Damaged Goods

Raymond J. Deneckere
Department of Economics
University of Wisconsin—Madison
Madison, WI 53706

R. Preston McAfee
University of Texas at Austin
Austin, TX 78712

Manufacturers may intentionally damage a portion of their goods in order to price discriminate. Many instances of this phenomenon are observed. It may result in a Pareto improvement.

1. Introduction

The 486SX processor of Intel Corporation was initially produced in a curious way. Intel began with a fully functioning 486DX processor, then *disabled the math coprocessor,* to produce a chip that is strictly inferior to the 486DX but more expensive to produce. Nevertheless, in 1991, the 486DX sold for \$588, and the 486SX for \$333, a little over half the price of the chip that is less expensive to produce (Frenkel, 1991).

We will argue in this paper that this is not an isolated incident, and that many manufacturers intentionally damage a portion of their production.[1] The obvious reason for doing so is to permit price discrimination. By producing an inferior substitute, the manufacturer can sell to customers who do not value the superior product so much, without decreasing demand for the superior product very much.[2] The

We thank Bruce Smith, Hal Varian, and seminar participants at Cornell, Harvard, Michigan State, Montreal, NYU, Northwestern, Princeton, Rice, SMU, Texas, Washington, and Yale for helpful discussions.

1. The phenomenon is sufficiently well known among marketing professionals that it has a name: *crimping the product.*

2. Most authors agree that it is possible to price discriminate with differentiated products, by charging distinct markups on the goods. For example, Jean Tirole (1988, p. 134) argues that "It should not be inferred that price discrimination does not occur when differentiated products are sold to different consumers," and specifically cites

Journal of Economics & Management Strategy, Volume 5, Number 2, Summer 1996, 149–174

novelty of this paper is not in noting that damaging a high quality good may be a less expensive way to produce a low quality good than directly manufacturing the low quality good. Indeed, we presume that damaging the superior product is the least expensive way to produce the inferior product. Our insight is that this may be a strict Pareto improvement: the manufacturer and *all* types of consumers strictly benefit from the price discrimination.

The simplest way to see the welfare effects is in the case of two types of consumers. Suppose that a manufacturer, selling only the high quality good, chooses to sell only to the high demand types. When the manufacturer price discriminates and sells a low quality good to the low demand types, the low demand types obviously benefit. Why may the high demand types benefit as well? Let M_H, M_L represent the monopoly prices for the high demand consumers purchasing the high quality good, and the low demand consumers purchasing the low quality good, respectively. Further suppose that at these prices, the high demand consumers would prefer the low quality good, but if the prices were slightly closer together, the high demand consumers prefer the high quality good. In order to introduce the low quality good, then, the manufacturer must reduce the gap between the two monopoly prices.

Since there is a zero first-order effect from reducing M_H on profits, the profit maximizing way to narrow the gap is to reduce M_H and increase M_L slightly. Note that low demand consumers still benefit from introducing the damaged good, because the increase in M_L is compared to not being served at all. Finally, as the manufacturer only introduces the low quality good when profits increase, the very existence of the good tells us that profits go up.

This paper adds to a standard result—that if price discrimination expands output, then a welfare improvement tends to occur—in two ways. First, the standard result assumes no resale. In contrast, we permit free resale, limiting the manufacturer to the more common

---

the example of quality differentiated services. In our context, the case for price discrimination is compelling. At lower cost, the manufacturer could have sold the high quality good, but instead chose to damage a portion of production, in order to be able to charge a lower price on the product. Thus, it is as if the manufacturer were selling the high quality product at two distinct prices—one without damage, and one with. This would not be price discrimination if the manufacturer's costs on the damaged good were lower, but since the manufacturer has incurred costs to damage the good, the manufacturer must be discriminating.

case of second degree (incentive compatible) price discrimination.[3] Second, incentive compatible price discrimination strengthens the case for a Pareto improvement, in some cases making price discrimination a strict Pareto improvement for all types.[4]

The paper is organized as follows. In Section 2, we will present a number of examples of manufacturers intentionally damaging a portion of their production. These examples divide naturally into two categories: those where there are two distinct uses for the product, and those where there is only a single use for the product. With two distinct uses for the product, the natural model is one where there are two types of consumers; for example, the educational and business markets correspond to distinct groups of software users. A model of the dual use environment is analyzed in Section 3. When there is only one use for the product, such as the 486 chip, the natural model involves a continuum of consumers with different reservation prices. We analyze this case in Section 4. We end the paper with some concluding remarks.

2. Damaged Goods

How common is it for manufacturers to intentionally damage a portion of their production? In this section, we will argue that, throughout history and across a broad variety of different industries, manufacturers damage some of their production solely for the purpose of enhancing their discriminatory abilities. We document four examples, and provide a brief summary of a variety of other examples.

3. The most general results known to date (Varian, 1985), while phrased in terms of third degree price discrimination, do allow for demand interdependencies, and hence could be applied to demand structures derived from self-selection models. However, our paper states conditions on the demand primitives (rather than conditions on endogenous variables) and focuses on Pareto improvements (rather than mere welfare increases).

4. Varian (1985) constructs an example where third degree price discrimination results in a Pareto improvement. However, because of the absence of demand interdependencies, the high demand market obtains the same price and utility as in the absence of price discrimination; the Pareto improvement arises because a new market is served. In contrast, we consider the case of second degree price discrimination, where the seller does not condition on observable characteristics of the buyers, but instead offers an incentive compatible price menu, so that buyers self-select into categories by their choice of good to purchase. While it is known that net welfare gains can be obtained in the case of second degree price discrimination, to our knowledge, no one has shown that Pareto improvements may arise, nor investigated the circumstances that tend to lead to Pareto improvements. The main theoretical contribution of our paper is to prove that incurring costs in order to damage production may make everyone better off, and that this is a quite plausible outcome in the dual use case, while a less plausible outcome in the single use case.

2.1 The 486SX

With the introduction of the 486 microprocessor, Intel provided a significant improvement in performance over its predecessor, the 386. The 486 microprocessor improved on the 386 in a number of ways. First, it makes more efficient use of internal clock cycles, thereby performing some operations in one step that would have taken several steps on a 386. Second, it contains an 8 kilobyte internal cache memory, allowing the 486 faster access to instructions than if it had to fetch them from slower external DRAM. The final advantage of the 486 is that it contains a 387 compatible math coprocessor, which handles floating point numerical computations. Installing the coprocessor in the same microprocessor eliminates time consuming communication between the processor and the coprocessor, as occurred with the 386 series. The combined effects of these improvements is that a 486 outperforms a 386-387 combination, even when running at lower clock speeds (operations per second).

In response to fast 386 based microprocessors produced by competitor Advanced Micro Devices, Intel decided to introduce a low cost, high performance alternative to the 386: the 486SX. Intel renamed the original 486 the 486DX. Unlike the 386SX processor[5] after which it was named,

> the 486SX is an exact duplicate of the 486DX, with one important difference—its internal math coprocessor is disabled. (Frenkel, 1991)

Although it is more costly for Intel to produce the 486SX, it sold in 1991 for substantially less: $333 as opposed to $588 for the 486DX. As with the 386, it is possible to improve numerical calculations on a 486SX by purchasing the 487SX math "coprocessor." Unlike the 387SX, however, the 487SX is not a real coprocessor.

