game 24 hours a day, 7 days a week.¹¹ The basic pricing is a simple two-part tariff: users pay a fixed fee to buy the game CD, usually around \$50, after which they are charged monthly subscription rates. (For example, the current monthly fee for Sony Online Entertainment's Everquest is \$12.99, and the one for Electronic Arts' The Sims Online is \$10.) However, the very successful MMPRPGs realized that they could profitably sell expansion packets and game enhancements.¹² Everquest, for example, started offering a premium server, Everquest Legends: for an extra \$30 per month, players gain access to additional content, guidelines, and events. Other MMPRPGs have pushed price discrimination even further by selling additional game characters and objects. Players of Electronic Arts' Ultima Online can get advanced characters (alchemists and magicians) for \$30. It was not long before a secondary market appeared on eBay, where players trade characters among themselves.¹³

Noticing these developments, console makers realized the revenue-generating potential they offered and sought to capture it. Xbox 360, with

10. This has a lot to do of course with the fact that PC online gaming has largely predated console online gaming, which has only become a key aspect of platform competition with the last two console generations.

11. IGDA Online Games white paper (2003).

12. A hit like Everquest reached 500,000 users, whereas more recently, Korea's NCSOFT shattered all records when it announced that its Lineage MMPRPG had an astounding 4 million paying customers.

13. "Patti Waldmeir: Cyber World Is Heading for Regulation," *Financial Times*, March 30, 2005.

(continued)

its centralized and proprietary online service, is in a particularly good position to create an online marketplace where players can trade game artifacts (levels, characters, weapons, and so on). PlayStation will find it more difficult to do this, since its online service is decentralized.

These practices open up the very interesting possibility of “piece-mealing” games, transforming them from unitary packaged goods into completely modular, mix-and-match collections of products. It is easy to imagine publishers selling bare-bones versions of their games on CDs and then price discriminating among users by offering additional levels, characters, features, weapons, and so on for sale individually. This would present console manufacturers with the opportunity of charging users each time they downloaded a game piece through their online services. This would be broadly akin to the per-data packet charges levied by mobile network operators such as Japan’s NTT DoCoMo on users downloading content on their mobile phones through the wireless network.

For video game consoles that connect to the Internet, one could imagine emulating MMRPG pricing and charging both an access fee (the purchase price of the console) and a variable (per-month or per-game) fee for games played at home on the console. In theory, using both forms of pricing would generally increase profits somewhat. In reality, collecting a usage fee for console-based games would certainly increase the seller’s costs, and that cost increase would almost certainly swamp any theoretical increase in profit. Moreover, charging a usage fee to consumers who have always been able to play “their” games as much as they want would almost certainly provoke a serious consumer backlash. As this example illustrates, the pricing instruments that each software platform can use on each side of its market depend to a large extent on the transaction costs involved, on the institutions of that particular market, and on the available technology.

Exclusion and Piracy

For a software platform to be able to charge a positive access or variable price to a certain side of its market, it first needs the means to exclude members of that side who don’t pay. On the end-user side, this is a relatively small problem when the software platform is integrated

into the hardware. Even if the software platform can be easily copied, it is useless without the basic hardware, and there are many ways (including using nonstandard components and nonpublic designs) to make cloning the hardware difficult. It is thus not surprising that we are unaware of any allegations of substantial piracy of Apple computer systems.¹⁴

On the other hand, it does not seem easy to exclude makers of peripheral devices. In any case, we are unaware of any serious attempt to do so in the industries studied here.

When the software platform is not bundled with the basic computer hardware, software piracy can become a major problem, particularly if (as is the case for PCs designed to run Windows) there are numerous third-party hardware producers using standardized components and working with well-documented hardware-software interfaces so that the necessary hardware is easy to buy or build. In this case, the software vendor almost certainly needs to devote resources to fighting piracy. It might license its systems only to hardware vendors who agree (on pain of heavy penalties) not to ship hardware without operating systems, for instance, and devote resources to enforcing that contract provision. It might also hire agents to attempt to buy pirated copies of its system and hand the sellers over to the authorities.

Because these and other sorts of antipiracy measures are costly, whereas copying software is essentially free, some amount of piracy of popular platforms is almost certain to occur. Since the higher the price of the platform, the more tempting it is to copy it illegally, a strategy of selling software separate from hardware fits best with a strategy of selling the software for a relatively low price.

