10

EXCLUSIONARY BEHAVIOR IN THE MARKET FOR OPERATING SYSTEM SOFTWARE: THE CASE OF MICROSOFT[1]

Glenn A. Woroch, Frederick R. Warren-Boulton
and Kenneth C. Baseman

1. INTRODUCTION AND SUMMARY

This chapter examines Microsoft's licensing practices for its MS-DOS and Microsoft Windows operating system software. Our main focus is on Microsoft's use of CPU (central processing unit, or per-processor) licenses under which an original equipment manufacturer (OEM) of personal computers pays a royalty for each machine it ships rather than for each unit of an operating system it installs. We also examine license provisions that require purchase of a minimum number of rights to install an operating system, Microsoft's tying of both technical support information and Windows to MS-DOS, and Microsoft's attempts to induce technical incompatibility between MS-DOS and its main competitor, DR-DOS.

We begin in section 2 with a brief description of the market for personal computer operating systems, and a history of Microsoft's licensing practices and technical design tactics. We also track the record of antitrust investigations of Microsoft, both here and abroad, that culminated in the signing of a consent decree with the Department of Justice.

In section 3 we examine the main efficiency argument for a CPU license, i.e., that it is a variant on the standard two-part tariff used to achieve first-degree price discrimination which is generally efficient and welfare-enhancing. Upon closer examination, however, we find that the CPU license is not equivalent to a two-part tariff. In this specific factual context, uniform linear prices may maximize profits for a secure monopoly, while a two-part tariff would be neither welfare enhancing nor—absent an exclusionary effect—profit maximizing. We conclude that Justice's attempt to eliminate CPU licenses was subverted by its own endorsement of volume discounts which can approximate lump-sum payments to any desired degree.

Section 4 turns to potential anticompetitive rationales for Microsoft's practices in the DOS market. We begin by observing that markets for many high technology products are characterized by a competitive process where a new product appears with a significantly superior technology or design and sweeps the field. By rapidly displacing the old product

D. Gabel et al. (eds.), *Opening Networks to Competition*

and its old technology, the new product achieves a large market share in a short time, earning high gross profit margins. This situation persists only until the dominant firm's product is itself displaced by another new product. This cycle of a new product with an innovative technology displacing an existing product with an old technology is a process of "creative destruction" in the race to be best. Firms achieve a dominant position, but that position is only transitory because, without artificial barriers to entry, today's dominant or monopoly firm and product can readily be dislodged by a new product developed by a competitor or a new entrant.

When the monopolist's position is protected by strategically erected barriers to entry, however, this displacement process can come to a halt. We examine the possibility that Microsoft has used a variety of exclusionary practices—notably nonlinear pricing and technical incompatibility—not to achieve its initial position but rather to retain that position against new competition. We conclude that, under the conditions present in the operating systems market, such practices can be, and in this instance have been, effective in limiting the growth and threatening the existence of entrants and rivals with small market shares. We conclude that Microsoft's anticompetitive behavior has reduced social welfare.

2. BACKGROUND

2.1 The Market for Personal Computer Operating Systems

2.1.1 Personal computer platforms

Our focus is on the market for packaged software that operates personal computers, and to a lesser extent, the software applications that run using those operating systems. To better understand the market for these products, we must delve into the economics and technology of the personal computer.

PCs can be decomposed into hardware and software components. Some components are essential: every computer system requires a microelectronic chip (usually called the central processing unit, or CPU) plus operating system (OS) software. The OS directs the stream of instructions requested by the applications software, while the CPU performs the numerical computations. Importantly, the CPU and the OS are almost always combined in fixed proportions: one of each is needed per system.

Once an OS is installed, a user can run many kinds of applications software.[2] The most popular packages do word processing, spreadsheet analysis, and database management. Increasingly popular is the use of a graphical user interface (GUI) that simplifies the management of the various applications. Both applications and GUIs are optional components of a personal computer system.

Personal computers are available in several "platforms" that differ in their hardware specifications. The so-called "IBM-compatible" PC is the predominant platform that evolved from the hardware and software specifications of the machine first introduced by IBM in 1981.

2.1.2 Industry structure

The supply of many components is highly concentrated. An overwhelming proportion of IBM-compatible PCs in use today are equipped with CPUs manufactured by the Intel

Corporation. The majority of existing PCs run on one version or another of the operating system sold by Microsoft Corporation. Sales of applications software and peripheral hardware components are far less concentrated.

Hundreds of OEMs assemble hardware components in various configurations (usually called "models"), distribute the machines through retail stores or mail order, and provide technical and repair service. In addition to a few large OEMs such as Compaq, Dell and AST Research in the U.S. and NEC, Toshiba and Hitachi in Asia, there is a host of small resellers. We can safely assume this segment of the market to be competitive.

In the early 1990s, the bulk of new PCs shipped in the U.S. (see table below) arrived loaded with some operating system, usually Microsoft's MS-DOS, and often with the Microsoft Windows interface as well. IBM ships its PCs with one of its own operating systems: PC-DOS or OS/2.[3] The only independent OS (i.e., compatible with, but not a clone or derivative of MS-DOS) were Digital Research Incorporated (DRI's) DR-DOS (which, with Novell's acquisition of DRI in 1991, became known as Novell DOS) and IBM's PC-DOS. Users could purchase OSs at retail stores or direct from the software publisher.

In 1992, it was estimated that the worldwide installed base of personal computers of all platforms totaled over 138 million (Bernstein Research, 1993). Of those, 72 percent were IBM-compatible. Less than a quarter of those machines were equipped with Microsoft Windows.