> The 487SX math coprocessor is really a 486SX with the floating-point unit (FPU) enabled. Keep in mind that the 486SX is actually a 486DX with the FPU disabled. So in reality, the 487SX is really a 486DX. . . . In fact, the 487SX doesn't coprocess at all. It simply disables the 486SX processor and performs like the 486DX that it really is. (Frenkel, 1991)

5. The 386SX was a 386 with a smaller internal data bus, which allowed PC manufacturers to use the 386SX as a "drop-in" replacement for the previous standard 286, without having to redesign the computer for the 386, yet increasing power and compatibility with the 386 generation of chips.

To obtain the full capabilities of a 486DX based machine, an owner of a 486SX based personal computer must therefore purchase the equivalent of *two* 486DX microprocessors: one with its internal coprocessor purposely disabled, and one that, while labeled a coprocessor, is actually a fully functioning 486DX that disables the operation of the 486SX.

It is not possible to purchase a 487SX and use it as a standalone processor, because Intel designed the 487SX so that it needs the presence of the 486SX to operate, although none of the 486SX's processing ability is actually used. Moreover, it would not be economic to purchase a 487SX alone even if it did work without the 486SX, for Intel sells the 487SX at $799, significantly more than the 486DX.

So why wouldn't the owner of a 486SX computer who finds a strong need for the math coprocessor not just scrap the 486SX and upgrade to a 486DX? Intel has reconfigured the pins on the 486SX, so that the SX socket won't accept the 486DX.

Intel ceased to manufacture 486SXs with disabled math coprocessors in 1991, and began removing the coprocessor. Intel could make the microprocessor small enough to be "surface mounted," that is, mounted without a socket, thereby freeing up space that is at a premium in the fast-growing segment of notebook computers (Seymour, 1991).

2.2 IBM LaserPrinter E

In May 1990, IBM announced the introduction of the LaserPrinter E, a lower cost alternative to its popular LaserPrinter. The LaserPrinter E was virtually identical to the original LaserPrinter, except that the E model printed text at 5 pages per minute (ppm), as opposed to 10 ppm for the LaserPrinter. According to Jones (1990), the LaserPrinter E uses the same "engine" and virtually identical parts, with one exception:

> The controllers in our evaluation unit differed only by virtue of four socketed firmware chips and one surface mounted chip. PC Labs' testing of numerous evaluation units indicated that the LaserPrinter E firmware in effect inserts wait states to slow print speed. . . . IBM has gone to some expense to slow the LaserPrinter in firmware so that it can market it at a lower price.

That is, IBM has added chips to the LaserPrinter E that serve as counters or idlers, chips that perform no function other than to make the machine pause and hence print more slowly. Moreover, this is the *only* difference in the two machines. In particular, the idling only applies to text printing, so that graphics comes out at the same speed.

It is interesting that *PC Magazine* (Jones, 1990) gave a good review of the LaserPrinter E, calling it "the obvious choice" over the Hewlett-Packard IIP. For an additional $1099, one can upgrade the LaserPrinter E to identical performance with the LaserPrinter, bringing the total cost of the upgraded LaserPrinter E to $200 more than the original LaserPrinter.

2.3 Sony MiniDiscs

Sony recently introduced a new digital recording-playback format intended to replace the analog audio cassette, but offering greater convenience and durability. To achieve the small form factor deemed necessary for success (the audio cassette's popularity derives from its small format, not its sound quality or durability) and still be able to provide 74 minutes of music, Sony's engineers devised a data compression algorithm that permits squeezing the content of an entire audio Compact Disc onto a disc which is only 2.5 in. in diameter. The MiniDisc is not only smaller than a regular CD but is also immune to the interruption of music caused by shock or vibration in portable applications. Sony accomplished this by inserting a memory buffer between the laser pickup and the digital decoding circuitry, a feature now being introduced on standard Compact Disc players.

MiniDiscs are similar in appearance to 3.5 in. computer diskettes, and come in two varieties: prerecorded and recordable. The prerecorded variety is essentially a miniature CD housed in a plastic shell: like its bigger brother, it uses a laser beam to read the information encoded on the surface of the disc; its principle of operation is therefore entirely optical. The recordable variety looks externally the same as the prerecorded discs, but uses a technology originally developed for computer data storage: magneto-optical recording.[6] Sony produces complementary hardware that either includes just a playback mechanism able to read both types of discs (intended primarily for portable or in-car use), or both a recording and a playback mechanism (intended primarily for home use).

Prerecorded MDs are priced in the same range, but slightly below CDs. Some 400 titles are currently available, mostly from Sony's own label. Blank MDs come in two varieties: 60-minute discs and 74-

6. For an in-depth discussion of the technology underlying MiniDiscs, see Harley (1992, 1994).

minute discs. The list prices for these discs are currently $13.99 and $16.99. Despite the difference in price and recording length, the two formats are physically identical:

> The 60- and 74-minute discs are identical in manufacture. A code in the table of contents identifies a 60-minute disc and prevents recording beyond this length, even though there's room on the media. (Harley, 1994)

One might think that a clever user could circumvent this scheme by constructing a device that alters the table of contents. However, Sony has made this nearly impossible:

> Blank MDs are polycarbonate substrates coated with very thin layers of magnetic material. A ring of polycarbonate at the inner radius is left uncoated. This area, called the "lead in," has pits impressed in it just as on a CD. The MD recorder reads the information in this un-erasable area, which includes the optimum laser power for recording and the disc playback time. . . . The only difference between a 74-minute disc and blanks of shorter playing time is the information encoded in the lead-in area: it tells the player how much recording time is available. (Harley, 1992)

Sony already has plans to make the technology available for computer storage. An MD data disk will have a maximum capacity of 128 Mb.

2.4 TONTINES

While the above examples are all of recent vintage, crimping the product is by no means a new phenomenon. In fact, the British and French governments widely used the practice during the seventeenth and eighteenth centuries. The frequent wars during this period produced large government deficits that were financed with a variety of different debt instruments. Interestingly, what differentiated the liabilities was not so much their maturity structure as their risk structure. In addition to relatively riskless government bonds yielding a normal rate of return, the governments also issued life annuities carrying much higher returns. The purchaser of such an annuity would name a nominee, and receive interest as long as the nominee remained alive. Unlike modern annuities, anyone could be named as the nominee. This effectively provided the annuitant with a random return uncontingent upon his own life. Meanwhile, by selling the annuities to a large subscriber base, the government faced little or no risk.[7] It might be argued that these annuities satisfied a desire for gambling, i.e., were effectively government sponsored lotteries. However, the fact that these securities paid *higher* interest rates and were only available in large denominations—approximately the average annual income of the time—indicates that they were inferior commodities. Interestingly, Smith and Villamil (1993) argue that tontines were created in an attempt to price discriminate between individuals who differed in their private investment opportunities.[8] Such random payout bonds disappeared as investment and saving opportunities increased, although Kurdistan is apparently offering them today.[9]

2.5 Other Examples

It is probably not surprising that manufacturers conceal damaged goods, since something seems wasteful about damaging a good in order to extract more revenue. Three out of the four examples above concern electronics or computer hardware. This is no accident: crimping the product is extremely pervasive in these industries, and easier to document. This subsection presents some brief summaries of other examples of damaged production.