One other aspect of licensing platform software to hardware makers deserves mention. Microsoft has long offered discounts on operating system licenses to computer makers to design and build machines that meet certain standards. These offers are a part of the Market Development Program, and their conditions typically include selling more than

14. Although Apple's June 2005 announcement that they would be switching to Intel architecture started rumors that such piracy might be a concern in the future.

a certain percentage of computers with certain minimum technical specifications (memory, CPU speed, and so on). The rationale is to provide incentives for improving the quality of the computers that reach end users, which in turn stimulates demand for Windows. This is also one method for coordinating innovation across corporate boundaries and, from Microsoft's point of view, of reducing the number of end users who are unhappy that advertised features of Windows don't work on their new computers.

Developers of applications, games, or media content cannot be automatically excluded from using operating systems. It is hard to keep APIs secret (they need to be documented for internal developers, for instance) or to prevent developer tools and programming languages from being copied, especially if they are being sold for a high price. All video game consoles obtain exclusion by using a security chip to prevent games produced by unauthorized developers from running, and one cannot use a mobile phone purchased from one carrier on another's network without the other carrier's SIM card. We are unaware of any other exclusion device that has been used successfully in the industries studied here at any significant scale.

In the smart phone industry, third-party vendors could initially supply applications freely, as in the PC industry. However, realizing that the quality of applications available had a significant impact on the overall user experience, network operators and handset manufacturers began to create *signing programs* for third-party software. These programs resemble the distinction DoCoMo makes between official and unofficial content: nobody is completely excluded, but not every application can obtain official approval. Users know that signed applications satisfy certain quality standards. In 2003, Symbian introduced what has become the most significant signing program, Symbian Signed, endorsed by Symbian's hardware licensees as well as network operators. In addition to granting signed applications a public seal of approval, handset manufacturers such as Nokia and Sony Ericsson, as well as operators such as Orange and T-Mobile, open their distribution channels only to applications that are Symbian-signed. (Developers pay modest fees to Symbian in this process: \$350 for registration and from €185 to €560 for

testing.¹⁵) We understand that it is technically possible on at least some phones for network operators to go a step farther and block end users from installing nonapproved applications on their handsets, and that some are at least considering this step.

Intensity of Use

If it is possible to exclude potential participants, the next choice is whether to do so and to charge something for access or usage. If a positive price is to be charged, access fees, which do not vary with intensity of usage, are typically (but not always) easier to charge than usage fees in the markets we've discussed here. This is because it is easy to monitor purchase, but purchase in these settings typically enables usage at widely different levels of intensity. (Arcade video games provide a counter example: it is much more natural to charge each time a game is played than to sell the rights to unlimited play.) Once I've bought a video game for my Xbox, I can play it every waking hour or toss it in the closet and forget it. The developer would like to charge more in the former case than in the latter, though he would generally like to charge something (an access fee) in the latter case as well. But, as we discussed above, it is simply not easy to monitor and charge for postpurchase usage in this case.

Things are changing with the advent of online gaming. Console online games have followed PC online games by adopting a subscription model; users pay a monthly subscription fee to play.¹⁶ This isn't pure usage pricing, though; that would involve charging for time spent online or at least for each login. This sort of pricing would seem to be feasible; perhaps it is not done for the same reason that most U.S. consumers pay a flat monthly fee for unlimited local (land-line) telephone calls: consumers value having a predictable monthly bill and don't like having to

15. <https://www.symbiansigned.com/app/page/faq>, <https://www.symbiansigned.com/app/page/testhouses>. Registration is done by VeriSign, Inc., and testing by either CapGemini, Mphasis, or NSTL; all are Symbian partners.

16. As we have seen, however, on Xbox Live, users pay a monthly fee to Microsoft and have access to all Xbox-supported online games, whereas the PlayStation online service leaves it to each individual game developer to charge subscription fees to users.

think about the cost consequences of their actions on a minute-by-minute basis.

The access/usage issue arises in slightly different form with applications developers. A software platform vendor could, in principle, charge each third-party application or game developer only a fixed fee for access to its system. One can argue that Apple, Microsoft, Symbian, and Palm-Source do this, but their fees for developer tools and related information at best just cover the costs involved. In fact, one can think of this policy as offering negative prices to credible developers, with more attention devoted (and thus, in effect, a larger subsidy given) to developers that produce the most popular applications—thus, roughly, a negative usage charge on that side of the market.