2.1.3 Supply conditions

Operating system software is very costly to develop and market. For instance, it has been estimated that IBM has spent over $2 billion developing OS/2. In comparison, reproducing and distributing operating system software is extremely cheap. As a result, fixed costs are enormous while marginal costs are negligible. The fixed costs are also largely sunk. The code itself is rarely of much value in other uses. Development teams accumulate expertise and reputation, only a portion of which can be redeployed into other projects.

Besides the irreversible investment in computer code, incumbents acquire sunk, or partially sunk, assets such as customer lists and brand name recognition. Furthermore, any new OS must be compatible with all the applications that were written to that "standard." User switching costs also limit the ability of new entrants to gain a toehold. Of course, these costs erect barriers only when the incumbent firm has a first-mover advantage. However, sunk costs ordinarily imply a first-mover advantage, at least for the current vintage of technology.[4]

2.1.4 History of PC operating systems

Dating back to 1976, Digital Research Incorporated sold a popular operating system, called "CP/M", for use on machines based on Intel's 8-bit 8080 chip. In 1980, in what has been called "the deal of the century," Microsoft paid a mere $100,000 for the rights to a CP/M derivative software package called "Disk Operating System," which, with minor modifications, became the initial MS-DOS. In 1981, when IBM launched its entry into the personal computer market, it selected Intel's new 16-bit 8088 chip as the CPU. It also chose to endorse Microsoft's MS-DOS as the operating system.

IBM's partnership with Microsoft later fell apart. But in the meantime, neither IBM nor DRI stopped developing their own operating systems.[5] Under the terms of the dissolution, IBM continued to develop MS-DOS, and eventually its own variant, PC-DOS, which it loaded on PCs bearing the IBM nameplate. In exchange, IBM agreed to pay Microsoft a royalty for a predetermined number of units.

Having been passed over by IBM, DRI went on to modify CP/M for the Intel 8086 chip, leading to its CP/M-86 product. Later it developed DOS PLUS and then DR-DOS. In April 1990, DRI introduced DR-DOS 5.0 to critical acclaim. Instantly, it began to make inroads into MS-DOS 4.0's market share. By year-end 1990, DR-DOS's share had increased to 10 percent of new OS shipments, leaving MS-DOS with 70 percent and IBM with 18 percent.[6]

Within a month of DR-DOS 5.0's inauguration, Microsoft announced development of MS-DOS 5.0. Curiously, it was to contain nearly all of the innovative features of the DRI product. Yet MS-DOS 5.0 was not commercially available until July 1991, more than a year after DR-DOS 5.0's release. Anticipation of the new Microsoft product, prolonged by continuous Microsoft statements indicating imminent availability, reined in growth of DR-DOS 5.0 sales (Sherer, 1990).

The emergence of the graphical interface played an important role in the events that followed. After repairing bugs in Microsoft Windows 3.0, Microsoft shipped Microsoft Windows 3.1 in April 1991. In that year, 18.5 percent of new PC shipments included Microsoft Windows along with MS-DOS. By 1992, that fraction jumped to 59.7 percent. Over that period, sales of MS-DOS (both with and without Microsoft Windows 3.1) rose 28.9 percent while sales of PC-DOS and DR-DOS fell 15.4 percent (see Table 1). By 1993, the market shares for operating systems on x86 PCs were 79 percent for MS-DOS, 13 percent for PC-DOS, 4 percent for OS/2, 3 percent for DR-DOS and 1 percent for UNIX.[7]

2. MICROSOFT'S PRACTICES

2.1 The CPU License

When first available, MS-DOS was sold to OEMs for a flat fee. Microsoft offered an unrestricted number of copies for $95,000 and, for a limited time, reduced that price by half.[8] Around 1983, Microsoft began to gear its license fees to the level of OEM sales. Each OEM contract was individually negotiated; an external price list never existed.

Over time, Microsoft phased in a new type of royalty contract. By 1992, the "CPU license" became the dominant sales arrangement, with 60 percent of Microsoft's operating system sales made under CPU licenses.[9] Under its terms, affiliated OEMs were required to pay a royalty for every CPU they shipped. Since each machine had a single CPU, the OEM paid for a copy regardless of whether the machine was preloaded with MS-DOS. Microsoft would sell DOS licenses to OEMs which refused the CPU license, but only at significantly higher prices.

Under the CPU license, an OEM usually had to commit also to a minimum "requirement" R that approximates its annual shipments. The one-time charge for this requirement is computed using a negotiated per-unit price p multiplied by R.[10] If an OEM

shipped a machine with a competing operating system, say PC-DOS or DR-DOS, it would receive no reduction in its payment to Microsoft. Consequently, an OEM which accepts a CPU license faces a zero marginal price for units of MS-DOS up to the minimum requirement.[11] In the event that an OEM exceeded its projected volume during the contract period, the per-unit fee p used to calculate the lump sum payment for the first R units would apply to each unit above R. Thus, once the contract is in place, the marginal price is 0 up to R units and p for additional units beyond R.

