Managing Digital Piracy: Pricing, Protection and Welfare

Arun Sundararajan

Department of Information, Operations and Management Sciences

Leonard N. Stern School of Business, New York University

44 West 4th Street, New York, NY 10012

asundara@stern.nyu.edu

May 2003

Center for Digital Economy Research
Stern School of Business
Working Paper IS-03-05

Managing Digital Piracy: Pricing, Protection and Welfare

Arun Sundararajan[1]

Leonard N. Stern School of Business, New York University
44 West 4th Street, KMC 8-93, New York, NY 10012
asundara@stern.nyu.edu

May 2003

Abstract: This paper analyzes the optimal choice of pricing schedules and technological deterrence levels in a market with digital piracy, when legal sellers can sometimes control the extent of piracy by implementing digital rights management (DRM) systems. It is shown that the seller's optimal pricing schedule can be characterized as a simple combination of the zero-piracy pricing schedule, and a *piracy-indifferent* pricing schedule which makes all customers indifferent between legal consumption and piracy. An increase in the level of piracy is shown to lower prices and profits, but may improve welfare by expanding the fraction of legal users and the volume of legal usage. In the absence of price-discrimination, the optimal level of technology-based protection against piracy is shown to be the *technologically-maximal* level, which maximizes the difference between the quality of the legal and pirated goods. However, when a seller can price-discriminate, it is always optimal for them to choose a strictly *lower* level of technology-based protection. Moreover, if a DRM system weakens over time, due to its technology being progressively hacked, the optimal strategic response may involve either increasing or decreasing the level of technology-based protection and the corresponding prices. This direction of change is related to whether the technology implementing each marginal reduction in piracy is increasingly less or more vulnerable to hacking. Pricing and technology choice guidelines based on these results are presented, and some social welfare issues are discussed.

[1] I thank Ravi Mantena for his careful reading of this paper and his detailed comments. I also thank seminar participants at New York University for their feedback on an earlier version of this paper. All errors and omissions remain mine.

1. Introduction

Over the last decade, sellers of digital products have actively fought the availability of pirated copies of their products. Nevertheless, digital piracy rates are still high and increasing in many markets, despite a continuous increase in the availability and sophistication of copy protection and digital rights management technologies. Piracy concerns have expanded following the emergence of file-sharing networks like Gnutella and Kazaa, which have substantially increased the availability and exchange of high-quality illegal versions of software, music and digital video.

The sustained presence of piracy complicates the design of pricing schedules for sellers of digital goods. It also poses the new challenge of choosing an appropriate level of technology-based protection, and of responding strategically to the hacking of existing digital rights management systems. These are the issues addressed by the model in this paper.

The first part of the paper studies optimal pricing strategy. When faced with rising digital piracy, a seller's pricing power is increasingly limited by the quality and availability of pirated copies, which are imperfect substitutes for the legally available product. This effect of piracy is analyzed in a model of monopoly price discrimination, when valuations of both legal products and illegal copies *differ* across customers. It is shown that the seller's optimal pricing schedule can be characterized as a simple combination of two contracts – the optimal pricing schedule in the absence of piracy (termed the *zero-piracy pricing schedule*), and a *piracy-indifferent pricing schedule,* which makes all customers indifferent between legal usage and piracy. An increase in the quality of pirated goods lowers prices and profits; however, there may be social benefits from piracy realized through the expansion of the fraction of legal users and the volume of legal usage.

The next part of this paper studies technology-based protection against digital piracy, which is typically achieved by implementing *digital rights management* (henceforth referred to as DRM) systems. Examples of DRM systems for digitally delivered products include Macrovision SafeCast, Microsoft Windows Media DRM Series, and Real Networks Helix DRM. Other DRM systems aimed specifically at protecting the physical sources of the digital files shared illegally over the Internet include the Windows Media Data Session Toolkit, and Macrovision's Cactus Data Shield.

Implementing a DRM system enables a seller to deter and control piracy to some extent. Unfortunately, as a DRM system becomes more effective, it often simultaneously places more restrictions on the flexibility of usage for a legal user, and thereby reduces the value of the *legal product.* For instance, in 2002, a number of music labels (most notably, Sony) introduced protection technology that would prevent their audio CD's from being played on personal computers[2]. This reduced value by restricting inter-device portability; moreover, a substantial fraction of discs did not work on regular CD and DVD players[3]. Many online music services implement DRM by limiting the rendering of their MP3 files to a single device, and by placing related restrictions on the portability of these files[4]. Highly restricted services like MusicNet and Rhapsody have not been well-received, and this is partly attributed to the fact that their protection schemes "...treat everyone like a potential criminal, and they take all the joy out of buying and playing music" (Mossberg, 2003). In contrast, the iTunes music service from Apple, which places substantially fewer restrictions[5] on their customers' ability to download, share and burn purchased MP3 files, has been a resounding success in its first few weeks.