Any attempt to make significant profit from access charges to developers would run into the exclusion problem discussed above. Moreover, such a policy would clearly be inefficient. Game developers, for instance, differ enormously in scale, and a fee that Electronic Arts would notice would likely exclude most of the firms in the industry. The pricing scheme actually adopted by video game platforms, a royalty for each game sold, is in effect a usage fee for game developers: it charges them more if they derive more value from access to the platform. But, as we discussed above, it leads to access pricing, not usage pricing, of individual games to end users.

Some platforms do have the potential to monitor intensity of usage, at least approximately. If the number of songs downloaded from iTunes were a good measure of iPod usage, Apple could use song pricing on iTunes to levy a usage charge on iPod owners. However, as we discussed in Chapter 8, Apple has decided not to do this and to sell songs on iTunes on roughly a break-even basis. This may be simply a continuation of Apple's long-standing strategy of seeking profit in hardware rather than software, or it may reflect worries that piracy would make higher prices for music unsustainable. RealNetworks' Rhapsody service, in contrast, charges users monthly fees for unlimited streaming access to a million songs, but it does not allow downloading, so that when they stop paying they no longer have access to any songs they have heard before.

In Chapter 7 we noted that one of the main features of i-mode is its sophisticated billing system, which allows NTT DoCoMo to charge users

based on the amount of data they receive and send to and from their i-mode phones. This is arguably a reasonable proxy for consumer usage of (and thus value derived from) the platform. While clearly imperfect, it seems superior to the much less accurate alternative that was used by WAP-based services such as Vizzavi that charged based on time spent on the network. Today, virtually all wireless service providers around the world charge usage fees based on the amount of data transferred or levy a flat fee for unlimited data transfers.

Particularly with ubiquitous access to the Internet, it is possible in principle to monitor intensity of end-user usage for essentially any software platform. One could, for example, imagine that every time a PDA using the Palm OS is linked to a computer connected to the Internet, it would report CPU usage to PalmSource, which would automatically generate a monthly charge to the owner's credit or debit card, just as some Internet service providers do now.¹⁷ But this sort of monitoring is not free, and in most cases suppliers seem to have found it more efficient on balance to let the cheapest method of transacting (per session at video game arcades and per game for home systems) determine whether access or usage is priced, and not to incur the extra cost of instituting a two-part tariff or other more sophisticated pricing system.

Price Discrimination

Like many, if not most, businesses, software platform vendors have employed a variety of forms of price discrimination, some traditional and some not, in order to enhance their profits. For instance, license fees for the Palm OS are negotiated separately for each licensee, while Symbian uses the same price schedule for all its licensees (who are also its owners). From 2002 to 2004 Symbian earned an average of almost half its revenue from providing consulting services to its licensees, helping them to adapt the Symbian operating system to their hardware, but consulting fees have become less important over time, and in 2004 they constituted only one-fourth of total revenue. We do not know how the use of these services varies across licensees or how, if at all, they are marked up. Like

17. After a certain amount of CPU usage, the system could be set up to deny the end user access to applications until it had been linked to the Internet.

Microsoft, both Palm and Symbian offer bulk discounts, so that larger licensees pay lower per-unit license fees.

In video games, it has been argued that console pricing reflects cream skimming, which is a form of intertemporal price discrimination: prices typically decline over the life spans of particular consoles so that on average, the most eager buyers pay higher prices. On the other hand, it is likely that costs decline over time as well, as learning occurs and component prices fall, so the extent of discrimination is not clear. It is our understanding that Sony and Microsoft charge lower per-unit royalties for games that sell many copies.¹⁸ This interesting form of nonlinear pricing strengthens game developers' incentives to focus on a few good games rather than many mediocre games.

Without a detailed knowledge of costs, it is difficult to know whether some pricing strategies have an element of discrimination. I-mode users pay on the basis of the volume of data transferred and iTunes users pay for each song downloaded, while subscribers to RealNetworks' Super Pass service pay monthly fees independent of usage. The first two seem better designed to measure individual differences in demand, while the subscription scheme, which resembles the way local phone service has traditionally been priced in the United States, is no doubt a response to consumers' preference for flat, known fees.