2.5.1 IBM 2319 Disk Drive:[10] A disk drive is composed of two major components. The *spindle* includes the disk platter, a motor to turn it, read/write heads, and the actual spindle that the platter turns on. The second major component is the *controller*, which connects the disk drive to the computer and controls the actions of the spindle. In the late 1960s, IBM began to face strong competition in the disk drive market, which provided memory for the IBM 360 mainframe. Five companies, notably Telex and Memorex, offered spindles superior to IBM's, which could be used with IBM controllers. In addition, Memorex offered a separate controller, which meant customers could pur-

7. In the case of a lottery bond, a bond with a random payout but without the annuity feature of tontines, offered by England, the bond was bundled with a lottery, and offered significantly higher average rates of return than bonds without lotteries, corresponding to a lower price. A tontine divides a fixed amount of money among the annuitants with surviving nominees, thereby being risk-free from the government's perspective, but risky from the individual's perspective.

8. Section 6 of Smith and Villamil (1993), upon which the above discussion is based, contains a detailed discussion of the various types of debt instruments in use during this period, as well as an extensive bibliography.

9. Source: conversations with Bruce Smith of Cornell University.

10. This material is derived from DeLamarter (1986, Chapter 12), who worked as an economist for the DOJ in the famous IBM case. For a contrasting view, see Fisher, McGowan, and Greenwood (1983).

chase Memorex disk drives and just plug them into the IBM mainframe. Worse still for IBM, it had introduced the IBM 370 mainframe, more powerful than the old 360s, but its new disk drive system, dubbed Merlin, would not be ready for introduction for several years, and thus IBM faced loss of disk drive sales for new customers as well as old customers.

To protect the market for 370 disk drives, IBM introduced a scheme internally called "Apricot," a name later changed to "Mallard." According to DeLamarter (1986) this scheme worked as follows: IBM renamed its existing 2314 disk drive the 2319, and integrated the controller, which was previously an outboard device (separate unit), into the single unit, thereby limiting the number of additional spindles that could be plugged into the unit, and undercutting the market for rivals' spindles. This, of course, didn't eliminate Memorex, which sold controllers as well. To undercut Memorex, IBM changed the controller interface as well, thus forcing any rival who wished to offer both a controller and a spindle to decode the controller communication language.

The extent of the price discrimination is summarized by DeLamarter (1986):

> IBM would favor 370 customers at the expense of 360 users, offering each group essentially the same product but at widely different prices. Where IBM charged $256,000 for eight unbundled 2314 spindles and a controller for use on the 360, a similar number of 2319 spindles along with the file adaptor [changed controller] on a 370/135 processor went for as little as $145,415.

IBM vice president P. W. Knaplund described the 2319 relabeling as a "gimicky tactic" to "buy time."

2.5.2 Consumer Electronics: Crimping the product is a popular tactic in consumer electronics. We have heard numerous accounts of how lower priced models of consumer electronics (such as pocket calculators, video equipment, VCRs, and multitesters) differ from their higher priced alternatives only by having some of the features disabled.[11] Unfortunately, the use of this strategy in consumer electronics products has proven harder to document. Nevertheless, Nagle (1987, p. 186) reports:

11. For example, in an April 1993 internet message posted to the rec.video discussion group, Terry Jeffery (UK) reported discovering undocumented features on his Cannon video-camera.

> A leading manufacturer of pocket calculators . . . sold a card programmable version of one calculator for much more than the nonprogrammable version. The only practical difference between the two was a slot in the plastic case of the programmable version where the cards could be inserted.

We have been told (but were unable to verify) that a consumer electronics magazine printed an article explaining how to convert the nonprogrammable version into a programmable version. Another leading calculator manufacturer, Sharp Electronics, produces a calculator that differs from the higher priced scientific version only in that its buttons do not have the alternate functions imprinted upon them.[12]

Robert Harley, one of *Stereophile's* technical editors, likens Sony's MD strategy to that of the hand-held multimeter industry:

> This is analogous to a trick of the hand-held multimeter industry. The same electronics are in every meter throughout the product line; the less expensive models merely have some of their features disabled. (Harley, 1992)

While deeply entrenched as a strategy in electronics and computers, crimping the product occurs in a broad range of other industries, as the next few examples demonstrate.

2.5.3 Educational Software: It appears that the normal way of producing "student versions" or educational versions of software is to put limiting factors into the full-featured versions, thus destroying some of their capability. We know of two examples.

Wolfram Research, Inc.[13] sells a student version of its popular mathematics program, Mathematica, for $180, less than a quarter of the normal price. The student version implements the complete Mathematica program with one exception: it does not use a math coprocessor, even if one is present on the student's computer. This disabling of the math coprocessor makes some kinds of numerical calculations significantly slower. Mathematica requires a fairly powerful microcomputer to operate; most student users are therefore likely to already have a coprocessor.[14]

12. It could, however, be argued that some people may prefer not having access to the scientific functions: square root buttons only serve to confuse them.

13. The source for this material is conversations with Hal Varian, April 22, 1993. Varian has edited a book entitled *Economic and Financial Modelling with Mathematica* and is familiar with the marketing practices of Wolfram Research.

14. Removing the calls to the coprocessor is a simple task, which directly incurs an insignificant additional cost. However, some additional cost would arise from marketing and supporting two versions of the program.

Data Desk is an exploratory data analysis and statistics package for the MacIntosh computer. The full retail version sells for $95.[15] The student version, which has "reduced data handling capabilities," sells for $69.95.

2.5.4 Buying Clubs: The proliferation of discount "buying club" stores such as Sam's, Costco, and The Price Club has segmented the market for many consumer items according to the quantity purchased. Buying clubs specialize in large quantities. Purchases in these outlets, of course, undercut the normal grocery store market, and manufacturers have responded in two ways: by bundling a number of units together, to produce a minimum purchase larger than would normally be demanded by even a large family, and by producing larger sizes specifically for this market.

Creating multipacks, or bundling, incurs additional cost directly. Manufacturers turn to contract packagers to create multipacks "because they don't want the expense of designing and building dedicated in-house lines for filling or multipacking larger-size packages" (Larson, 1993). Other manufacturers, including Chinet (disposable tableware) and Mrs. Paul's (fish sticks), design and manufacture separate production runs for the warehouse market.

Grocers are quite concerned about the growth of the warehouse club market siphoning off demand from retail grocery stores, and have responded by obtaining larger sizes from manufacturers. Traditionally, Chinet sold its products in packages of fifteen. Having introduced packages of 125 units for warehouse clubs, it introduced a 40-count package for "economy aisles" in grocery stores. Nonetheless, manufacturers are segmenting the market:

> With larger sizes, manufacturers are creating stumbling blocks for wholesalers [who sell to retail grocers]; making it difficult for us to buy those items with a low price per ounce.[16]

Interestingly, packagers note the inefficiency of these large bundles associated with consumer sales:

15. Data Desk is produced by Data Description of Ithaca, NY. Prices are as of December 1991, as reported in *The Higher Education Product Companion*, Vol. 1, No. 1, p. 18.