There are more direct ways in which implementing technology-based protection may necessitate degrading the value of a legal product. Currently, textual content that is protected by Adobe's DRM partners can be electronically scanned by OCR software that takes PDF files as direct inputs and produces near-perfect scanned versions. Countering this form of piracy will necessitate degrading fonts in rendered legal files, thereby lowering quality for legal users. Protecting digital video or music files using encryption makes the corresponding files larger, which may lower value by increasing download times for digitally delivered content.

These observations suggest that when choosing their DRM system and the corresponding level of technology-based protection against piracy, the seller of a digital good needs to trade-off the effectiveness of *deterring piracy* with the *value reduction* of the legal product that is caused by the implementation of the DRM system. To study this trade-off, the model of pricing with digital piracy developed in this paper is extended to incorporate endogenous choices of technology-based protection. The *technologically-maximal level* of protection, which is the level of technology-based protection at which the quality *difference* between the legal good and the pirated good is maximized, is contrasted with the *profit-maximizing level* of protection. When the seller can price-discriminate, the profit-maximizing level of protection is shown to be *strictly lower* than the technologically-

[2]Sony's Key2Audio system worked by adding a bogus data track to each audio CD. Since computer hard drives are programmed to read data files first, the technology relied on PC's continuously attempting to play the bogus track, and never getting to the music files. The system failed, however, when it became known that simply blacking out the outermost track of the CD with a marker could override the protection.

[3]For instance, an early release made under the Cactus format in Germany reportedly had a 4% return rate from people who found that these CDs didn't work on their normal CD players.

[4]The most recent version of MusicNet, launched on AOL in February 2003, does not permit transferring MP3 files to portable players, and limits the number of files that can be burnt onto a CD.

[5]iTunes allows users to burn their MP3 files to an unlimited number of CD's, copy them to an unlimited number of iPod MP3 players, and play them on upto three Macintosh computers.

maximal level. The economic drivers of this result are explored in some detail, since it indicates that even after accounting for the quality degradation imposed on legal products, a DRM system that maximizes the quality gap between the legal good and the pirated good is overprotecting the digital product. As the effectiveness of a DRM system weakens over time (which typically occurs due to its technology being progressively hacked), the optimal technological and pricing responses for the seller are examined. It is shown that the seller's optimal response may either be to decrease or to increase their level of technology-based protection, and conditions under which each of these responses is optimal are characterized.

Related work: A distinguishing aspect of this paper is its formal analysis of the economic effects of DRM technologies, and their role as technological deterrents to piracy that are controlled explicitly (and strategically) by a seller. Recent related work on optimal deterrence includes Chen and Png (2001) who study monopoly investment levels in piracy detection (an indirect form of deterrence), and Yoon (2001) who provides a simple model of the optimal level of protection chosen by a monopolist. Related policy issues are also examined by Png and Chen (2003) who consider how a regulator should balance taxing copying and subsidizing legitimate purchases when a monopolist invests in piracy detection, and by Gopal and Sanders (1998) who examine how governmental incentives to enforce copyright laws are related to the size of the domestic software industry. Other work on piracy-related P2P technologies include Duchene and Waelboeck (2001) who model the effects of these technologies on new-product introduction, and Gayer and Shy (2002) who explore their role as a marketing vehicle to spur in-store sales. An early model of piracy deterrence is Conner and Rummelt (1991) who established that increases in protection technology always increase firm profits, unless the product displays positive network effects.

This paper builds on and adds to this literature significantly, providing a richer model of the economics of technology-based piracy deterrence, and deriving new results on choosing optimal protection strategies. For instance, it is shown that even in the absence of network effects, increases in protection technology can reduce firm profits, and the right response to an increase in the level of piracy may be a weakening of protection levels, rather than a tightening of them.

This paper also contributes to the literature on the economics of copying[6] and piracy, which has seen a significant resurgence in the wake of Napster and other Internet-based piracy threats. Some of the recent research in this area is surveyed and extensively supplemented in Watt (2000).

[6]This is part of a larger body of work on the economics of intellectual property and copyright. A good guide to the literature on the economics of copyright can be found in Besen and Raskind (1991). The papers by Liebowitz (1985) and Johnson (1985) are notable as early examples of the economics of piracy and copying technology, in contrast with much of the literature's focus on socially optimal copyright policy.

More recently, Alvisi, Argentisi and Carbonara (2002) analyze whether the presence of a pirated good may make it optimal for a monopolist to introduce a lower-quality substitute for their product. Belleflamme (2002) studies the interdependence between different producers' incentives to accommodate/deter the presence of a pirated good. Ben-Shahar and Jacob (2002) examine the entry-deterring properties of piracy, characterizing a monopolist's incentives to encourage piracy as a form of predatory pricing. Chellappa and Shivendu (2002) model pricing when buyers bear a 'moral cost' from using a pirated good derived from an evaluation version of the legal good.