Two relatively unusual pricing practices encountered in these industries reflect the multisided nature of most platform businesses. First, as we noted in Chapter 5, video game console vendors generally offer lower royalties and joint marketing arrangements to developers that develop exclusively for that console.¹⁹ Note that there is no forcing: developers are free to choose whether or not they wish to develop exclusively for a

18. Email correspondence with Nanako Kato and Gerald Cavanagh of SCEI, May–August, 2005, and Dean Takahashi, *Opening the Xbox* (Roseville, Calif.: Prima, 2002). SCEI.

19. Authors' interviews at SCEI. In the economic literature on multisided markets, developers who focus on a single platform are said to single-home, while others are described as multihoming. For a discussion of the strategic implications of these behaviors, see Rochet and Tirole "Platform Competition in Two-Sided Markets"; Jean-Charles Rochet and Jean Tirole, "Tying in Two-Sided Markets and the Impact of the Honor All Cards Rule" (mimeo, IDEI University of Toulouse, 2003); Jean-Charles Rochet and Jean Tirole, "Two-Sided Markets: An Overview" (mimeo, IDEI University of Toulouse, 2004).

particular platform. (This differs from the exclusive contracts used by Nintendo in the late 1980s and early 1990s, which left independent developers with no choice but to support NES exclusively if they wanted to be granted any access at all.) The motivation for offering better prices to exclusive developers is simply that exclusivity offers the console maker a competitive advantage over rival platforms in competition for the end users that constitute the other side of the market. This is especially true for killer games, whose sole presence on a console is sometimes sufficient to entice many users to purchase that console. Another software platform offering better deals in exchange for exclusivity is RealNetworks, which paid Major League Baseball \$20 million in exchange for making game coverage available exclusively through RealPlayer.²⁰

In the video game industry, game developers also seek exclusivity. As we noted in Chapter 5, movie and sports tie-ins are important sources of value to game publishers. Accordingly, significant licensing fees are paid to Hollywood studios and professional sports leagues in exchange for the rights to feature their images, characters, and players in games. And publishers pay significantly higher amounts for exclusive rights: Electronic Arts' exclusive NFL license was rumored to be close to \$200 million for five years, whereas their NBA license cost well below that amount, as it was shared with rival game publishers (the NBA sold nonexclusive rights to five game makers for a total of \$400 million).²¹

The second unusual pricing practice is discrimination among complementors based on "quality." Microsoft has given more favorable license terms to computer makers that offer machines meeting certain design standards. These arrangements do not reflect Microsoft's costs or differences among computer manufacturers; rather, they reflect the greater value of "quality" computers as complements to Microsoft's operating systems. In addition, as we noted above, NTT DoCoMo offers better deals in the form of additional services to "official" i-mode content providers (those endorsed by DoCoMo) than to others, and the Symbian

20. "RealNetworks pays \$20M for baseball audio rights," *UPSIDE Today*, March 27, 2001.

21. "NBA Grants Videogame Rights to 5 Publishers for \$400 Million," *The Wall Street Journal*, March 22, 2005.

Signed program is administered by some network operators and handset providers in a similar fashion. Finally, video game consoles simply exclude poor-quality games completely by denying them the necessary security chips. Multisided platforms have strong interests in raising the quality of the products supplied by complementors to end users.

Price Structures of Multisided Platforms

A fundamental decision facing all multisided platform businesses is choice of a price structure: How much should the platform vendor charge each side relative to the others? Since transactions involving some sides may have significant associated variable costs (the production and distribution costs of video game consoles, for instance), the most illuminating way to analyze observed price structures is to look at the contributions of each side to gross margin or variable profits: revenue minus side-specific variable cost. Should a two-sided platform derive most of its gross margin from one side of the market, and if so, which side, or should it choose a more balanced structure, with both sides making significant contributions to gross margin?

Like all multisided platforms, the pricing structures of the software platforms we have encountered in this book reflect the need to get all unintegrated sides on board: end users, application/game/content developers, and manufacturers of hardware and peripheral equipment. The structures we have examined have three remarkable features. First, all of them are extremely skewed: almost all earn a disproportionate share of their variable profits on only one side of the market, either end users or developers. Second, for all but video games, the platform earns the bulk of its net revenues from end users. The third remarkable feature, which we consider in the next section, is that these structures have been stable over time.

Main Characteristics

As we have seen in Chapter 4, PC operating system vendors such as Microsoft and Apple make virtually all of their profits on the end-user side of the market. Since applications developers tend to multihome and end users tend not to, this is somewhat at odds with the theoretical

prediction, noted above, that *all else equal*, pricing tends to favor the side of the market that does not multihome. Of course, all else is never equal.