Table 1 New Shipments of Personal Computer Operating Systems (000s of units)

Company	Operating System	1990	1991	1992
Microsoft	MS-DOS	11,648	13,178	18,525
	w/ Windows	490	2,440	11,056
	w/o Windows	11,158	10,738	7,469
IBM	PC-DOS	3,031	3,003	2,315
DRI/Novell	DR-DOS	1,737	1,819	1,617
DOS Subtotal		28,064	31,178	22,847
Apple	Macintosh	1,411	2,204	2,570
UNIX	UNIX	357	582	797
IBM	OS/2	0	0	409
Other	NEC, etc.	5,079	4,628	4,458
Totals		23,450	25,702	31,080

Sources: Bernstein Research, International Data Corporation

Regardless of whether an OEM ends up shipping more or less than R PCs during the contract, the terms of the CPU license commit the OEM to pay for one unit of MS-DOS for each PC it ships. As a result, customers view themselves as paying double if they use other OSs. If the supplier of a competing OS offers to sell at a per-unit price m, the OEM will only buy the second OS if that OS has a quality advantage over MS-DOS valued at m or more.

Table 2 shows the marginal cost of a PC under various scenarios facing a PC manufacturer, which has signed a CPU license with Microsoft. Let z be the marginal cost of producing the machine excluding an operating system, let X be the number of PCs produced regardless of which operating system is installed, if any at all. When the OEM

ships less than the requirement of R units, its marginal cost of using MS-DOS on the next PC is zero, compared with a marginal cost of m if it chooses another OS. After R PCs have been shipped, the marginal cost is p if the OEM uses MS-DOS, and it is $p + m$ if the OEM uses the other OS. In each case, the increment to marginal cost from using an alternative OS is m. If the machine is shipped "naked," then the total marginal cost is just the production marginal cost, z.

Table 2 Marginal Cost of Different PC Configurations

CPUs shipped by OEM	MC with MS-DOS		MC with Alternative OS	
	OS	PC	OS	PC
X<R	0.00	z	m	z+m
X>R	p	z+p	p+m	z+p+m

In 1992, the average license fee per copy of MS-DOS to a hardware OEM under these CPU licenses has been estimated at $15, far below the average retail price of an upgrade of $49 (Bernstein Research, 1993). All together in that year, Microsoft grossed $399 million on worldwide sales of 18,525,000 units of MS-DOS to OEMs and as upgrades.[12]

The typical CPU license ran for a period of 2 years. It was quite likely an OEM will finish the contract period with unused licenses, in which case the customer does not necessarily receive a credit for its unused units. Microsoft exercised its discretion as to when the OEM could carry forward its unused licenses from the prior year.[13]

In addition to the price incentives for exclusivity embodied in the CPU license, Microsoft was alleged to have responded with a variety of direct penalties should an OEM ship some of its machines with a competing operating system. First, the OEM could be prohibited from carrying forward unused MS-DOS licenses, or be required to renew the CPU license at equal or higher volumes to retain the carry-forward option. In this way, Microsoft's policy on carry forwards could establish a "tie" between each year's sales and the next year's sales.

Second, Microsoft's technical service and support could be withheld from an OEM which installed a competing OS. This practice can disadvantage an OEM which needs this information to match its hardware configuration (especially the choice of the microprocessor, the amount of RAM, and the graphics card) with the demands of the operating software.

Third, the price of Microsoft Windows was allegedly increased to OEMs which purchase OSs from someone other than Microsoft. As far back as the days when Microsoft Windows was called "Interface Manager," Microsoft established a connection between the terms of sale of MS-DOS and its graphical interfaces.[14] Microsoft cautioned OEMs against bundling competing multitasking interfaces (such as Quarterdeck's DESQview, VisiCorp's VisiOn and DRI's GEM) with PC hardware components such as hard disks. Discounts on Microsoft Windows were extended to OEMs which agreed to accept a CPU license for MS-DOS. Those who refused the CPU license or who did not use MS-DOS exclusively, could still purchase Windows, but at a much higher per-unit price.

2.2 Technical Incompatibilities

Coordination on technical standards is crucial between the OS developer and applications developers. Nowhere is this coordination more important than with the publication of the Applications Program Interfaces (APIs) which contain the technical specifications that permit applications programs to communicate with the operating system. Microsoft has left undocumented some of these interfaces. In principle, access to these APIs would allow Microsoft to write applications (such as for its MS Word word processor or its Excel spreadsheet) that work faster and with greater functionality. Furthermore, even if an applications developer were to discover and use these undocumented interfaces, Microsoft could, as long as they remain "unofficial," remove or alter them in later versions of the operating software, rendering parts of the applications useless.

Compatibility is also be crucial to the success of operating system software when it must work with programs that function as an intermediary between operating systems and applications programs such as Microsoft's Windows program. Competitors to MS-DOS need to be aware of the functionality of Microsoft Windows so that their products remain compatible with applications written for Windows. In several instances, Microsoft made it difficult for competitors, especially DRI/Novell's DR-DOS, to achieve compatibility with Microsoft Windows. Concerns over possible incompatibility between DR-DOS and Microsoft Windows resulted in significant declines in DR-DOS sales.

One way for applications programmers to insure compatibility with an operating system is to receive copies of the preliminary version of the software. Known as "beta testing," this gives applications developers an opportunity to fine tune the interaction between the two programs.