16. Marty Rodgers, quoted in *Progressive Grocer*, May 1992, p. 86. Much of this issue of the magazine is devoted to grocers grousing about competition from warehouse clubs. Another adds "We have seen items in competitors' stores that we were not shown. When we tracked it to the manufacturer, we were told that those items were not made for our class of trade.... I might not need a three-pack of toothpaste banded together, but I want to know it's available" (p. 92).

> Many warehouse shoppers have no intention of using, say, a bundle of 24 rolls of paper towels themselves. They just don't have room, or their family isn't big enough. So they "team-buy" with a friend. Then they split up the multi-pack.[17]

This evidence is suggestive that manufacturers create larger sizes than are efficient to assist in segmenting the market.

2.5.5 Chemicals: Distinct uses for a given chemical will generally offer an opportunity for price discrimination, provided the manufacturer can reduce arbitrage. One method of reducing arbitrage is to add an adulterant to the chemical sold for a low value use, which seriously compromises the high value use. For example, Brazil added gasoline to ethanol sold as automobile fuel to prevent people from drinking the ethanol. The newspapers often contain stories of medicines with vast price differences between human and veterinary use. Even within veterinary use, medicines used for distinct animals may have different values, and the medicine aimed for low value use may come bundled with vitamins, to deter the high value users (whose animals don't need the vitamins) from switching.[18]

In the following two examples, there is no evidence available to us that the manufacturer actually damaged the product. Instead, there is evidence that the manufacturer *contemplated* damaging the product, and in one case, expended significant resources to compute the best means of damaging the product.

Methyl methacrylate (MM) is a plastic with a variety of industrial uses. It is also used to make dentures. According to Stocking and Watkins (1947), the two manufacturers of MM, du Pont and Rohm & Haas, followed a uniform price policy and acted as a cartel. Their pricing policy certainly corroborates this claim: they sold the powdered version of MM (polymer) for industrial uses at 85 cents per pound, and a prepared mixture consisting of powder and liquid (monomer) MM for $22 per pound to licensed dental laboratories. The price difference was evidently too great, and attracted

> bootleggers who found they could crack the powder *back to liquid,* and sell the polymer and monomer together at a profit to the dental trade. (Stocking and Watkins, 1947, p. 403; italics added)

17. John Berkeley, quoted in Larson (1993).

18. An imperfect example is cooking wine: this is ordinary wine with sufficient salt added to make it undrinkable. Cooking wine is generally sold to avoid paying taxes rather than to screen out high value users. Originally, cooking wine was developed to solve the moral hazard problem associated with cooks drinking the wine.

Rohm & Haas considered adulterating the powdered version so that it would be unsuitable for use in dentures and would be prohibited by the Food and Drug Administration for that use. A licensee of Rohm & Haas suggested

> A millionth of one percent of arsenic or lead might cause them [the FDA] to confiscate every bootleg unit in the country. There ought to be a trace of something that would make them rear up. (Quoted by Stocking and Watkins, 1947, p. 403)

There is no evidence that Rohm & Haas put this policy into effect, although they called it "a very fine method of controlling the bootleg situation." However, Rohm & Haas did resort to the less effective strategy of planting a rumor that they had adulterated their powdered MM (Nagle, 1987).

According to Stocking and Watkins (1947), two chemical companies, du Pont and General Aniline, possessed exclusive U.S. rights to market "Monastral" colors, used for both paints and textile dyeing. The use in paints required low prices, while the use in textiles permitted quite high prices. Both companies held conferences and ran experiments to determine the feasibility of adding contaminants to the colors that would render them suitable for paint but not for textiles. Three distinct strategies for contamination were considered. Ground glass would damage painting rolls used in textiles, but have an insignificant effect on paints. Compounds that would cause cotton to deteriorate rapidly, but not affect paint, were also considered. Finally, compounds that would irritate skin and cause dermatitis could be added, again to prohibit use in textile dyeing.

As we have seen, manufacturers in many circumstances disable features, degrade performance, or otherwise damage products to create a lower quality good, which they may sell at a lower price without significantly reducing demand for the high quality good. The examples naturally divide into two categories: those where there is a secondary, low value use for the product (such as educational use in software, or paint dyeing for pigments), which we call the dual use case, and those where there is a single use for the product, such as laser printers and microprocessors. We now turn to the welfare implications of damaged goods in the dual use case.

3. The Dual Use Case

There are two types of consumers, denoted X and Y, and a monopoly producer of two qualities, L and H, for low and high. Consumers buy one good or the other, but not both, and base their purchase decision on which good yields the highest net surplus. Let x_L (respectively x_H) denote consumer X's demand for the low (respectively high) quality good, when this is the only version available for purchase. Similarly let y_L and y_H represent consumer Y's respective demands. We denote the monopoly prices on these demand curves by M_i^z, for $i \in \{L, H\}$ and $z \in \{x, y\}$. For example, M_L^x is the monopoly price associated with demand x_L. The monopolist's marginal cost of production of good L and good H are constant, and denoted by c_L and c_H.

Our first set of assumptions ensures that the model fits the applications discussed in Section 2:

$$0 \leq c_H \leq c_L, \tag{1}$$

$$y_H(p) \geq y_L(p) \quad \text{and} \quad x_H(p) \geq x_L(p), \tag{2}$$

and

$$(\forall p_L, p_H) \qquad 0 < \int_{p_L}^{\infty} x_L(p)\, dp \leq \int_{p_H}^{\infty} x_H(p)\, dp$$
$$\Rightarrow \int_{p_L}^{\infty} y_L(p)\, dp < \int_{p_H}^{\infty} y_H(p)\, dp. \tag{3}$$

Inequality (1) just formalizes the notion that L represents an altered, and hence more costly, version of H. Inequality (2) guarantees that H is indeed the high quality good, with increased demand by both types of consumers. Jointly, conditions (1) and (2) imply that if the monopolist were required to sell only one quality, he would offer high quality:[19]

$$(\forall p \geq c_H) \qquad (p - c_L)[x_L(p) + y_L(p)] \leq (p - c_H)[x_H(p) + y_H(p)].$$

Recall that the consumer surplus associated with demand q and price p is $\text{CS} = \int_p^{\infty} q(z)\, dz$. Thus, assumption (3) says that whenever consumer X weakly prefers purchasing H to purchasing L, consumer Y strictly prefers purchasing H.[20] This ensures that if the low quality

19. This is the only place in our argument where we use assumptions (1) and (2). The conclusions of Theorem 1 remain valid if we impose the above condition directly, increasing the range of applicability beyond situations where (1) and (2) are satisfied. However, while the theory applies to cases where $c_L < c_H$, these are economically not very interesting, for it is then no longer obvious that price discrimination is occurring, and less surprising that introducing L leads to a Pareto improvement.

20. There are various more primitive assumptions that can be imposed to imply condition (3). In particular, if either $x_L = x_H$ or $y_L = x_L$ and $y_H > x_H$ whenever $x_H > 0$, then (3) holds. Also, if $\inf \{p \mid x_L(p) = 0\} = \inf \{p \mid x_H(p) = 0\}$ and if for all p_L, p_H such that $x_L(p_L) > 0$ we have $x_H(p_H)/x_L(p_L) < y_H(p_H)/y_L(p_L)$, then (3) is satisfied. None of these appear to improve on (3) directly, which [in conjunction with assumption (4) below] is interpretable as stating that Y is more H loving than X.

good is introduced, it will be targeted towards the X segment of the market.