The model in this paper builds on the approach of each these papers, by preserving their notion of the pirated good as an inferior (vertically differentiated) substitute for the legal good. However, it generalizes their pricing analysis significantly, by modeling and deriving a continuous pricing schedule (rather than a single variable price, or a pair of prices for two quality-differentiated products) which explicitly takes into account the *differing* value of pirated products to *different* customers. This generalization substantially alters results relating to the optimal level of technological protection, and to post-implementation protection and pricing trends – differences that would not be evident in a model with unit consumption and no price discrimination. Furthermore, results from this paper are more likely to provide managerially relevant pricing guidelines, by prescribing a straightforward way to actually design pricing schedules in the presence of digital piracy. For instance, different users of pay-per-download services like the iTunes music service from Apple download different numbers of songs, and have a different willingness to pay for each additional download, but pay a fixed (linear) price per download. If Apple wants to use price discrimination to generate higher revenues from their service, they would need guidelines for the design of a usage-based pricing schedule in the presence of a piracy threat.

The model also adds to the growing literature on price screening in technology markets, in which theories developed by Naor (1969), Mussa and Rosen (1978) and Maskin and Riley (1984) have been adapted and applied to problems unique to technology markets. Nault (1997), Jones and Mendelson (1998), Konana, Gupta and Whinston (2000), Bhargava and Choudhary (2001), Bhargava and Sundaresan (2002), Jing (2002), Sundararajan (2002), and Hosanagar et al. (2003) are good examples. Unlike the piracy models in Conner and Rummelt (1991), Takeyama (1994) and Shy and Thisse (1999), positive network externalities are not considered. These externalities are quite important in software industries; however, it is unlikely that there are substantial direct network effects in industries more recently threatened by digital piracy[7] – music, video and content.

[7]In some cases, there may be indirect network effects (the increase in the value of portable MP3 players that has resulted from the wide availability of pirated music, for instance), though it seems likely that these are appropriated

The rest of this paper is organized as follows. Section 2 provides an overview of the model, and describes the seller's optimal pricing schedule in the absence of piracy. Section 3 characterizes optimal pricing strategy in the presence of different levels of digital piracy. Section 4 models the economic effects of using digital rights management systems, and prescribes optimal technology-based protection levels, as well as strategic responses to changes in the effectiveness of an implemented DRM system. Section 5 discusses some results of the paper, and concludes with an outline of the open research questions.

2. Model

2.1. Seller and customers

The model involves an information good which may be used by consumers in continuously varying quantities[8]. The seller of this good (termed the *legal good*) is assumed to be a monopolist, by virtue of owning a copyright over the information good. Any fixed costs of production or IP protection are assumed to be sunk. Since the product is an information good, variable costs of production are zero. In addition to the legal good, there is also a *pirated good* available, which is a lower-quality substitute for the legal good, and is free.

Customers are heterogeneous, indexed by their type $\theta \in [0, \overline{\theta}]$. The preferences of a customer of type θ for the *legal good* are represented by the utility function $vU(q, \theta)$, where q is the quantity of the legal good used by the customer, and $v \leq 1$ is a measure of the quality of the legal good. The function $U(q, \theta)$ is assumed to take the following form:

$$U(q, \theta) = (\theta + w)q - \frac{1}{2}q^2. \tag{2.1}$$

Based on equation (2.1), it is clear that a higher customer type θ gets higher marginal value from the legal good, at every level of usage. The differences in value across customer types are mediated to some extent by the variable w, which is termed the common marginal value of the good.

Numbered subscripts of functions represent partial derivatives with respect to the corresponding variable. For instance, the partial derivative of $U(q, \theta)$ with respect to q is denoted $U_1(q, \theta)$, and

by parties other than the seller. Recent results from Sundararajan (2003) have presented a framework for analyzing network effects in nonlinear pricing models – a possible extension is discussed in Section 5.

[8] This could either be a homogeneous good, or a bundle of related goods (such as a library of digital information, music or video), from which different customers use different goods. For instance, Apple's iTune music service offers access to a library of over 200,000 songs. Each user downloads a small fraction of this library, and different subscribers download different songs. In the context of this model, the unit of quantity would therefore be number of downloads.

the cross-partial of $U(q,\theta)$ with respect to q and θ is denoted $U_{12}(q,\theta)$. This notation is preserved throughout the paper.