In a well-publicized episode, DRI was excluded from the beta testing of Microsoft Windows 3.1 and, subsequently, from the beta testing of Microsoft's Windows for Workgroups product. The importance of compatibility testing with the beta version became evident when applications developers using DR-DOS received error messages warning them of a potential incompatibility with Microsoft Windows. Upon installation, Microsoft Windows 3.1 checked whether the source of the underlying system and the extended memory manager were Microsoft products. If they were not, the user was informed that a problem was detected, and was asked to contact Microsoft's beta technical support for Microsoft Windows 3.1. This message appeared on the screen even though no actual compatibility problem was detected. Indeed, if users continued past the alleged error message, they would discover that Microsoft Windows 3.1 would run in conjunction with DR-DOS.[15] The error messages raised fears of incompatibility among developers and users who contemplated running Microsoft Windows with non-Microsoft OSs. Finally, Microsoft Windows disks included a "Readme" text file that cautioned users that "running Microsoft Windows 3.1 with an operating system other than MS-DOS could cause unexpected results or poor performance."[16]

2.3 Antitrust Action

Microsoft's practices first came to the attention of antitrust authorities in Korea. The Korean Fair Trading Commission launched an investigation that centered on use of the CPU license in Asia. In May 1992, the Korean FTC banned the use of CPU licenses in that

country (Phang, 1992). That action was not very effective, however, because Microsoft then began offering customer-specific price schedules with steep "cliffs" (sharp average price reductions) at volumes close to the customer's requirements.

In June 1990, the U.S. Federal Trade Commission initiated a nonpublic (sic) investigation of Microsoft's practices. The investigation eventually focused on Microsoft's marketing practices for DOS and Windows.[17] Without ever acknowledging the investigation, the Commission twice voted on whether to seek a preliminary injunction requiring Microsoft to cease and desist from its marketing practices. Both times the outcome was a 2-2 tie, resulting in no FTC action.

But then, in an unprecedented move, the Antitrust Division of the Department of Justice (the Department) took up the case and, after extensive further investigation, negotiated a consent decree with Microsoft. On July 15, 1994 the Department filed a civil antitrust complaint along with a proposed Final Judgment to which Microsoft had consented (the Consent Decree).[18] Simultaneously, Microsoft consent to a settlement filed by the European Commission. Next, a Competitive Impact Statement (CIS) was filed as required under the Tunney Act.[19]

The case then took an even more startling twist when Judge Sporkin of the DC District Court refused to play the role of a "mushroom"[20] and rejected the decree as inadequate under the Tunney Act. The U.S. government and Microsoft jointly appealed Judge Sporkin's decision. The U.S. Court of Appeals for the District of Columbia has since upheld the consent decree.

3. THE CPU LICENSE, FIRST-DEGREE PRICE DISCRIMINATION AND QUANTITY DISCOUNTS

3.1 The CPU License as First-Degree Price Discrimination

At first glance it may appear that the CPU license is just a means to provide discounts to OEMs that purchased large volumes of MS-DOS. This is, however, not the case: because the OEM's average royalty payment for MS-DOS is based on the share of its machines shipped with MS-DOS, an OEM that purchases *more* MS-DOS could pay a *higher* per-unit price than one that purchases *fewer* units. This would happen if an OEM purchased more units of MS-DOS than some other OEM, but then proceeded to ship many machines that were loaded with an alternative OS.[21]

Nor can the CPU license be characterized as first-degree price discrimination in any meaningful sense. First-degree price discrimination occurs when a seller charges a two-part fee, consisting of a lump-sum payment for the right to purchase the product and a price for each unit equal to the marginal production cost. Where, as here, marginal cost is essentially zero, first-degree price discrimination requires a per-unit price of zero.[22] It is correct that an OEM which signs under the CPU license (or a take-or-pay license with X > output) has agreed to a lump-sum payment, with an (expected) zero marginal price for one year. However, since the size of the lump-sum payment is based on expected sales multiplied by a per-unit royalty, the OEM knows that if its sales increase, the (apparent) lump-sum payment next year will also increase proportionately (based on the per-unit

royalty Microsoft will be charging in the next year). Thus, for any time horizon longer than one year, the CPU license is a tax on output; it is not first-degree price discrimination.[23]

This finding should not be surprising, however, since first-degree price discrimination would not be profitable to Microsoft (nor would it be welfare-enhancing) when compared to a per-unit royalty. Economists have long recognized the strong efficiency advantages of first-degree price discrimination when customers are final consumers so that their demands are independent. But these results do not carry over to intermediate inputs sold to competing downstream firms. In that case the demands of such customers (the firms in the downstream industry) are clearly not independent (i.e., if my rivals pay less for an input than I do, the price of the final product falls, reducing my demand for the input). Ordover and Panzar (1982) state the issue quite clearly.

> . . . we recast the welfare analysis of the simple two-part tariff using the classical model of perfect competition in which all firms are identical and free entry and exit ensures that the equilibrium output price is equal to minimum average cost. In this context we discover that two-part tariffs are not generally desirable from a welfare standpoint . . . . This is due to the fact that the entry fee, instead of acting as a "lump sum levy," affects both the equilibrium number of firms and their output level. This new distortion must be balanced against the losses due to a unit price in excess of marginal cost.

Where, as here, the input (the OS) is used in fixed proportions with the output (the PC) and the downstream industry is a classic competitive industry with U-shaped average cost curves, Ordover and Panzar find that a strong theoretical result obtains: a monopoly seller of the input would find any two-part tariff, including an all-or-nothing offer where marginal cost to the buyer is zero, *less profitable than a uniform per-unit fee.* In addition, the uniform per-unit fee results in *higher economic welfare than any two part-tariff.* As Ordover and Panzar (1982) put it,

> Most surprisingly, for the empirically relevant class of production processes in which the purchased input is required in fixed proportion to output, we discover that a two-part tariff is never optimal from either a profit or welfare maximizing standpoint. (p. 660.)