Our next set of assumptions serves to guarantee that the introduction of good L is profitable, and produces a Pareto improvement. Let $M_H^{xy} = \arg\max_p (p - c_H)[x_H(p) + y_H(p)]$. We assume that the X market is not served if the firm sells only one quality. That is,

$$x_H(M_H^{xy}) = 0. \tag{4}$$

If (4) fails, then typically we have the monopoly price for both markets falling between the monopoly prices for the markets individually, and thus introducing the L good will tend to increase the price charged for H. This effect need not dominate the effects of the incentive constraints, but such a consideration does not appear to lead to an economically meaningful characterization.

It is possible that, at any price for which the X consumer is willing to purchase a positive quantity of the L good, the Y consumer prefers the L good over the monopoly price for the H good: This will tend to make introducing the low quality good unprofitable. To rule this out, we assume

$$\int_{M_H^y}^{\infty} y_H(p)\, dp \geq \int_{\overline{p}_L}^{\infty} y_L(p)\, dp. \tag{5}$$

The inequality (5) ensures that the Y consumer would prefer purchasing H at the monopoly price to purchasing L at the price $\overline{p}_L$, the lowest price for which the X demand for L is zero.

Under these assumptions, the firm will always choose to introduce the low quality good L, and its introduction is a Pareto improvement.

THEOREM 1 *Suppose that* (1)–(5) *hold, and suppose that* $x_L(c_L) > 0$. *Then the introduction of the good L is a Pareto improvement. If, in addition,*

$$\int_{M_H^y}^{\infty} y_H(p)\, dp < \int_{M_L^x}^{\infty} y_L(p)\, dp, \tag{6}$$

then the improvement is strict: all three agents strictly benefit.

All proofs are contained in the Appendix.

Theorem 1 captures the intuition provided in the introduction. If the low demand L type is not served when only the high quality good is introduced, then the introduction of the low quality good benefits the low demand type and the firm. If, in addition, the incentive constraints on the high demand type bind at the monopoly prices, then the high demand type benefits as well, because his price is decreased to deter him from buying the inferior good.

The assumptions (1–3) and (5–6) are invariant to rescaling both X demand curves or both Y demand curves. As a result, provided M_H^{xy} stays sufficiently high that the X types remain excluded from the H market, the existing theory applies to the case where there are many distinct agents of the X type and many agents of the Y type.

Theorem 1 seems relevant for several of the examples discussed in the previous section. The key question is whether the low demand type would be served in the absence of the lower quality good. Thus, chemical products, pharmaceuticals, and software used for both business and education would seem to fit this model. In contrast, the 486 microprocessor and the IBM LaserPrinter E seem best modeled by a continuum of consumer types, rather than two distinct markets.

4. The Single Use Case

Index the consumers by the value v of the high quality good. We assume v has cumulative distribution function F, with continuous density f, and that F has support $[a, b]$. The value of the low quality good to a type v consumer is $\lambda(v)$. We assume

$$\lambda(a) \leq a \quad \text{and} \quad (\forall v) \quad 0 \leq \lambda'(v) < 1. \tag{7}$$

The monopolist has constant marginal cost c_H for the high quality good and constant marginal cost c_L for the low quality good. In keeping with the applications previously discussed, we also assume[21]

$$a \leq c_H \leq c_L < b. \tag{8}$$

Note that the demand for the high quality good at price p is given by $1 - F(p)$, and that the demand for the low quality good at price p is given by $1 - F(\lambda^{-1}(p))$. Consequently, as in Section 3, assumptions (7) and (8) imply that if the firm produces only one quality, it produces high quality:

$$(\forall p \geq c_H) \qquad (p - c_H)[1 - F(p)] \geq (p - c_L)[1 - F(\lambda^{-1}(p))].$$

When offering only high quality, the profit maximizing price p_1 must satisfy

$$0 = p_1 - c_H - \frac{1 - F(p_1)}{f(p_1)}. \tag{9}$$

21. Forcing $c_H \geq a$ is convenient to ensure an interior solution to the firm's maximization problem, but not necessary. In particular, it is possible to place assumptions directly on the inverse hazard rates used below.

To ensure uniqueness of a solution to eq. (9), we employ the usual hazard rate assumption, familiar in all adverse selection models:

$$(\forall x \in (a, b)) \qquad x - \frac{1 - F(x)}{f(x)} \text{ is increasing.} \tag{10}$$

We also assume that $\lim_{x \to b} [1 - F(x)]/f(x) = 0$. This is satisfied is F is analytic or if $f(b) > 0$. Conditions (8) and (10) then imply that a solution to eq. (9) exists and satisfies $a < p_1 < b$.

The condition for the low quality good, analogous to (10), will also prove useful. Selling only the low quality good, the firm earns

$$\pi_L(p) = (p - c_L)[1 - F(\lambda^{-1}(p))]. \tag{11}$$

Thus,

$$0 = \frac{\partial \pi_L}{\partial p} = -\frac{f(\lambda^{-1}(p))}{\lambda'(\lambda^{-1}(p))}\left(p - c_L - \lambda'(\lambda^{-1}(p))\frac{1 - F(\lambda^{-1}(p))}{f(\lambda^{-1}(p))}\right). \tag{12}$$

To guarantee that the profit function π_L has a unique maximum for every $c_L \in (a, \lambda(b))$, we assume

$$\lambda(v) - \lambda'(v)\frac{1 - F(v)}{f(v)} \text{ is increasing.} \tag{13}$$

We now turn to the case of two qualities. Note that when the firm offers only high quality, some segment of the market remains unserved, since $p_1 > a$. Introducing L can draw some of these consumers into the market, and hence is potentially profitable. However, by introducing L the seller necessarily cannibalizes some of his high quality market. Whether or not introducing L is profitable depends upon the strength of these two opposing forces. We will now set up the seller's optimization problem, and provide necessary and sufficient conditions under which he is willing to introduce L.

By (7), the premium a consumer is willing to pay for the increase in quality of H over L, $v - \lambda(v)$, is increasing in v. This ensures that if both goods are offered for sale, high quality will be targeted towards high valuation consumers, and low quality towards low valuation consumers. More precisely, let v_H be the consumer type who is indifferent between purchasing either good,[22]

$$v_H - p_H = \lambda(v_H) - p_L, \tag{14}$$

22. The critical value v_H may lie outside the range $[a, b]$, but profit maximization ensures that this will never happen in equilibrium.

and let v_L be the type who is indifferent between purchasing L and not purchasing at all,[23]

$$p_L = \lambda(v_L). \tag{15}$$

Then consumers in the interval $[v_L, v_H]$ purchase L, and consumers in the interval $[v_H, b]$ purchase H.[24]

It is useful to express the monopolist's profits in terms of the types of consumers making purchases rather than the prices directly:

$$\begin{aligned}\pi &= (p_H - c_H)[1 - F(v_H)] + (p_L - c_L)[F(v_H) - F(v_L)]\\ &= [v_H - \lambda(v_H) + \lambda(v_L) - c_H][1 - F(v_H)] + [\lambda(v_L) - c_L]\\ &\quad \times [F(v_H) - F(v_L)]\\ &= [v_H - \lambda(v_H) + c_L - c_H][1 - F(v_H)] + [\lambda(v_L) - c_L][1 - F(v_L)].\end{aligned} \tag{16}$$

That is, we can view the firm's maximization problem as maximizing π subject to $a \leq v_L \leq v_H \leq b$. If $v_L = v_H$, then π gives the one quality outcome.