The following properties of the utility function follow from equation (2.1):

1. For every θ, $U(q,\theta)$ has a finite maximum usage $\alpha(\theta) = \arg\max_q U(q,\theta) = (\theta + w)$.
2. $U_1(q,\theta) > 0$ for $q < \alpha(\theta)$, and $U_1(q,\theta) < 0$ for $q > \alpha(\theta)$.
3. $U_{11}(q,\theta) = -1$, and $U_{12}(q,\theta) = 1$. Therefore, $U(q,\theta)$ is strictly concave in q (diminishing marginal value from usage), and has the Spence-Mirrlees single crossing property.

The preferences of a customer of type θ for the *pirated good* are represented by the utility function $sU(q,\theta)$, where q is the quantity of the pirated good used by the customer, and $U(q,\theta)$ is as defined in (2.1). The parameter s is a measure of the quality of the pirated good, and is also referred to as the *level of piracy* (based on the fact that at a higher level of piracy, the quality of the pirated good is higher[9]). s is assumed to be strictly less than v, implying that the pirated good is always strictly inferior to the legal good, and therefore, the seller can make a non-zero profit[10].

The maximum value that a customer of type θ can get from a pirated good of quality s is denoted $\hat{u}(\theta, s)$:

$$\hat{u}(\theta, s) = sU(\alpha(\theta), \theta) = \frac{s(\theta + w)^2}{2}. \tag{2.2}$$

Since the pirated goods are free, $\hat{u}(\theta, s)$ is the *reservation utility* of customer type θ.

The monopolist *does not observe* the type θ of any customer, but knows $F(\theta)$, the probability distribution of types in the customer population, which has the following properties:

1. $f(\theta) > 0$ for all θ, where $f(\theta)$ is the density corresponding to the distribution $F(\theta)$.
2. $\frac{\partial}{\partial\theta}(\frac{1-F(\theta)}{f(\theta)}) \leq 0$ for all θ: the inverse hazard rate is non-increasing in θ.

For expositional simplicity, and since the hazard rate of the customer type distribution plays a significant role in subsequent analysis, we define the *inverse hazard rate* function $h(\theta)$:

$$h(\theta) = \frac{1 - F(\theta)}{f(\theta)}, \tag{2.3}$$

[9]The 'level' of piracy referred to here is distinct from the *piracy rate*. The former measures the extent to which the legal good can be pirated (and therefore determines the quality of the pirated good, perhaps indirectly through its availability as well), while the latter is a measure of how many customers are actually using the pirated good.

[10]Depending on the context, the parameter v may be normalized to 1, or may be replaced by the function $v(\rho)$. In addition, the parameter s may be replaced by the function $s(\rho)$, or the function $s(\rho, t)$. Each of these functions will be explained in the appropriate section.

and the *cumulative inverse hazard rate* function $H(\theta)$:

$$H(\theta) = \int_0^{\theta} \frac{1-F(x)}{f(x)} dx. \tag{2.4}$$

Each customer knows their own type θ. Without any loss in generality, the total number of customers in the market is normalized to 1.

2.2. Customer choice and pricing schedules

The seller offers a nonlinear pricing schedule (sometimes referred to as either a *contract* or a *pricing schedule*) which assigns a non-negative price to each feasible level of usage for the *legal good.* Rather than considering all possible pricing functions, the revelation principle ensures that we can restrict our attention to *direct mechanisms* – menus of quantity-price pairs $q(x), \tau(x)$, indexed by $x \in [0, \overline{\theta}]$, which are *incentive-compatible.* A pricing schedule $q(x), \tau(x)$, $x \in [0, \overline{\theta}]$, for the legal good is said to be incentive-compatible if it satisfies:

$$\theta = \arg\max_x [vU(q(x), \theta) - \tau(x)], \text{ for all } \theta.$$

Given a pricing schedule $q(x), \tau(x)$ for the legal good, customers of type θ choose to purchase the legal good if their surplus from doing so is at least as much as the value they would derive from the (free) pirated good. Mathematically, if:

$$\max_x [vU(q(x), \theta) - \tau(x)] \geq \hat{u}(\theta, s), \tag{2.5}$$

then customers of type θ purchase the legal good. Therefore, an incentive-compatible pricing schedule $q(x), \tau(x)$ is said to *induce participation* from customer type θ if all customers of this type (weakly) prefer the legal good to the pirated good:

$$[vU(q(\theta), \theta) - \tau(\theta)] \geq \hat{u}(\theta, s). \tag{2.6}$$

The constraint (2.6) above is often referred to as the *piracy constraint* for type θ, since it becomes progressively harder to satisfy as s increases. Note that the piracy constraint is type-dependent[11]. In the special case of $s = 0$, since the pirated good has no value, $\hat{u}(\theta, s) = 0$, and constraint (2.6)

[11] This is in contrast with standard price discrimination or adverse selection models, where the RHS of the participation (individual rationality) constraint is normalized to zero across all types.

above reduces to the standard individual rationality constraint of price discrimination problems.

