

ONLINE PIRACY AND THE “LONGER ARM” OF ENFORCEMENT*

Debabrata Dey
(ddey@uw.edu)

University of Washington, Foster School of Business, Seattle, WA 98195–3226, USA

Antino Kim
(antino@iu.edu)

Indiana University, Kelley School of Business, Bloomington, IN 47405, USA

Atanu Lahiri
(atanu.lahiri@utdallas.edu)

University of Texas at Dallas, Jindal School of Management, Richardson, TX 75080, USA

Last revised: October 11, 2017

Abstract

Controlling digital piracy has remained a top priority for manufacturers of information goods, as well as for many governments around the world. Among the many forms taken by digital piracy, we focus on an increasingly common one—namely, *online piracy*—that is facilitated by torrent sites and cyberlockers who bring together consumers of pirated content and its suppliers. Motivated by recent empirical literature which makes a clear distinction between anti-piracy efforts that restrict supply of pirated goods (*supply-side* enforcement) and ones that penalize illegal consumption (*demand-side* enforcement), we develop a simple economic model and discover some fundamental differences between these two types in terms of their impacts on innovation and welfare. All in all, supply-side enforcement turns out to be the “longer arm”—it has a more desirable economic impact in the long run. Our results have clear implications for manufacturers, consumers, and policymakers.

Keywords: Online piracy, supply-side enforcement, demand-side enforcement, innovation, welfare.

1 Introduction

Piracy of digital goods—ranging from movies, music, and TV shows to video games and computer software—has become an important issue facing their manufacturers and, at times, consumers and

*We thank the editors and reviewers of *Management Science* for their constructive feedback; the paper has greatly benefited from their efforts. Earlier versions of this research have been presented at the following workshops and conferences: Conference on Information Systems and Technology (CIST) 2013, Institute for Operations Research and the Management Sciences (INFORMS) Conference 2013, Theory in Economics of Information Systems (TEIS) Workshop 2014, Production and Operations Management Society (POMS) Conference 2014, and Workshop on Information Systems and Economics (WISE) 2015. We express our gratitude to the organizers of these workshops and conferences for the opportunity to present our work there and to the participants for their valuable feedback. We also thank the seminar participants at the following institutions: Indian School of Business (Hyderabad), Indian Institute of Management (Bangalore and Calcutta), Indiana University, University of Connecticut, and Seattle University.

policymakers. Digital piracy can take many forms. For example, when two friends share copyrighted music or software, such sharing is indeed a form of piracy (Duhigg 2006). Similarly, when a small neighborhood computer shop throws in a copy of Microsoft Windows or Office for free to attract potential consumers, chances are that the free software product is actually pirated (Diola 2013, Kan 2012). And, of course, when a movie enthusiast downloads a popular movie from a torrent site or a cyberlocker, it is also an act of piracy (Liebelson 2014). Among the many forms of digital piracy that exist today, it is this last type—often referred to as *online piracy*—that is the focus of our research.

To counter online piracy, significant steps—we call them anti-piracy measures or enforcement efforts—have been taken globally. Traditionally, they have been centered mostly around the consumers of pirated products, and have typically included monitoring, auditing, and levying penalties against illegal use (Danaher and Smith 2014, de Beer and Clemmer 2009, Farivar 2013, Mills 2012). In recent times, however, governments have also started scanning for sites that distribute, or aid in the distribution of, pirated content. Often, governments have attempted to arrest the actions of these sites, by shutting them down and prosecuting the site operators (Danaher et al. 2014, Epstein 2012, Horwitz 2013, Masnick 2012).

In capturing the essence of enforcement efforts, prior theoretical works on piracy have typically associated anti-piracy measures with a *piracy cost*—the expected loss resulting from potential legal liabilities, that is, the probability that piracy gets detected times the expected penalty assessed on detection (cf. August and Tunca 2008, Lahiri and Dey 2013). Such a conceptualization is certainly applicable to instances where penalties are imposed on illegal consumption. A higher level of enforcement increases either the probability of detection or the expected penalty on detection, or both, effectively making piracy a more costly option for consumers.

However, the above conceptualization does not capture all anti-piracy measures seen in practice today. In particular, it does not capture the ones that only restrict the supply of pirated content and not directly involve prosecution of pirate consumers. A case in point is the situation described in Figure 1. This figure illustrates how, based on a request from Viacom, the Digital Millennium Copyright Act (DMCA) has forced YouTube to stop delivering the content sought by an end user. Evidently, in addition to the class of anti-piracy measures that curbs the *demand* for piracy by making pirated goods less attractive to consumers, there is also a class that impacts the *supply*