If $v_L < v_H$, the first order conditions for maximizing π are

$$0 = \lambda(v_L) - c_L - \lambda'(v_L)\frac{1 - F(v_L)}{f(v_L)}, \tag{17}$$

$$v_H - c_H - \frac{1 - F(v_H)}{f(v_H)} = \lambda(v_H) - c_L - \lambda'(v_H)\frac{1 - F(v_H)}{f(v_H)}. \tag{18}$$

If $v_L = v_H$, then the right-hand side of (17) must be nonpositive, and the left-hand side of (18) is zero, in accordance with (9). In order to ensure a unique solution for v_H, we need the following regularity condition, which also ensures that a solution to the first order conditions (17) and (18) yields a global profit maximum:[25]

$$v - \lambda(v) - [1 - \lambda'(v)]\frac{1 - F(v)}{f(v)} \text{ is increasing.}^{26} \tag{19}$$

23. Without loss of generality, we can assume that $a \leq v_L \leq v_H \leq b$, since values outside this range produce zero quantities.

24. Note that $[v_L, v_H]$ is a nontrivial interval if and only if $p_L < \lambda(p_H)$.

25. When $c_L < a$, a variety of other cases emerge. The possible solutions to these cases are as follows: (i) $a = v_L = p_1 = v_H$, (ii) $a = v_L = p_1 < v_H$, and (iii) $a = v_L < p_1 \leq v_H$. In case (i), the low quality good is not introduced. In case (ii), $p_H > p_1$, and high v types are worse off when the low quality good is introduced. Case (iii) requires further assumptions to make a Pareto comparison.

26. Provided f is differentiable, the assumption (19) is equivalent to the remarkably weak condition $(\partial/\partial v)\,[1 - \lambda'(v)][1 - F(v)]^2/f(v) < 0$. By (10), then, (19) holds if λ is convex, or not too concave.

Under these conditions, we have:

LEMMA 1 *Suppose eqs.* (7)–(19) *hold. Then introducing L is profitable if and only if*

$$\lambda(p_1) - c_L - \lambda'(p_1)\frac{1 - F(p_1)}{f(p_1)} > 0. \tag{20}$$

Henceforth, we will therefore assume that (20) holds. We now turn to the conditions under which introducing L produces a Pareto improvement. The next lemma provides some immediate insights into the welfare consequences of introducing good L.

LEMMA 2 *Suppose* (7)–(20) *hold. Then* $v_L < p_1 < v_H < b$.

Lemma 2 shows that, selling two qualities, the monopolist sells to more consumers but sells fewer units of the high quality good, relative to selling only the high quality. Note that introducing the low quality good makes consumers with valuations in (v_L, p_1) always strictly better off. Obviously, by (20), the monopolist benefits as well. A Pareto improvement therefore occurs if and only if $p_H \leq p_1$, for then high valuation customers are made better off as well.

The next result provides conditions sufficient to ensure that $p_H < p_1$, so that all market participants other than consumers with valuations below v_L (who do not get to purchase under either scenario) are strictly better off.

THEOREM 2 *Suppose that* (7)–(20) *hold, that* $[1 - F(v)]/f(v)$ *is nonincreasing, and that* $\lambda'(v)\,[1 - F(v)]/f(v)$ *is nondecreasing. Then* $p_H < p_1$, *that is, introducing the low quality good is a Pareto improvement.*

The hypotheses of Theorem 2, in conjunction with (7)–(21), are not vacuous, as the following example demonstrates.

Example 1: Let $F(v) = 1 - e^{-(v-a)/\alpha}$ for $v > a$, with $b = \infty$, and $\lambda(v) = \beta v + e^{-\beta v} + (1 - \beta)a - e^{-\beta a}$. The parameters are assumed to satisfy $a < c_H + \alpha$, $(1 + \alpha\beta)e^{-\beta a} < 1$, and

$$c_H \leq c_L < \beta[c_H + (1 + \alpha\beta)e^{-\beta(c_H+\alpha)}] + (1 - \beta)a - e^{-\beta a}.$$

These are satisfiable if β is near zero and $\alpha > 1$. All of the assumptions are strictly satisfied, and thus are robust to perturbations in a smooth C^1 metric.

Nevertheless, the conditions guaranteeing a Pareto improvement for the single use case are much more stringent than for the dual use case. Indeed, assumption (20) fails for many specifications of the environment, as the next result shows.

Lemma 3 *Suppose $\lambda(v)/v$ is nondecreasing. Then (20) fails.*

A sufficient condition for (20) to fail is that $\lambda(0) = 0$ and λ is convex, for that implies the hypothesis of Lemma 3. Recall that a convex λ is a sufficient condition for (19) to hold (see footnote 26). This makes (20) seem somewhat unnatural. However, when (20) fails, we won't ever observe a good damaged to produce a lower quality good. In spite of the large number of examples of the use of this strategy, clearly most goods sold are not intentionally damaged by the manufacturer. Thus, it may be that (20) only holds for a small percentage of all goods sold, but still on a large number of goods.

It might be countered that the conditions of Theorem 2 are sufficient rather than necessary, and hence that the above conclusion is premature. That this is not the case is shown by our next result. For "large" c_L, we have an exact characterization of the conditions under which a Pareto improvement occurs. Suppose f is continuously differentiable, and let $\bar{c}_L$ just make (20) fail, that is,

$$\bar{c}_L = \lambda(p_1) + \lambda'(p_1)\frac{1 - F(p_1)}{f(p_1)}.$$

Then we have

Theorem 3 *For c_L close to $\bar{c}_L$, one has $p_H < p_1$ if and only if λ is convex.*

Thus, convexity of λ is necessary for a general result. However, as shown in Lemma 3, convexity of λ tends to make introducing L unprofitable. The reader may wonder why the conditions for a Pareto improvement are so much more stringent in the single use case than in the dual use case. The intuition is as follows. In the dual use case, assumption (4) implies that lowering the price of good H below its monopoly price will not cannibalize profits from the L market. In the single use case, whenever the L good has positive market share, lowering the price of good H necessarily results in cannibalization.[27] This has two consequences. First, introducing the L good is now no longer necessarily profitable. Technically, the first order effect of a reduction in the price of good H is no longer zero (it is negative). Secondly, because of the cannibalization, the optimal response to a lower price of good L is now more likely to be a price increase, making a Pareto improvement harder to achieve.

27. The single use model also differs from the dual use model in that individual demand is inelastic (up to the reservation price). Introducing unit demand into the dual use model produces qualitatively similar results to the general dual use case; thus the distinction between the two cases appears to have more to do with discrete types versus a continuum of types than with downward-sloping demand.