Finally, the *optimal pricing schedule* $q^*(x,v,s), \tau^*(x,v,s)$ is the incentive compatible pricing schedule that maximizes the seller's profits[12] .

Notation used most frequently is summarized in Table 2.1. In general, the sequence of events is as follows: the seller announces their pricing schedule (and technological choices, if any), the customers make their purchase decisions (whether to use the legal good or the pirated good, and at what usage level) based on the pricing schedule, and each party gets their payoffs. An exact timeline is specified separately in each of the following sections.

2.3. Optimal pricing schedule in the absence of piracy

The optimal pricing schedule in the absence of piracy, termed the *zero-piracy pricing schedule*, is specified in this section. The zero-piracy pricing schedule benchmarks the analysis of pricing in the presence of piracy, and is also used in constructing the corresponding optimal pricing schedules.

Lemma 1. *The zero-piracy pricing schedule $q^{ZP}(\theta,v), \tau^{ZP}(\theta,v)$, which is the optimal pricing schedule for the seller when $s=0$, takes one of the following two forms.*

(a) If $h(0) \leq w$, then the pricing schedule is designed to include all customer types The optimal contract is:

$$q^{ZP}(\theta,v) = \theta + w - h(\theta); \tag{2.7}$$

$$\tau^{ZP}(\theta,v) = \frac{v[w^2 - h(\theta)^2]}{2} + vH(\theta), \tag{2.8}$$

for all $\theta \in [0,\overline{\theta}]$.

(b) If $h(0) > w$, then a set $[0,\theta_{ZP}]$ of customer types are priced out of the market, where θ_{ZP} is defined as:

$$\theta_{ZP} = \theta : h(\theta) = w + \theta,\ \theta \in (0,\overline{\theta}). \tag{2.9}$$

The optimal contract is:

$$q^{ZP}(\theta,v) = \theta + w - h(\theta); \tag{2.10}$$

$$\tau^{ZP}(\theta,v) = \frac{v[(\theta_{ZP} + w)^2 - h(\theta)^2]}{2} + v[H(\theta) - H(\theta_{ZP})]. \tag{2.11}$$

[12]The necessary details about mechanism design and the revelation principle can be found in any standard graduate-level textbook on game theory. For example, chapter 7 of Fudenberg and Tirole (1991) describes the revelation principle and a non-linear pricing application.

Symbol	Explanation
θ	Index of customer types. $\theta \in [0, \overline{\theta}]$.
$f(\theta), F(\theta)$	Density and distribution functions of the customer types θ.
$vU(q,\theta)$	Value that customer type θ gets from a usage level q of the *legal good*.
$\alpha(\theta)$	Usage level which maximizes $U(q,\theta)$. $\alpha(\theta) = \arg\max_{q} U(q,\theta)$.
$sU(q,\theta)$	Value that customer type θ gets from a usage level q of the *pirated good*.
$\hat{u}(\theta, s)$	Maximum value that customer type θ can get from the pirated good. By definition, $\hat{u}(\theta, s) = sU(\alpha(\theta), \theta)$.
$q(x), \tau(x)$	Generic representation of a pricing schedule as a continuum of quantity-price pairs, indexed by $x \in [0, \overline{\theta}]$. Under this pricing schedule, a customer who buys a quantity $q(x)$ is charged a total price $\tau(x)$.
$h(\theta)$	Inverse hazard rate function. $h(\theta) = \frac{1-F(\theta)}{f(\theta)}$.
$H(\theta)$	Cumulative inverse hazard rate function. $H(\theta) = \int_0^\theta h(x)dx$.
$q^*(\theta, v, s), \tau^*(\theta, v, s)$	Optimal pricing schedule for the legal good,. This is the profit-maximizing incentive compatible pricing schedule which induces participation from all customer types allocated a non-zero usage level.
$q^{ZP}(\theta, v), \tau^{ZP}(\theta, v)$	Zero-piracy pricing schedule, which is the optimal pricing schedule in the absence of piracy
$q^{PI}(\theta, v, s), \tau^{PI}(\theta, v, s)$	Piracy-indifferent pricing schedule. Under this incentive-compatible pricing schedule, every customer of type θ gets equal value from the legal product and the pirated product.
$s(\rho)$	Quality level of the *pirated good* as a function of the level of digital rights management technology ρ.
$v(\rho)$	Quality level of the *legal good* as a function of the level of digital rights management technology ρ.
$\hat{\theta}(\rho)$	At a level of protection ρ, the customer type in $[0, \overline{\theta})$ which defines the transition from the piracy-indifferent portion of the contract. Types $\theta \leq \hat{\theta}(\rho)$ have usage levels defined by $q^{PI}(\theta, v(\rho), s(\rho))$, while types $\theta \geq \hat{\theta}(\rho)$ have usage levels defined by $q^{ZP}(\theta, v(\rho))$
$s(\rho, t)$	Quality level of the *pirated good* at time t, as a function of the level of digital rights management technology ρ.