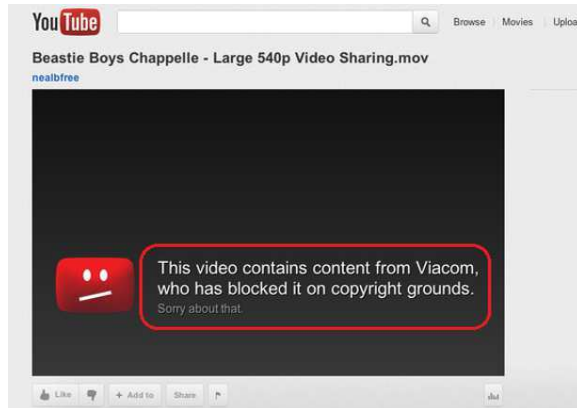


Figure 1: DMCA-Induced Filtering of Pirated Content by Provider

side and diminishes the visibility of pirated content. Following Danaher et al. (2014), we term the former type “demand-side” enforcement, and the latter, “supply-side” enforcement. Table 1, which categorizes different real-world instances of anti-piracy measures, clearly illustrates that this demand- versus supply-side classification is not only simple and intuitive but is also exhaustive. Recently, Danaher and Smith (2014) and Danaher et al. (2014) have empirically investigated the

Table 1: Examples of Demand- and Supply-Side Enforcement

Demand Side	Supply Side
New Zealand’s “three strikes” law against illegal downloaders	MegaUpload.com shutdown and prosecution of Kim Dotcom for online piracy
Copyright Alert System or CAS (also called the “six strikes” program) in the US	China shutting down two major sites involved in online piracy
Canipre, a Canadian intellectual property rights company, tracking down illegal downloaders	French government penalizing “sites that profit from pirated material”
Mass scanning of IP addresses in Australia and the US	New anti-piracy law in Russia to tackle pirate sites
Enactment of the French HADOPI law (Creation and Internet law)	Founder of illegal movie streaming site sentenced to 4.5 years in jail, \$4.7 million in fine
Boston University graduate student fined \$675,000 for illegally downloading 30 songs	Google voluntarily playing copyright cop, suppressing violators in search results
Minnesota woman fined \$220,000 for illegally downloading 24 songs	Google dropping <i>Pirate Bay</i> from auto-complete results
Enactment of Japan’s new law punishing illegal downloaders to a jail-term of up to 2 years	Stop Online Piracy Act (SOPA) and Protect IP Act (PIPA) of 2012
Several Canadian firms and one school division fined \$270,000 in piracy-related damages	Combating Online Infringements and Counterfeits Act (COICA) of 2010
Random audits in S. Korea for software piracy at companies, universities, and government agencies	“Operation In Our Sites” initiative in the US resulting in the seizure of 125 websites

impacts of demand- and supply-side enforcement on legal sales and found both types to have

favorable impacts. Thus, although there are useful empirical findings about private profit, the impact on innovation as well as public welfare—consumer and social—has not received adequate attention and remains an important open issue.¹

Rooted in the unsettled nature of this issue is our desire to address a set of questions, which have become increasingly prominent in the backdrop of online piracy fueled by the closure of the likes of `PirateBay.com` and `MegaUpload.com`. For example, should anti-piracy measures be directed at illegal downloaders, or should they target sites that distribute pirated content? How do these measures impact consumer and social welfare? And, if indeed a shift in enforcement efforts were to take place, what would be the short- and long-run implications? In this work, we bring both enforcement types into one consolidated framework and investigate whether there are reasons to suspect the two approaches to have different economic implications, not only from the manufacturer’s perspective, but also from consumers’ and policymaker’s points of view. In doing so, we theoretically scrutinize previous empirical findings on legal sales, and also examine, if the two approaches indeed have different economic consequences, which direction has a “longer arm”—a better reach, a more desirable impact—within a socioeconomic system.

Intuitively, either type of enforcement can curb piracy and may thus be beneficial to a manufacturer. However, it is not obvious exactly how the manufacturer will, or even should, react to these two types of enforcement, and whether that will eventually translate to gains or losses in welfare. Recent research shows that a manufacturer may, in fact, respond to a higher level of demand-side enforcement by decreasing the quality of its product and lowering welfare in the process (Lahiri and Dey 2013), contradicting the common argument that stronger enforcement is necessary to foster innovation (Jain 2008). We seek to verify whether the impact of supply-side enforcement is also qualitatively similar and, if not, why, when, and how it differs.

In developing our model, we inherit much of the setup from prior literature. Specifically, we consider a profit-maximizing monopolist serving consumers heterogeneous in their taste for quality (e.g., Moorthy 1984, Mussa and Rosen 1978), and assume that there exists a quality difference between the pirated and legal versions (e.g., Lahiri and Dey 2013, Sundararajan 2004). Furthermore, we assume that higher levels of demand-side enforcement make piracy more costly—and hence less attractive—to consumers (e.g., August and Tunca 2008). Our model, in essence, is an

¹To be sure, Rob and Waldfogel (2006) do estimate the welfare implications of piracy from survey data, but they do it only for the short-run case where quality is fixed. Also, they do not consider supply-side enforcement.

extension of the typical setup used in prior literature, the new elements being the ad-supported pirate suppliers and the presence of supply-side enforcement.