> The intuition behind this result is rather straightforward. It is well known that under fixed proportions an upstream uniform pricing monopolist can extract all the profits which an integrated uniform pricing monopolist could reap. Since competition downstream ensures that a uniform price prevails in the final product market, there can be nothing to gain from introducing a two-part tariff; optimal choice of [the per-unit price] allows the monopolist to earn the maximum possible under such circumstances. There is something to lose, however, since an entry fee $e > 0$ causes the downstream firms to operate at an inefficiently large scale. Total (upstream plus downstream) costs are not minimized and a portion of this dead-weight burden falls on the monopolist. Viewed another way, this result reveals the futility of attempting to impose a seemingly nondistortionary

> lump-sum levy e on a perfectly competitive industry with free entry and exit. (pp. 666-67)

In short, even if Microsoft's CPU license (or equivalent volume discounts) did impose a true lump-sum payment, there would be no efficiency or welfare gain that could provide a defense for such a pricing system. Nor would such a licensing system be profitable for Microsoft to impose on OEMs even if those OEMs somehow did not recognize the link between their sales and the lump-sum royalties they paid. Both theory and the available evidence would indicate, therefore, that Microsoft's CPU license (or its equivalent in the form of a volume discount) is not a form of first-degree price discrimination.

While the CPU license does not produce a positive output effect (i.e., encourage efficient utilization of a zero-marginal-cost input), it does have a significant substitution effect. The CPU license induces substitution of MS-DOS for an alternative OS. While this may be privately profitable, the social gain is zero even if it did not induce the exit of rival operating systems such as DR-DOS, with its attendant expected effects on raising the MS-DOS license fees. MS-DOS and any other OS have a near-zero social marginal cost in use. Thus, to the extent that the CPU license induces substitution of MS-DOS for an alternative OS, no cost saving results. Indeed, if, as appears to be the case, other OSs offer greater value than that of MS-DOS, the substitution of MS-DOS for alternative OSs actually reduces efficiency and total welfare even in the short run. The adverse effects on social welfare are even greater in the long run, since the exclusionary nature of the CPU license will deter investments in competing OSs.

Finally, we should note that two other efficiency defenses for CPU licenses were raised, at least during the course of the FTC investigation. CPU licenses, it was argued, might be an effective way both to deter unscrupulous OEMs from engaging in under-reporting the number of units of MS-DOS installed and to reduce software piracy by OEMs, retailers and/or computer users. By reducing the number of "naked" machines shipped by OEMs, a CPU license could eliminate the incentive to engage in piracy and fraud. An examination of the historical record leads us to conclude, however, that the prevention of piracy and fraud is not a plausible explanation for why CPU licensing was introduced. Even more telling, however, is that the CPU license is no more effective at deterring piracy or fraud than are other available but unused non exclusionary alternatives such as a "credited-CPU" licenses.[24]

3.2 The DOJ Consent Decree and Quantity Discounts

Microsoft's practices did not allow an OEM to reduce its total payments to Microsoft if it installed a competing operating system on some of its machines. The Department of Justice's complaint and the CIS clearly state that such contracts are illegal and explains their exclusionary and anticompetitive nature. The consent decree does define and ban three types of contracts—per-processor licenses, lump-sum pricing, and minimum commitments—under which there is no reduction whatsoever in an OEM's total payments to Microsoft when the OEM installs a competing operating system on some of its machines. Nevertheless, the consent decree explicitly permits schemes that amount to *near* per-processor pricing, i.e., extreme quantity discounts that can have the same effect, or an effect sufficient to exclude a competitor.[25]

Thus, if Microsoft sets a royalty of $2.5 million to an OEM with a projected output of 100,000 machines, this would be lump-sum pricing. But if Microsoft sets a royalty of $2.499 million plus $0.01 for each unit of MS-DOS installed, this is not lump-sum pricing and would not be banned by the decree. To eliminate any possible confusion on this issue, the consent decree explicitly allows for license arrangements that embody volume discounts.[26] Thus, our hypothetical sales contract ($2.499 million for the first unit of MS-DOS, one cent for each additional unit) is explicitly legal.

The Department was aware of the potential for anticompetitive uses of quantity discounts.[27] Lacking evidence that Microsoft used volume discounts to foreclose competitors, the Department permitted this practice.[28] However, as long as CPU licenses are available to Microsoft, using quantity discounts to achieve exclusion would be redundant and unnecessary so that one should hardly expect to see them used. Only when CPU licenses are prohibited would we expect to see Microsoft turn to sales practices with an equally exclusionary. This is just what had occurred in Korea in 1992, after the Korean FTC investigated and banned the use of CPU licensing by Microsoft.[29] And, even if the Department did not believe when it entered into the consent decree that Microsoft would turn to exclusionary volume discounts, they should question their belief after the first report of Microsoft turning to such discounts.[30]

4. MARKET-POWER RATIONALES FOR CPU LICENSES

4.1 Workable Competition in Technology Markets with Rapid Technological Change

Economic theory would predict highly volatile market shares under a set of conditions that, to varying degrees, have often characterized PC software markets. Consider a market where numerous potential entrants face no *ex ante* barriers to entry into the development of a new technology: entrepreneurs, usually scientists or engineers themselves, put together teams of scientists and engineers, financed internally from their past successes or from venture capital, with access to a common pool of basic technology and to learning acquired at their previous firms. These new ventures incur significant sunk costs to develop a higher-quality technology that (we shall assume) is protected by laws that cover intellectual property to the optimal extent.[31] The new technology may be simply licensed to users (as to OEMs in the case of software) or embodied in a new product using manufacturing facilities available from competitive firms in a number of markets (e.g., software duplicators, or packagers for shrink-wrapped sales of software at retail). The products embodying these alternative technologies are mutually exclusive in the sense that a customer will almost always use only one operating system on any PC.