Table 2.1: Summary of key notation

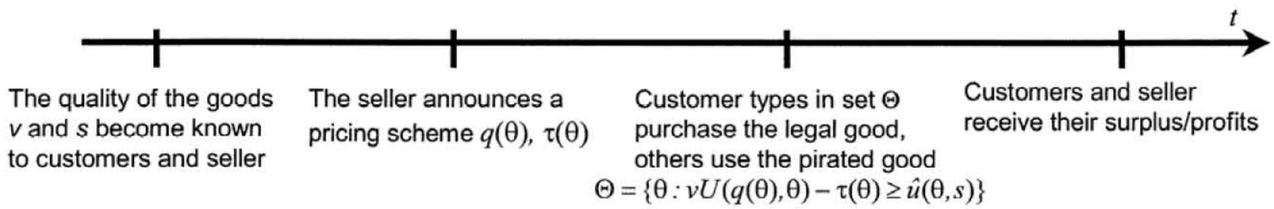

Figure 3.1: Timeline of events for Section 3

for $\theta \in [\theta_{ZP}, \overline{\theta}]$, *and*

$$q^{ZP}(\theta, v) = 0,\ \tau^{ZP}(\theta, v) = 0, \tag{2.12}$$

for $\theta \in [0, \theta_{ZP}]$.

Unless otherwise specified, all proofs are available in Appendix A.

3. Pricing with digital piracy

This section analyzes pricing strategy when the seller faces digital piracy. The sequence of events modeled in this section is summarized in Figure 3.1.

3.1. Piracy-indifferent pricing schedule

This section specifies the incentive-compatible pricing schedule that implements *piracy-indifference*. Under this pricing schedule, all customer types are exactly indifferent between the legal good and the pirated good. This pricing schedule is important because it often forms a building block for the optimal pricing schedule.

Proposition 1. *The unique incentive-compatible piracy-indifferent pricing schedule* $q^{PI}(\theta, v, s), \tau^{PI}(\theta, v, s)$ *for the legal good takes the following form:*

$$q^{PI}(\theta, v, s) = \frac{s(\theta + w)}{v}; \tag{3.1}$$

$$\tau^{PI}(\theta, v, s) = \left(\frac{s(v-s)}{v}\right)\frac{(\theta + w)^2}{2}. \tag{3.2}$$

Under this pricing schedule, each customer type get the same surplus from their optimal usage of the legal good of quality v *and their maximal usage of the pirated good of quality* s.

Proposition 1 establishes that there is a unique piracy-indifferent pricing schedule, under which each customer type gets a net surplus exactly equal to their reservation utility – the value $\hat{u}(\theta, s)$ that they would get from their maximal usage of the pirated good. From equation (3.1), all customer types purchase positive quantities of the legal good under this pricing schedule. Moreover, their usage levels of the legal good are strictly increasing in the level of piracy s. In addition, (3.2) indicates that so long as $s < v$, the total payment $\tau^{PI}(\theta, v, s)$ from each customer type θ is strictly positive. This establishes that the piracy-indifferent pricing schedule is strictly profitable for the seller.

3.2. Optimal pricing schedule for lower levels of digital piracy

The structure of the seller's optimal pricing schedule depends on the level of digital piracy s. This section describes how to design the optimal pricing schedule when the level of piracy is relatively low. Subsequently, Section 3.3 describes the solution to the corresponding problem at higher levels of digital piracy.

Proposition 2. *At lower levels of digital piracy – that is, when $s \leq \dfrac{v(w - h(0))}{w}$ – the sellers optimal pricing schedule is a modified version of the zero-piracy pricing schedule, with total prices adjusted downwards by the same amount across all usage levels. The optimal contract is:*

$$q^*(\theta, v, s) = q^{ZP}(\theta, v); \tag{3.3}$$

$$\tau^*(\theta, v, s) = \tau^{ZP}(\theta, v) - \frac{sw^2}{2}, \tag{3.4}$$

for all $\theta \in [0, \overline{\theta}]$, where $q^{ZP}(\theta, v)$ and $\tau^{ZP}(\theta, v)$ are as defined in equations (2.7) and (2.8).

Proposition 2 establishes that at lower levels of piracy[13], the optimal pricing schedule is simply the zero-piracy pricing schedule, with a constant reduction in total price across all usage levels. The resulting usage levels of all consumers are unaffected by the presence of piracy, and the reduction in total price across all customers is proportionate to the level of piracy s. An immediate corollary is that as the level of piracy s increases, prices are strictly lower at all usage levels.