Apple makes profits from end users directly by selling Apple computers based on the Mac OS. Microsoft, on the other hand, charges most end users indirectly, through the licensing fees it levies on OEMs, which the latter pass through in the final prices of their computers. (Some Windows users do buy upgrades themselves.) As one would expect in this highly competitive industry, these fees appear to be passed through roughly dollar-for-dollar.²² Apple's variable costs on the user side are the marginal costs of producing each Macintosh computer and installing its software on it, whereas Microsoft has essentially zero marginal costs. (Microsoft distributes master CDs; the computer manufacturers do the copying.)

On the developer side, Apple, Microsoft, and most other operating system vendors devote significant resources to supporting application developers through development tools, conferences, and direct assistance. The prices charged for these services are set to at most cover costs. In fact, some development tools are available for download for free from the software platforms' Web sites.²³ Somewhat less attention is paid to makers of peripheral equipment, in part because they need less information to produce compatible devices. But neither Apple nor Microsoft seeks profits from the provision of this information.

The pricing strategies of PalmSource and Symbian are very similar to Microsoft's strategy in PCs, which in turn is very similar to Microsoft's strategy in PDAs and smart phones. PalmSource and Symbian, like Microsoft, make most of their profits in the form of licensing fees charged to manufacturers of devices running their operating systems, and both offer a great deal of support to third-party developers in exchange for fees that are generally set just to cover costs.

NTT DoCoMo's i-mode mobile Internet platform also earns a disproportionate share of its profits from end users, in this case through

22. Interview with Microsoft.

23. Many of Apple and Microsoft's SDKs are available free of charge on their developer Web sites, and there are some other tools that are sold, but we believe that these are sold at close to cost. See <http://developer.apple.com/>; <http://msdn.microsoft.com/>.

network usage charges. We have seen that DoCoMo also earns some revenues from “official” content providers who choose to use DoCoMo’s billing system, but this accounted for only around 1 percent of total revenue from users in 2004. Furthermore, although these revenues are designed to cover the costs of providing the billing service to official content providers, it is unlikely that they cover the overhead costs DoCoMo incurs in connection with the teams that monitor and select official content.

Net profits from i-mode’s hardware side are negative, as DoCoMo buys the handsets from manufacturers and resells them at a significantly lower price to consumers in order to encourage adoption of the i-mode platform and thus to generate more revenues from network usage charges.²⁴ This practice, known in the mobile phone industry as “handset subsidies,” is yet another version of the cheap razor/expensive blades or two-part-tariff policy we’ve discussed before. Selling a handset both produces profits directly and generates future profits by increasing network traffic. The presence of this second effect makes the optimal handset price lower than it would otherwise be. If this effect is strong enough, it can drive the optimal price below marginal cost, as it apparently does for DoCoMo and other network operators.

As we discussed in Chapter 8, the leading digital media platforms employ markedly different business models. But in all these models, end users are the primary source of variable profit: in Microsoft’s case through licensing of Windows, in RealNetworks’ case through licensing access to content, and in Apple’s case through sales of iPods. None of these vendors extracts profit from content owners, and indeed, Apple and RealNetworks pay for content.

Video game console manufacturers are the single, striking exception to the general “end users pay” pattern in these industries. These firms derive most of their variable profits from games, both by charging royalties to independent or third-party game developers and through sales of games produced in-house (so-called first-party games). Ever since Atari introduced the VCS 2600 and the cheap razor/expensive blades business model in 1977, game consoles have been most often priced at

24. Interview with Takeshi Natsuno of DoCoMo, March 2005.

or below marginal cost.²⁵ However, due to falling costs of components and learning effects, there commonly exist periods of time over a console's life cycle when price exceeds marginal cost.

For example, a Sony executive in 2004 stated that there is a "positive gross margin" on PlayStation 2 sales,²⁶ a statement that our interviews at Sony Computer Entertainment Inc. (SCEI) have confirmed. It is most likely that Sony originally sold the PlayStation consoles below marginal cost, but over time it has been able to make manufacturing more efficient and derive positive gross margins. Nonetheless, even now the largest share of SCEI's PlayStation variable profits come from royalties levied on third-party publishers of video game software and sales of first-party games—between 60 and 70 percent, according to our interviews. In addition, game developers must also pay a fixed fee for the necessary game development tools, but this fee is very small relative to royalty revenues. The published price for a three-year Tools and Middleware license for the PlayStation 2 was approximately \$12,000.²⁷ That is equivalent to typical \$8 royalties levied on only 1,500 copies of a PlayStation game, when hits like *Tomb Raider* sell millions of copies.