When two other conditions also hold, we would expect to observe "competitive" or "socially optimal" performance. First, firms in this market take their competitors' prices as given and unaffected by their own actions, and try to undercut their rivals' (quality adjusted) prices as long as that price exceeds their own marginal cost. Second, customers can costlessly switch among the products of rival suppliers.

Given these two conditions, we would expect to observe that (1) a new technology or product will be developed if (and only if) the expected value of the cost of development is

less than the expected value of the increase in the value to consumers of this technology over the prior technology; (2) the price of the old technology (e.g., the license or royalty fee) will fall to zero upon introduction of the new technology; (3) the price of the new technology will equal the difference in value between the old and the new technologies; and (4) market share will rapidly shift from 100 percent for the old technology to 100 percent for the new technology.

While such competition may seem tough on the players, it can still generate very large rewards to the winners needed to cover the risks and costs of development, and results in even greater benefits to consumers since as each new generation appears, the value added by the prior generation is passed on directly to consumers.[32] It is efficient in terms of production and distribution: a technology is developed if and only if it adds more value than it costs to develop, and that technology is priced, like all products in a competitive market, just below the marginal cost of its next best substitute (the prior technology) plus the value of the quality differential. The case of "perfect" competition thus provides a benchmark for evaluating performance in any particular case.

When the two above conditions do not hold, performance may suffer. For instance, if it is costly for consumers to switch to the new technology, and heterogeneous consumers face different costs and benefits from switching, the old technology will retain market share at a positive price. The new technology will sell at a markup higher than its quality differential over the old technology. Similarly, if both technologies are owned by the same firm, the implicit price of the old technology will not fall all the way to zero. Again, the new technology will sell at a higher price than the quality differential, although it may still be profitable for the firm to set relative prices so as to encourage migration to the new technology.

The resulting deviation from the perfect competition model is not necessarily inefficient to the extent it reflects real costs of learning and equipment. But if owners of the current technology are allowed to erect artificial barriers to the entry of a new technology, those suppliers will earn too much, opportunities for technical change will suffer, and consumers will be harmed.

One might expect something close to the result of the competitive model in operating systems because the industry appears characterized by *ex ante* barriers to entry that are low enough for these industries to be workably competitive (absent exclusionary practices). Given the combination of high fixed development costs and low marginal production and distribution costs, the competition resulting from entry can have a dramatic effect on the profits of the first mover. Not surprisingly, therefore, the incumbent has a strong incentive to make life difficult for subsequent entrants, either by directly increasing their costs or by reducing the attractiveness of their product to consumers, and to do so as soon as possible.

Under certain conditions, it may be possible for a first mover to maintain or even extend its dominant position through certain price and nonprice strategies that seek to exclude or handicap its smaller rivals in dealing with its immediate customers. The goal of such a strategy, rather than to assist in achieving the original large market share which requires having, at least for a while, the first-best technology, would be to artificially preserve that status. The four conditions described below appear to hold in the market for operating systems, where Microsoft successfully maintained an overwhelming market share against competition from a product regarded by many software experts as technically superior. The conditions are:

(1) Buyers at the next level downstream (i.e., OEMs), can be presented with an all-or-nothing choice by the dominant firm that compels them to deal either exclusively or not at all with the dominant firm;[33]

(2) While buyers would be interested in purchasing rivals' products for some of their requirements, they are unwilling to rely exclusively on those rivals' products: at least some of the dominant firm's product is important or even essential to many or even all the downstream firms;

(3) The substitute product requires significant fixed sunk costs to develop, maintain or expand, so that some significant minimum market share is essential for entry or expansion and the market is not contestable (substantial sunk costs are lost in a failed entry attempt); and

(4) The costs to the dominant firm of forcing exclusivity on the downstream firms are relatively low.

4.2 Microsoft's Pricing and Marketing Strategies

Let us now turn to each of the four conditions for exclusivity to be an effective strategy against smaller rivals. Our first condition was that immediate buyers can be presented with an all-or-nothing choice by the dominant firm that compels them either to deal exclusively with the dominant firm or not at all. Here, Microsoft can induce OEMs that wish to incorporate MS-DOS in any of their PCs to use MS-DOS exclusively through either of two policies. First, Microsoft can set per-unit MS-DOS prices that are so high relative to CPU rates as to make selecting the per-unit "option" economically infeasible: the OEM that wishes to use any MS-DOS will in effect be required to sign a CPU contract.[34] The CPU license (or a policy of inducing large carry forwards) then provides a strong economic incentive (a zero cost to the OEM for using MS-DOS at the margin) for the OEMs to use MS-DOS exclusively. Second, Microsoft can refuse to sell Windows to an OEM that purchases any alternatives to MS-DOS, and can cut off the OEM from technical information and other services provided to "favored" OEMs. This imposes a direct penalty on the OEM for using an alternative DOS in addition to the pricing incentive created by the CPU contract.

Our second condition was that, while buyers would be interested in purchasing rivals' products for some of their requirements, they are unwilling to rely exclusively on rivals' products: at least some of the dominant firm's product is very important or even essential to many or even all the downstream firms. In this case, OEMs are very reluctant to purchase OSs exclusively from sources other than Microsoft, at least in the short run, for several reasons. First, requiring a sudden and complete switch from one OS to another imposes real costs that could be avoided under a more gradual transition. Second, actual or threatened technical incompatibility between other Microsoft products, such as Microsoft Windows, and competing versions of DOS results in at least some of the OEM's customers insisting on MS-DOS. Third, withdrawal of Microsoft support services to any OEM that does not enter into a CPU contract (or that purchases DOS from a source other than Microsoft) would impose what is in effect a lump-sum penalty for switching. Finally,

Microsoft can refuse to sell Microsoft Windows to an OEM unless that OEM also purchases MS-DOS through a CPU contract.