The functions $q^{ZP}(\theta, v)$ and $\tau^{ZP}(\theta, v)$, derived in part (a) of Lemma 1, form the optimal pricing schedule when the seller is unconstrained by piracy. As the level of piracy s increases, the

[13]Clearly, the condition $s \leq \dfrac{v(w - h(0))}{w}$ of Proposition 2 does not just depend on the level of piracy s, but also depends on v, w, and $h(0)$. The statement 'lower levels of piracy' is meant to indicate that for a fixed distribution, and fixed values of v and w, the proposition is more likely to apply at lower levels of s.

reduction in total price lowers the seller's profits. In addition, under the optimal pricing schedule, all customer types get a surplus level which is *higher* than their reservation utility $\hat{u}(\theta, s)$ (which was also the surplus to each customer under the piracy indifferent pricing schedule). This surplus increase is due to the seller's desire to increase profits beyond the level obtained under the piracy-indifferent contract, by inducing higher usage across all customer types. Higher usage is necessarily accompanied by an increase in surplus for all types, in order to ensure incentive-compatibility.

3.3. Optimal pricing schedule for higher levels of digital piracy

In contrast with the previous result, this section describes the structure of the optimal pricing schedule at higher levels of digital piracy.

Proposition 3. *At higher level of piracy s – that is, when $s > \dfrac{v(w-h(0))}{w}$ – the seller's optimal pricing strategy is as follows:*

(a) Customer types are partitioned into two sets $[0, \hat{\theta}]$ and $[\hat{\theta}, \overline{\theta}]$, where the transition type $\hat{\theta}$ is defined by:

$$\hat{\theta} = \theta : vh(\theta) = (v - s)(w + \theta), \; \theta \in (0, \overline{\theta}). \tag{3.5}$$

(b) The optimal pricing schedule for the lower set of customers is simply the piracy-indifferent pricing schedule:

$$q^*(\theta, v, s) = q^{PI}(\theta, v, s); \tag{3.6}$$

$$\tau^*(\theta, v, s) = \tau^{PI}(\theta, v, s), \tag{3.7}$$

for $\theta \in [0, \hat{\theta}]$.

(c) The optimal pricing schedule for the higher set of customers is an adjusted version of the zero-piracy pricing schedule, with total prices adjusted downwards by the same amount across all usage levels:

$$q^*(\theta, v, s) = q^{ZP}(\theta, v); \tag{3.8}$$

$$\tau^*(\theta, v, s) = \tau^{ZP}(\theta, v) - \left(\frac{s(\hat{\theta} + w)^2}{2} + v[H(\hat{\theta}) - \frac{\hat{\theta}^2 + 2\hat{\theta}w}{2}] \right), \tag{3.9}$$

for $\theta \in [\hat{\theta}, \overline{\theta}]$, where $\tau^{ZP}(\theta, v)$ is as defined in part (a) of Lemma 1, in equation (2.8).

Note that the pricing schedule $q^*(\theta, v, s), \tau^*(\theta, v, s)$, while specified independently for each cus-

tomer set, is incentive-compatible across the entire set of customer types $[0, \overline{\theta}]$.

Proposition 3 establishes that at higher levels of piracy, the portion of the optimal pricing schedule which is relevant to a lower set of customer types $[0, \hat{\theta}]$ is simply the piracy-indifferent pricing schedule. Since $q^{PI}(\theta, v, s) > 0$, all these customer types purchase positive levels of the legal good. It can be shown that $\hat{\theta} > \theta_{ZP}$, and therefore any customer type who did not purchase in the absence of piracy is now a legal user, at their piracy-indifferent usage level $q^{PI}(\theta, v, s)$.

Therefore, the presence of piracy can have the socially beneficial effect of inducing *legal* usage from customers who may have otherwise been excluded by the seller's optimal price discrimination. While counter-intuitive, this result has a straightforward economic explanation. In the absence of piracy, there is no imperfect substitute for the seller's legal good, since the seller is a monopolist, and the reservation utility of all customers is zero. Under this scenario, when optimally price-discriminating, the seller finds it more favorable to capture a higher level of surplus from the customer types $\theta > \theta_{ZP}$, at the cost of excluding customer types $[0, \theta_{ZP}]$ from the market. The only reason why customer types $\theta \in [0, \theta_{ZP}]$ are excluded is because the seller's optimal surplus extraction from higher customer types would not be feasible if there was any positive usage level affordable for customer types $\theta \in [0, \theta_{ZP}]$. In the presence of piracy, each customer type who purchases the legal good must be provided with positive surplus of at least $\hat{u}(\theta, s)$ – their value from maximal usage of the pirated good. Since the seller is forced to provide this surplus level to the higher set, customer types in the lower set can now be offered positive and affordable usage levels, without affecting incentive-compatibility (and the seller's price-discrimination objectives).