The Xbox console also has had negative gross margins. The average selling price of an Xbox console has been \$160 since 2002, yet the average cost of producing it for the same period has been \$304. Microsoft's newest console, the Xbox 360, is also being sold at a loss. The company released the Xbox 360 in November 2005, but does not expect any profits until 2007. By contrast, from its 2001 launch and through December 2003, the company received \$961 million in revenues from software sales—direct sales of its own games plus \$7 per-unit royalties levied on third-party games—and \$313 million from sales of

25. Leonard Herman, *Phoenix: The Fall and Rise of Videogames* (Union, N.J. Rolenta Press, 1997); Southwest Securities "Interactive Entertainment Software: Industry Report," Fall 2000; Peter Coughlan, "Competitive Dynamics in Home Video Games," *Harvard Business Online*, June 13, 2001.

26. "kelly s i." *Mirror*, March 4, 2000; quotation from Takao Yuhara, Sony Corporation Senior Vice President and Group CFO, Sony Corporation earnings conference call, January 28, 2004.

27. Tools & Middleware License for PlayStation(R)2, Sony Corporation, September 8, 1999.

peripherals such as game controllers, memory cards and other plug-ins, and remote controls. First-party games accounted for roughly 70 percent of total game software revenues.²⁸

Some Explanations

What determines these pricing structures? In particular, how can one make sense of the fact that video game consoles have chosen to earn the bulk of their profits on the game developer side of their market, whereas all other software platforms studied here make most of their profits from end users? It is clear that transactions costs can't be the driving force, since extra costs for a security system must be incurred to exclude games from unlicensed developers.

We offered one explanation in Chapter 5, based on the assumption that the number of games purchased by an individual console owner is correlated with the value he or she places on the video game system. Under this assumption the optimal pricing policy is that first introduced with the Atari VCS 2600: price the console low to generate penetration and build demand for games, and make most profit on the games. Once this pricing model has become established, it is difficult for any firm to depart from it by charging a high price for the console and a low price for games, as 3DO learned to its sorrow, since it is hard for a high-priced console to get penetration unless it somehow manages to launch with an unusually large number of great games, and without penetration it won't get great games in the first place.

We have encountered a variety of related alternative explanations that deserve discussion. Many of our interviewees stress that one important reason why PC operating system vendors do not generate variable profits from their application developers is that PC platforms are open. In principle, anyone can develop applications for Windows or Mac OS without explicit consent from Microsoft or Apple. (You don't even need development tools, which are relatively expensive in any case.) In other words, the openness of these software platforms means that their sponsors have

28. Arik Hesseldahl, "Microsoft's Red-Ink Game," *Business Week Online*, November 22, 2005; Andrew Hendley, Adam Holt, Phil Michelson, and Derek Wong, J.P. Morgan North American Equity Research, "Microsoft Corporation: Patience Is a Virtue," January 6, 2004, table 12.

forfeited the ability to exclude developers and therefore the ability to charge them for access and usage of the platform. The same is true for Palm and Symbian and for Microsoft's PDA and smart phone operating systems. In contrast, video game consoles have always been closed platforms and have maintained the ability to exclude through the use of a security system that locks out unauthorized developers. These security systems are necessary to be able to charge royalties to third-party game developers.

Of course, simply having the ability to charge game developers does not explain why they should in fact be charged. Moreover, video game platforms are closed because their owners spent money to close them, and this begs the question of why other platforms have remained open. It is not that hard to imagine Microsoft, Apple or Palm devising software security systems to lock out unauthorized application developers. (Note, in particular, that video game developers pay royalties only for video games that run on consoles; royalties are not charged for games that run on PCs and other open platforms.) If these companies have chosen not to do this, the reason is unlikely to be technological; it is most probably because the costs exceed the benefits. In particular, if it were optimal for operating system vendors to charge independent software developers substantial royalties, similar to those charged by video game consoles, the potential revenues created would likely justify the fixed cost of developing a security system. If, on the other hand, it is not optimal to charge much anyway, then leaving their platforms open is the most cost-effective solution.