Our third condition was that the substitute product requires significant fixed sunk costs to develop, maintain or expand, so that some significant minimum market share is essential for entry or expansion. In this case, given the large non-sunk fixed costs of remaining in the DOS market, any alternative to MS-DOS must either achieve a critical minimum market share, exit the market, or be subsidized indefinitely through other operations of the rival firm.

Our final condition was that the costs to the dominant firm of forcing exclusivity on the downstream firms are relatively low. Here, the cost to Microsoft of excluding rivals from the DOS market is small as long as the share of those rivals remains small. By requiring a CPU contract, Microsoft runs the risk of losing an OEM's entire purchases to a competitor. As long as MS-DOS remains essential, however, no OEM would refuse the CPU contract, and the cost to Microsoft is minimal.

The cost to Microsoft of tying Microsoft Windows to MS-DOS is also very small. Microsoft might have to sacrifice some sales of Microsoft Windows to customers for whom the value of Microsoft Windows is very low, but who would buy it to use with a rival's OS but not with MS-DOS. But until a rival achieves a significant share of the OS market, tying of Windows to other Microsoft products or services (or simply making Microsoft Windows technically incompatible with any rival OS) would, again, impose minimal costs on Microsoft.

Our analysis thus concludes that, as compared with other strategies for maintaining market share—such as cutting prices or merging with entrants—implementing exclusionary practices can be a relatively cost-effective strategy against an entrant which has a superior technology but whose market share is very small. It can thus be characterized as a "fight them on the beaches" or (less kindly) as "economic infanticide." The higher the market share of the entrant, the greater the costs and the less the benefits of this strategy to the established firm. Once, or if, the entrant reaches a critical market share, however, the incumbent can be expected to switch to the alternative defensive strategies or, if the entrant's technology is strictly superior and user switching costs are not significant, to simply abandon the field.

REFERENCES

Baseman, Kenneth C., Frederick R. Warren-Boulton, and Glenn A. Woroch (1995). "Microsoft Plays Hardball: Use of Nonlinear Pricing and Technical Incompatibility to Exclude Rivals in Operating Software." *Antitrust Bulletin* 40, 265-315.

Bernstein Research (1993). "FTC Investigation of Microsoft." New York: Sanford C. Bernstein & Co.

Clark, Don and Laurie Hays (1994). "Microsoft's New Marketing Tactics Draw Complaints." *Wall Street Journal* (12 December), B5.

Goldman-Rohm, Wendy (1993). "Will the FTC Come to Its Senses About Microsoft's Mischief?" *Upside* 5, 10-27.

Manes, Stephen and Paul Andrews (1993a). *Gates: How Microsoft's Mogul Reinvented an Industry and Made Himself the Richest Man in America.* New York: Simon and Schuster.

Manes, Stephen and Paul Andrews (1993b). "Microsoft Monopoly." *Upside* 5, 10-18.

Ordover, Janusz and John Panzar (1982). "On the Nonlinear Pricing of Inputs." *International Economic Review* 6, 659-675.

Phang (1992). "Microsoft Deals Not Fair: Korea." *Asia Computer Weekly* (May), 11-17.

Sherer (1990). "Microsoft Outlines DOS 5.0 to Ward Off DR-DOS." *PC Week* 7, 4.

U.S. Department of Justice (1994). "Proposed Final Judgment and Competitive Impact Statement." 59 Fed. Reg. 42845 (proposed 19 August).

U.S. v. Microsoft Corp. (1995). No. 94-1564 (D.D.C. filed 15 July), unpublished.

U.S. v. Microsoft Corp. (1995). 56 F. 3rd 1448 (C.A.C.D.).

Warren-Boulton, Frederick R. (1977). "Vertical Control by Labor Unions." *American Economic Review* 67, 309-322.

Warren-Boulton, Frederick R., Kenneth C. Baseman, and Glenn A. Woroch (1995a). "Current Copyright Protection Can Make Economic Sense." *The Computer Lawyer* 12, 18-28.

Warren-Boulton, Frederick R., Kenneth C. Baseman, and Glenn A. Woroch (1995b). "The Economics of Intellectual Property Protection for Software: The Proper Role for Copyright." *StandardView* 3, 68-78.

NOTES

1. This paper has been adapted from a previously published article, see note 28 *infra*. We would like to express our appreciation for helpful comments and other assistance to Paul Dennis, David Gabel, Linnet Harlan and Sturge Sobin. We are also grateful to the participants of the Columbia Institute for Tele-Information's "Seminar on Sustaining Competition in Network Industries through Regulating and Pricing Access," especially Janusz Ordover and Bobby Willig, as well as to participants of a July 1995 session of the Western Economic Association conference, especially Ben Klein and Bob Levinson.

2. A PC also requires a layer of software that stands between the CPU and the OS. Called the BIOS, or basic instruction operating system, this code is burned into the machine's ROM (read only memory) chip.

3. OS/2 combines OS software and a GUI in one program.

4. If there are no cost complementarities across vintages of technology, then the requirement to sink substantial investment in software development will not convey an advantage to the firm that succeeded in the first generation of a technology when competes to develop subsequent generations.