In addition, since it is straightforward to establish that $q^{ZP}(\theta, s) < q^{PI}(\theta, s)$ for $\theta \in [0, \hat{\theta}]$, total usage either remains constant or goes up for all customer types, relative to the usage levels under the zero-piracy contract, and this increase is more pronounced at higher levels of piracy s. As a consequence, the total value $vU(q^*(\theta, v, s), \theta)$ created by the usage of each customer type also increases, which in turn implies that total surplus is higher at higher levels of s. These observations are illustrated further in Figure 3.2.

3.4. Example

In this section, the optimal pricing schedule is derived explicitly for a specific family of customer type distributions. This example further illustrates the effects of digital piracy on pricing, usage, and welfare, and highlights the effect of some properties of the customer type distribution.

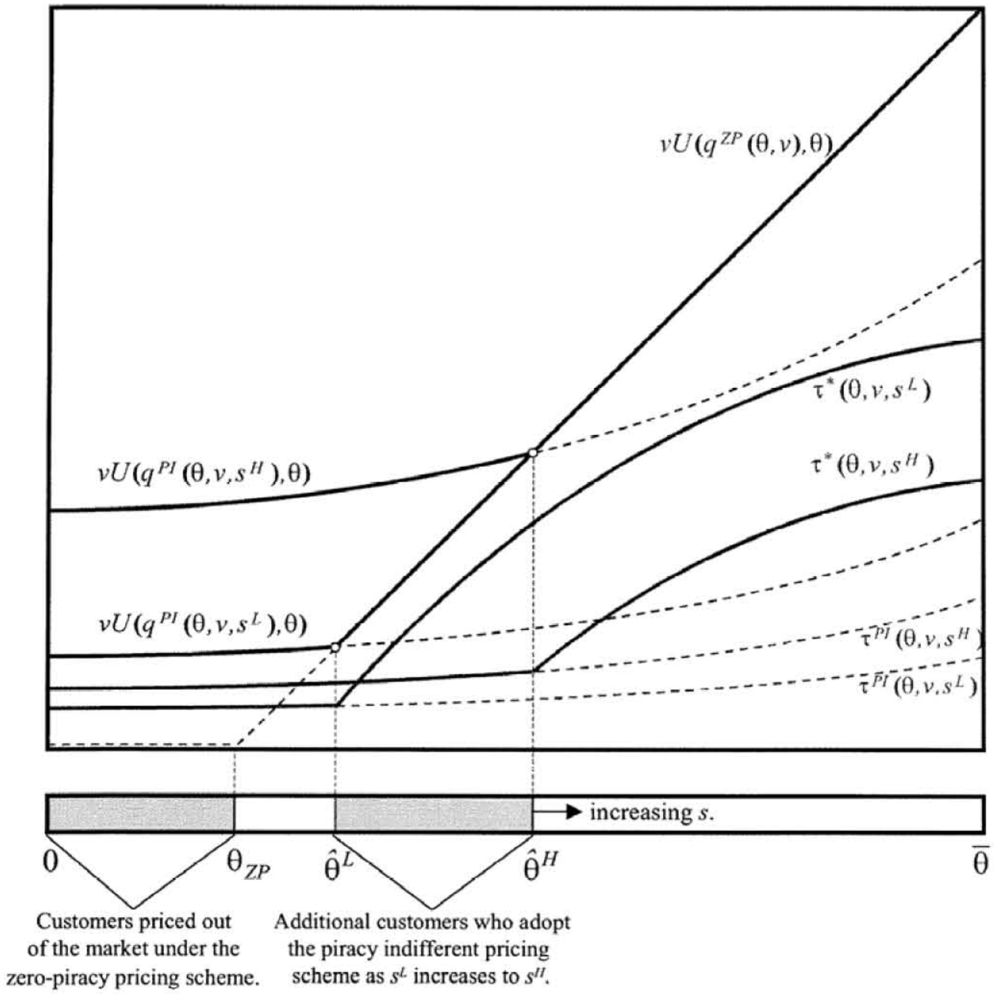

Figure 3.2: **(This figure is clearer when viewed electronically, or in color, since the curves are color-coded).** Illustrates the changes in pricing and surplus as the quality of the pirated good s changes For $s = s^L, s^H$, the piracy-indifferent pricing schedules are $q^{PI}(\theta, v, s), \tau^{PI}(\theta, v, s)$, and the total surplus generated by the usage of customer type θ is $vU(q^{PI}(\theta, v, s), \theta)$. The difference between $vU(q^{PI}(\theta, v, s), \theta)$ and $\tau^{PI}(\theta, v, s)$ is the minimum surplus that type θ must be provided in order to induce them to purchase the legal good, rather than using the pirated good. The thicker curves