5. CONCLUSION

For many products, it appears that the cost effective way to segment the market and price discriminate is to damage an existing product to produce a lower quality product, rather than improve the quality of an inferior product or produce the product separately. Although many of our examples are in chemicals and electronics, we think the practice is much more widespread. For example, outlet malls are often located much farther away from major cities than land prices would seem to dictate, as the rent gradient would appear to bottom out once low value use, such as farming, commences. It is often impossible to find the defects in apparel products labeled as "seconds" and sold in discount stores, as we recall from our graduate student days.[28]

We have argued that the phenomenon of manufacturers damaging goods naturally divides into two distinct cases based on customer characteristics. In one case, there are two uses, a high value use and a low value use, and this case seems best modeled as two distinct markets (Sec. 3). Most of the chemicals and pharmaceuticals fit this category, along with business versus educational uses of software. In this case, the conditions for price discrimination to be a Pareto improvement are not severe, and boil down to the assumption that, absent price discrimination, the high value market is more profitable than the low value market.

In the second case, there are not two separate markets for the products, but rather a group of consumers with distinct use values for the two products. Most of the electronics examples would seem to fit this category. This case seems more naturally modeled with a continuum of use values. The restrictions necessary for price discrimination to produce a Pareto improvement seem more severe and unnatural in this case.

In modeling the phenomenon, we have not endogenized the quality of the inferior good, mainly because of the resulting complexity of the mathematical description of preferences, which must be defined over all possible qualities, and firms' costs, which are now a function not only of quantity but also of quality. It is clear from some of the examples that quality is endogenous. For example, the slowdown of

28. The obliteration of manufacturers' tags for high fashion clothes by discount houses might seem like an example of damaging goods, but has an alternative, compelling explanation. Many upscale retailers, such as Neimann-Marcus, accept returns of high fashion clothes without a receipt. If the same item could be purchased unblemished at a discounter, it would not be possible for the upscale retailers to accept returns without proof of purchase, to the annoyance of their customers. Consequently, manufacturers rip labels or mark through them, as a signal that the item was sold by a discounter.

the IBM LaserPrinter was chosen by IBM. Characterizing the conditions for a Pareto improvement when quality is endogenous would appear to be an interesting, if daunting, research goal.

At least two of the damaged goods, the Intel 486SX and the IBM LaserPrinter E, appear to have been introduced in response to competition by another producer. This is difficult to explain, particularly in the LaserPrinter case. In this case, IBM's regular LaserPrinter was significantly faster than the Hewlett-Packard IIP, and thus the products were significantly differentiated. By introducing a product comparable to the IIP, IBM makes the market for the slower printers more competitive, thereby reducing prices for the slower printers. This should have the effect of further undercutting the market for the faster printer, as Hewlett-Packard responds to IBM's LaserPrinter E with a price cut.

It is possible as a theoretical matter that IBM's ability to punish Hewlett-Packard for competitive pricing is enhanced by the introduction of the similar product. Not only does IBM gain the ability to take a large percentage of HP's sales by aggressive pricing, but it limits the cost of such punishment by permitting higher prices on the faster machine. Thus, the set of equilibria to the repeated pricing game may grow, admitting some equilibria with higher profits than existed absent the introduction of the LaserPrinter E. Nevertheless, this seems an unlikely explanation for IBM's behavior.

An alternative explanation involves the perceived need of manufacturers to offer a full line of products. Offering a full line clearly makes consumers feel more comfortable, perhaps because the firm is less likely to exit, more likely to support the products with technical help and product updates, and more likely for the products themselves to be of high quality because of experience with the industry. Further research into price discrimination by imperfectly competitive firms seems warranted.

Appendix: Proofs

Proof of Theorem 1. From (3), if the firm sells both goods, it sells H to Y and L to X. Let p_H be the price of H and p_L be the price of L. The firm solves the following maximization problem:

$$\max_{p_H, p_L} (p_L - c_L)x_L(p_L) + (p_H - c_H)y_H(p_H) \tag{P}$$

subject to

$$\int_{p_H}^{\infty} y_H(p)\,dp \geq \int_{p_L}^{\infty} y_L(p)\,dp, \tag{IC$_y$}$$

$$\int_{p_H}^{\infty} x_H(p)\,dp \leq \int_{p_L}^{\infty} x_L(p)\,dp. \tag{IC$_x$}$$

Let (p_H^*, p_L^*) solve (P). First, note that $x_L(p_L^*) > 0$. Suppose by way of contradiction that $x_L(p_L^*) = 0$. Then $p_H^* = M_H^x = M_H^{xy}$, and by (5) we may assume that $p_L^* = \bar{p}_L$. Consider the following deviation: $p_L = p_L^* + \Delta p_L$, $p_H = p_H^* + \Delta p_H$. Choose $\Delta p_L < 0$, and let Δp_H be determined as follows. If IC_x does not bind, let $\Delta p_H = 0$. If IC_x binds, choose Δp_H so that IC_x holds with equality. In either case, by (3) and (5), IC_y is satisfied. This deviation increases profits when Δp_L is sufficiently small, for

$$\begin{aligned}\left.\frac{\partial \pi}{\partial p_L}\right|_{p_L=\bar{p}_L, p_H=M_H^y} &= [(M_H^y - c_H)y_H'(M_H^y) + y_H(M_H^y)]\frac{dp_H}{dp_L} \\ &\quad + (\bar{p}_L - c_L)x_L'(\bar{p}_L) + x_L(\bar{p}_L) \\ &= (\bar{p}_L - c_L)x_L'(\bar{p}_L) < 0.\end{aligned}$$

This contradicts the hypothesis that $(p_H^*, \bar{p}_L)$ solved (P).

That $x_L(p_L^*) > 0$ implies IC_x does not bind. For suppose it does. Then (3) implies IC_y does not bind, implying that $p_H^* = M_H^y$, which implies that IC_x does not bind by (4). There are two remaining possibilities: either IC_y does not bind, or it does [which occurs when (6) holds].

If IC_y does not bind, then $p_H^* = M_H^y$, so both X and the firm are better off, and Y obtains the same (monopoly) price as when the firm only offers one quality.

Now suppose IC_y binds, that is, (6) holds. It must therefore be the case that $p_H^* < M_H^y$, for $p_H^* \geq M_H^y$ and (6) imply $p_L^* > M_L^x$, and lowering both prices increases profits. Hence if (6) holds, a strict Pareto improvement occurs and all three agents benefit. □

Proof of Lemma 1. Suppose that

$$\lambda(p_1) - c_L - \lambda'(p_1)\frac{1 - F(p_1)}{f(p_1)} \leq 0,$$

and that $v_H > v_L$. Then by (17) and (13) we have $v_L \geq p_1$, so $v_H > p_1$. Now

$$\frac{\partial \pi}{\partial v_H} = f(v_H)\left[-\left(v_H - c_H - \frac{1 - F(v_H)}{f(v_H)}\right) + \left(\lambda(v_H) - c_L - \lambda'(v_H)\frac{1 - F(v_H)}{f(v_H)}\right)\right],$$

and by (19)

$$\frac{1}{f(v_H)}\frac{\partial \pi}{\partial v_H}(v_H) < \frac{1}{f(p_1)}\frac{\partial \pi}{\partial v_H}(p_1).$$

The hypothesis and (9) then imply that

$$\frac{\partial \pi}{\partial v_H}(p_1) \le 0,$$

contradicting that v_H is chosen optimally.