When discussing why the price of consoles is set low, manufacturers often argue that their prime customers are particularly price-sensitive and reluctant to pay too much for a platform mainly designed for playing games. They also stress the importance of obtaining a large installed base of users right away in order to reward game developers and give them incentives to keep writing games for that particular console.²⁹ On the other hand, the price of consoles is generally highest at launch. Manufacturing cost is also highest at launch, of course, and consoles are typically priced at or below marginal cost at launch.³⁰ As we have suggested

29. Interview with Nanako Kato and Gerald Cavanagh of SCEI, April 2005.

30. Coughlan, "Competitive Dynamics in Home Video Games."

above, subsequent price cuts reflect both cost reductions and, plausibly, intertemporal price discrimination: eager early adopters pay higher prices and, plausibly, contribute more to variable profits than less interested, late adopters.

The recent economic literature on two-sided markets implies that optimal price structures are determined by side-specific marginal costs, price elasticities of demand on both sides of the market, and the relative intensities of externalities between the two sides.³¹ One general result that has emerged is that the side that “cares” more about the other side should pay more, all else equal. It’s rather difficult to compare indirect network effects or side-specific price responsiveness in software platform industries, however, and examining side-specific costs doesn’t discriminate between, say, Apple’s computers and Sony’s video game consoles.

The only economic modeling framework that has been specifically designed for studying two-sided software platforms is that of Hagiu.³² He shows that the greater is user demand for variety of applications/games/content, the greater is the optimal share of platform profits contributed by developers (as opposed to end users). When demand for variety is higher, products (applications, games or content) are less substitutable, so there is less competition among developers. This allows developers to charge higher prices to end users and some of them to earn high profits, making it harder for the software platform to earn profits directly from end users and easier to extract them from developers via royalties.

Although user demand for product variety is difficult to quantify precisely, it is quite clear that video game users care about product variety

31. Rochet and Tirole, “Platform Competition in Two-Sided Markets,” “Tying in Two-Sided Markets and the Impact of the Honor All Cards Rule,” and “Two-Sided Markets: An Overview”; M. Armstrong and J. Wright, “Two-Sided Markets, Competitive Bottlenecks and Exclusive Contracts” (mimeo, University College, London, and National University of Singapore, 2004); Wilko Bolt and Alexander F. Tieman, “Skewed Pricing in Two-Sided Markets: An IO Approach” (DNB working paper no. 13, October 2004); Wilko Bolt and Alexander F. Tieman, “A Note on Social Welfare and Cost Recovery in Two-Sided Markets” (DNB working paper no. 24, December 2004).

32. Andrei Hagiu, “Two-Sided Platforms: Pricing and Social Efficiency” (<http://www.princeton.edu/~ahagiu/job%20market%20paper%204%202.pdf>).

more than users of computers, PDAs, or smart phone applications. One reason is durability: consumers quickly grow tired of one video game and frequently demand new ones, while we have been using the same basic word-processing program since around 1995. Video game console users buy an average of 3.5 games per year.³³ On the other hand, when it comes to computers, PDAs, smart phones and even mobile Internet services such as i-mode consumers use a remarkably low number of different products, and they stick to the same ones for long periods of time. Put otherwise, video games are more substitutable (interchangeable) from the point of view of consumers than applications for computers, PDAs, or mobile phones. Hence, one would expect to see video game platforms make a larger share of profits on developers relative to the other software platforms. And this is precisely what we observe.

Evolution of Pricing Strategies

It is natural to ask how software platforms' pricing strategies evolve over time. Based on our case studies, the surprising answer is, to a first approximation, they don't. Both what is priced and the basic pricing structures tend to remain constant over time. The only major shift in pricing strategy that we have observed occurred in 1977, when Atari began selling its new VCS game console below manufacturing cost, planning to make its money selling games. This razor/blades strategy has persisted in video games ever since. This continuity in pricing strategies is somewhat surprising, since the environments that software platforms face when they are established in the market differ dramatically from the ones they faced at their inception.