5. There were two other significant MS-DOS derivatives. For a while, Compaq Computer had shipped its machines with its own Compaq DOS, and NEC developed NEC-DOS, a proprietary operating system that, until recently, dominated the Japanese market.

6. See Bernstein Research (1993), Exhibit 2, p. 10.

7. See Bernstein Research (1993), Exhibit 2.

8. Manes and Andrews (1993a). Citations are to the edited and condensed version in Manes and Andrews (1993b).

9. The percentage of Microsoft's operating system sales made under CPU agreements rose from 20 percent in FY 1989 to 22 percent in 1990, 27 percent in FY 1991 and 50 percent in FY 1992. By FY 1993, 60 percent of MS-DOS sales to OEMs and 43 percent of Windows sales to OEMs were covered under CPU agreements. See note 22 *infra*.

10. It is paid to Microsoft over the course of the year with an initial installment at the beginning of the year.

11. From every indication, the implicit per-unit charges and requirement levels vary across the contracts signed by different OEMs.

12. See *Ibid.*, Exhibit 3. In that year, Microsoft's sales of Windows through OEM and upgrades totaled $599 million.

13. Whether the unit is marginal or inframarginal, its value to the OEM equals the reduction in license fees when that unit is applied to next year's purchases. Of course, to assess its current value, one must discount for time and for the likelihood that the additional unit will actually be needed.

14. See *supra* note 10.

15. "The only error was that the customer was running Microsoft Windows on a competitor's version of DOS;" Manes and Andrews (1993b).

16. Microsoft refused to address compatibility problems with DRI. Microsoft boldly defended its action, claiming it had no responsibility to assist an operating systems competitor. Microsoft's actions went beyond refusal to assist a competitor, however, as it had engaged in commercial sabotage. See Rohm (1993).

17. At one time, the FTC staff was also investigating whether the relationship between Microsoft's operating systems and applications divisions created remediable competitive problems in markets for applications software.

18. *United States v. Microsoft Corp.*, No. 94-1564 (D.D.C. filed 15 July 1995). Amended versions of the Proposed Final Judgment and the Competitive Impact Statement were filed with the court on July 27, 1994.

19. Proposed Final Judgment and Competitive Impact Statement, 59 Fed. Reg. 42845 (1994) (proposed August 19, 1994).

20. The court noted that "Tunney Act courts are not mushrooms to be placed in a dark corner and sprinkled with fertilizer." Microsoft, 1995 U.S. Dist. LEXIS 1654, p. 42.

21. The actual price paid per unit could be thus higher even if the royalty fee itself incorporated volume discounts. Moreover, if units beyond the required minimum R are sold at a per-unit charge, the marginal price jumps from zero to a positive level once PC production exceeds R, so that purchases beyond the requirements level incur a quantity *premium* (see Table 2, above). Average price is the more typical yardstick for measuring nonlinearity of prices, and in the case of a CPU license, average price falls through the range up to the minimum requirements. Thereafter it may rise or fall depending on whether the average price at the requirements level is lower or higher, respectively, than the per-unit charge for additional sales.

22. It also requires that the lump-sum payment from each OEM be tailored so as to be less than the incremental profit that OEM earns from substituting MS-DOS in place of the next-best alternative.

23. For a discussion of the incentives of an input monopolist to substitute an output tax for above-marginal-cost pricing of the input when inputs can be used in variable proportions, see Warren-Boulton (1977).

24. For an extensive analysis of these two efficiency defenses for Microsoft's CPU license, see Baseman, Warren-Boulton, and Woroch (1995).

25. The core provisions are found in sections IV(H) and II(F) of the Proposed Final Judgment.

26. See Sections IV(F) and IV(H) of the Proposed Final Judgment.

27. See the section of the CIS on "Alternatives to the Proposed Final Judgment."

28. "The Department . . . does not have evidence that Microsoft has, to date, in fact structured its volume discounts to achieve anticompetitive ends." The Competitive Impact Statement.

29. The resulting pricing schedule left the Korean OEMs with essentially no option but to deal exclusively with Microsoft.

30. On December 12, 1994 *The Wall Street Journal* reported that in August, just after the consent decree was signed, Microsoft proposed a contract to Vobis, the German PC maker, that estimated its annual shipments of 88 models at about 475,000 and quoted a Windows price of $28 a copy based on that total. When the chairman of Vobis tried to negotiate a discount based on lower estimated sales, in order to accommodate customers that might ask for OS/2, Microsoft's response was that Vobis would have to pay $83 for each machine under a per-copy license.

31. The optimal degree of protection for intellectual property—in particular, the optimal scope for patent or copyright protection in the computer hardware and software industries—is a matter of considerable debate that we can only touch on here. This article is focused exclusively on the horizontal effects of Microsoft's practices, and so we do not express an opinion here as to the merits of the vertical aspects of the antitrust case against Microsoft. We have dealt with similar issues (i.e., network externalities, sunk investments by users, *de facto* standards and interface specifications) in an analysis of the proper role for copyright in software. See Warren-Boulton, Baseman, and Woroch (1995a) and (1995b).

32. In effect, firms earn a normal (i.e., competitive) risk-adjusted return on their investment, while the value of the underlying opportunity is passed on to consumers.

33. For this condition to hold, arbitrage among OEMs must be uneconomic.

34. Microsoft can also structure its Windows pricing to an OEM in such a fashion as to make it very difficult for OEMs to avoid a Windows CPU contract.