Conversely, suppose that

$$\lambda(p_1) - c_L - \lambda'(p_1)\frac{1 - F(p_1)}{f(p_1)} > 0,$$

and that $v_H = v_L$. Then $v_L = v_H = p_1$, and so

$$\frac{\partial \pi}{\partial v_L}(v_L) = -f(v_L)\left(\lambda(v_L) - c_L - \lambda'(v_L)\frac{1 - F(v_L)}{f(v_L)}\right) < 0,$$

contradicting that v_L is chosen optimally. □

Proof of Lemma 2. If $v_L = v_H$, then $v_H = p_1$, and by (20)

$$\frac{\partial \pi}{\partial v_L} = -f(p_1)\left(\lambda(p_1) - c_L - \lambda'(p_1)\frac{1 - F(p_1)}{f(p_1)}\right) < 0.$$

Thus $v_L < v_H$. By (13) and (17), $v_L < p_1$. Therefore, by (13) and (18),

$$v_H - c_H - \frac{1 - F(v_H)}{f(v_H)} = \lambda(v_H) - c_L - \lambda'(v_H)\frac{1 - F(v_H)}{f(v_H)} > \lambda(v_L) - c_L - \lambda'(v_L)\frac{1 - F(v_L)}{f(v_L)} = 0.$$

By (9) and (10), $v_H > p_1$. Finally, if $v_H = b$ and if $f(b) > 0$, then by (7) and (8)

$$\frac{\partial \pi}{\partial v_H}(b) = f(b)\left[-b + c_H + \lambda(b) - c_L\right] < 0.$$

If $v_H = b$ and $f(b) = 0$, then $\lim_{x \to b} [1 - F(x)]/f(x) = 0$ implies that

$$\frac{\partial \pi}{\partial v_H}(v_H) < 0$$

for some neighborhood around b. In either case, this contradicts the optimality of v_H. □

Proof of Theorem 2. By (9), (15), (17), (18),

$$\begin{aligned} p_H - p_1 &= v_H - \lambda(v_H) + \lambda(v_L) - p_1 \\ &= \frac{1 - F(v_H)}{f(v_H)}[1 - \lambda'(v_H)] + \lambda'(v_L)\frac{1 - F(v_L)}{f(v_L)} - \frac{1 - F(p_1)}{f(p_1)} \\ &= \left(\frac{1 - F(v_H)}{f(v_H)} - \frac{1 - F(p_1)}{f(p_1)}\right) \\ &\quad - \left(\lambda'(v_H)\frac{1 - F(v_H)}{f(v_H)} - \lambda'(v_L)\frac{1 - F(v_L)}{f(v_L)}\right) < 0, \end{aligned}$$

by the monotonicity assumptions and Lemma 3. □

Proof of Theorem 3. First note that (17) implies

$$\frac{\partial v_L}{\partial c_L} = \frac{1}{\lambda'(v_L)\left[2 + \frac{1 - F(v_L)}{f(v_L)^2} f'(v_L) - \frac{\lambda''(v_L)}{\lambda'(v_L)} \frac{1 - F(v_L)}{f(v_L)}\right]}$$

and

$$\frac{\partial v_H}{\partial c_L} = \frac{-1}{[1 - \lambda'(v_H)]\left[2 + \frac{1 - F(v_H)}{f(v_H)^2} f'(v_H) + \frac{\lambda''(v_H)}{1 - \lambda'(v_H)} \frac{1 - F(v_H)}{f(v_H)}\right]}.$$

Thus, evaluating at $c_L = \bar{c}_L$, we have $v_H = v_L = p_1$ and

$$\begin{aligned} \left.\frac{\partial p_H}{\partial c_L}\right|_{c_L = \bar{c}_L} &= [1 - \lambda'(p_1)]\frac{\partial v_H}{\partial c_L} + \lambda'(p_1)\frac{\partial v_L}{\partial c_L} \\ &= \frac{-1}{2 + \frac{1 - F(p_1)}{f(p_1)^2} f'(p_1) + \frac{\lambda''(p_1)}{1 - \lambda'(p_1)} \frac{1 - F(p_1)}{f(p_1)}} \\ &\quad + \frac{1}{2 + \frac{1 - F(p_1)}{f(p_1)^2} f'(p_1) - \frac{\lambda''(p_1)}{\lambda'(p_1)} \frac{1 - F(p_1)}{f(p_1)}}. \end{aligned}$$

By (13) and (19), this is positive if and only if $\lambda''(p_1) > 0$. That is, near $\bar{c}_L$, p_H falls as c_L is decreased if and only if λ is convex. □

Proof of Lemma 3. $\lambda(v)/v$ is nondecreasing if and only if $v\lambda'(v) \geq \lambda(v)$. Thus,

$$\lambda(p_1) - c_L - \lambda'(p_1)\frac{1 - F(p_1)}{f(p_1)} \leq \lambda'(p_1)\left(p_1 - \frac{1 - F(p_1)}{f(p_1)}\right) - c_L$$

$$= \lambda'(p_1)c_H - c_L < 0. \qquad \square$$

REFERENCES

DeLamarter, T., 1986, *BIG BLUE: IBM's Use and Abuse of Power*, New York: Dodd, Mead & Company, 161–175.

Fisher, F.M., J. McGowan, and J.E. Greenwood, 1983, *Folded, Spindled and Mutilated: Economic Analysis and U.S. vs. IBM*, Cambridge, MA: The MIT Press.

Frenkel, G., 1991, "For Intel, a 486 Chip by Any Other Name is Still the Same," *PC Week Supplement*, supplement to *PC Week*, 8(5), S37–S38.

Harley, R., 1992, "Industry Update: Japan," *Stereophile*, 15(11), 50–63.

Harley, R., 1994, "Sony MDS-501 Home Minidisc Recorder/Player & MZ-E2 Portable Minidisc Player," *Stereophile*, 17(9), 115–129.

Jones, M., 1990, "Low-Cost IBM LaserPrinter E Beats HP LaserJet IIP on Performance and Features," *PC Magazine*, 8, No. 10, May 29, 33–36.

Larson, M., 1993, "Packaging for Success in Warehouse Clubs," *Packaging*, 38(7), 26–28.

Nagle, T., 1987, *The Strategy and Tactics of Pricing: A Guide to Profitable Decision-Making* (p. 168). Englewood Cliffs, NJ: Prentice-Hall.

Seymour, J., 1991, "Intel: Is it Still 'The Computer Inside'?," *PC Week*, November 11, 73.

Smith, B. and A. Villamil, 1993, "Efficiency, Randomization and Commitment in Government Borrowing," Mimeo, Cornell University, June.

Stocking, G. and M. Watkins, 1947, *Cartels In Action: Case Studies in International Business Diplomacy* (pp. 402–405). New York: The Twentieth Century Fund.

Tirole, J., 1988, *The Theory of Industrial Organization*, Cambridge, MA: The MIT Press.

Varian, H., 1985, "Price Discrimination and Social Welfare," *American Economic Review* 75(4), 870–875.