In principle, at least, all two-sided platforms face a rather difficult chicken-and-egg problem at launch. Application/game/content developers, along with third-party hardware and peripheral equipment manufacturers, are naturally reluctant to invest in supporting a new software

33. Schelley Olhava, "Worldwide Videogame Hardware and Software 2004–2008 Forecast and Analysis: Predicting the Future," (IDC report no. 31260), May 2004. Installed base for consoles is 123,368 and software shipments were 434,715.

platform unless they expect it to have a substantial installed base of end-users, and end users are generally reluctant to adopt a platform unless they expect it to be supported by an attractive array of applications, hardware, and peripherals. Economic theorists have argued that platform vendors' initial pricing structures should be designed to overcome this startup problem. The divide-and-conquer strategy calls for subsidizing the participation of one side of the market through fees low enough (possibly negative) to attract it, regardless of the participation of the other side (divide), and then charge positive prices to the latter, who knows that the first side will participate no matter what (conquer).

Once both sides are on board and the platform is clearly viable, of course, there is no need to subsidize the participation of either side in this fashion. Thus, a platform following a divide-and-conquer strategy would be predicted to subsidize participation of at least one side but to do so only temporarily; price should rise substantially to the initially subsidized side. But we have seen no such behavior by any of the platform software businesses we have studied.

On the other hand, we *have* seen changes in platform businesses' scope and integration that may play a similar role. In the extreme case, Palm removed all doubts about the availability of hardware, applications, and peripherals by being completely integrated into all these market sides at the launch of the PalmPilot. Over time, as it became established and these sectors matured, it was able to withdraw and focus on the software platform. Similarly, in video games both Sony and Microsoft acquired several high-profile game developers before releasing their consoles so that both end users and other game developers could reliably expect their consoles to have a number of high-quality games. It may be that because integration decisions are less easily changed than price policies, they serve as more credible devices to affect expectations when products are launched. Price policies, in contrast, seem to be selected with the long run in mind, though it is a bit surprising that so many firms apparently managed to get those policies right from the beginning. In any case, the fact is that price policies have been more stable than integration and scope strategies in the computer-based industries studied in this book.

INSIGHTS

- In principle, the pricing problem for software platforms (and other two-sided businesses) is complex, since it must consider the interdependencies of costs and demands (particularly indirect network effects) linking all sides.
- There is much variety in what software platform vendors charge for, and this is expanding as technology progresses. A general rule of thumb is that even though it is more profitable *in theory* for software platforms to charge both access and usage fees, they generally charge only one or the other.
- Like most businesses, software platform vendors use a variety of forms of price discrimination in order to ignite the market for their products and services. This helps firms target the most profitable customers and enables the most efficient and profitable bundling of services.
- Software platform vendors generally earn the bulk of their profits from only one side of the market, typically end users. The exception is video game console producers, who subsidize end users and incur extra costs to enable them to charge royalties to game developers. A plausible explanation is that the number of games purchased correlates with end users' demand for video game systems, so that making money on games enables console vendors to earn more from those who have the highest demand.
- Pricing strategies of software platforms have been remarkably stable. There are no examples of a software platform-based business pricing low to one side at first and then raising prices after getting that side on board. The pricing structure that ignites the business is generally the pricing structure that persists over time.

This is a section of [doi:10.7551/mitpress/3959.001.0001](https://doi.org/10.7551/mitpress/3959.001.0001)

Invisible Engines

How Software Platforms Drive Innovation and Transform Industries

By: David S. Evans, Andrei Hagiu, Richard Schmalensee

Citation:

Invisible Engines: How Software Platforms Drive Innovation and Transform Industries

By: David S. Evans, Andrei Hagiu, Richard Schmalensee

DOI: 10.7551/mitpress/3959.001.0001

ISBN (electronic): 9780262272421

Publisher: The MIT Press

Published: 2008



The MIT Press

© 2006 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Sabon by SNP Best-set Typesetter Ltd., Hong Kong. Printed and bound in the United States of America.

An electronic version of this book is available under a Creative Commons license.

Library of Congress Cataloging-in-Publication Data

Evans, David S. (David Sparks)

Invisible engines : how software platforms drive innovation and transform industries / David S. Evans, Andrei Hagiu, and Richard Schmalensee.

p. cm.

Includes bibliographical references and index.

ISBN 0-262-05085-4 (alk. paper)

1. Application program interfaces (Computer software). 2. Industries—Data processing. I. Hagiu, Andrei. II. Schmalensee, Richard. III. Title.

QA76.76.A63 E93 2006

005.3—dc22

2006046629

10 9 8 7 6 5 4 3 2 1