Our research questions and modeling approach put the spotlight squarely on how suppliers of pirated content actually operate, unraveling the mystery surrounding their business models. The underlying ecosystem shows a curious interdependence among aggregator sites, online ad agencies, and pirate suppliers (Aguilar et al. 2015, Seidler 2011). In practice, many aggregator sites providing online access to pirated content operate on revenues generated through online advertisements. These sites include, among others, content-hosting and streaming sites known as *cyberlockers* (e.g., MegaUpload.com, FileServe.com, and RapidShare.com), *torrent indexers* such as BitSnoop.com, as well as individual blogs with clickable links to pirated content. Since the ad revenue increases with the traffic, such a site incentivizes individual pirate suppliers to share content that will be in high demand, as measured by the number of downloads. Pirate suppliers leveraging such a site accordingly offer free pirated content to lure hoards of downloaders to their pages. The traffic attracted by a supplier, in turn, has a direct impact on its “revenue” funneled through the aggregator site. It is the explicit modeling of this ecosystem that sets our work apart from prior literature.

We find that supply-side enforcement, when compared to its demand-side counterpart, has a more favorable impact on innovation and welfare in the long run. This superiority of supply-side enforcement may seem puzzling a bit—after all, either type of enforcement essentially makes piracy more difficult. How could their impacts then be so different? A closer examination reveals a subtle interplay between the enforcement type and the piracy ecosystem. Consistent with prior literature (Lahiri and Dey 2013), a higher level of demand-side enforcement makes the pirated version less appealing to consumers, and the diminishing shadow competition from the pirated product, in turn, allows the manufacturer to respond with a lower quality. Stronger supply-side enforcement, on the other hand, has a more direct impact on pirate suppliers and makes pirated content less available, but its indirect impact on the relative appeal of pirated content ought to be much weaker, which limits the manufacturer’s ability to reduce quality. In fact, the reduction in supply effectively expands the legal consumer base, yielding better returns on investments in product development and inducing the manufacturer to invest more in quality. It is this positive impact on quality that eventually reflects itself more favorably on welfare.

2 Literature Review

The main point of our work is comparing the two forms of enforcement in terms of their impacts on a monopolist’s incentive to innovate and the resulting impacts on consumer and social welfare. In this regard, existing research provides some interesting cues. For example, Danaher et al. (2014) scrutinize the effects of *HADOPI*—a demand-side anti-piracy law used in France to punish repeat offenders more severely (see Table 1)—and find that it resulted in a 25% increase in legal sales. In another work, Danaher and Smith (2014) study the effects of shutting down `MegaUpload.com`—clearly a case of supply-side enforcement—and estimate that this shutdown resulted in a 6.5–8.5% boost in sales of digital movies. In two other related studies, Danaher et al. (2015, 2016) find that blocking of piracy sites, also a supply-side measure, can be effective in boosting legal sales only when several such sites are blocked simultaneously, but the impact on legal sales is insignificant when only one site is blocked; they also find a diminishing return from blocking more piracy sites. Although empirical studies can spot such immediate impacts on sales, the welfare implications of enforcement are far more elusive. Our focus in this work, therefore, is complementing this line of research by qualitatively distinguishing various anti-piracy efforts and comparing their welfare implications.

Some researchers have conducted field experiments to study the influence of anti-piracy measures directed at the demand side. For example, Baird et al. (2016) conduct experiments to investigate how versioning, in the face of increasing piracy cost, affects consumers’ inclination towards piracy, their surplus, and the manufacturer’s profit. Hashim et al. (2014) study how anti-piracy campaigns aimed at consumers through direct and indirect channels can have differing impacts.

There is also a vast literature on the economics of piracy using quantitative models. A branch of this literature identifies several situations in which a manufacturer may find it profitable to tolerate or support piracy. One such situation occurs in the presence of a positive network effect that translates illegal usage into a higher willingness to pay for the legal product (Conner and Rumelt 1991). Just as positive network effects alter the incentive to tolerate piracy, negative network effects do, too—August and Tunca (2008) find that denying pirates critical security patches can be counterproductive when doing so makes legal users vulnerable to security attacks, lowering their willingness to pay. Another branch examines when and how certain tools—digital rights management (DRM), nonlinear pricing, versioning, patch distribution, bundling, content delivery

technology, or free trials—may help a manufacturer combat piracy (e.g., Chellappa and Shivendu 2005, Cho and Ahn 2010, Gopal and Gupta 2010, Johar et al. 2012, Kannan et al. 2016, Sundararajan 2004, Vernik et al. 2011, Wu and Chen 2008). In order to identify how various anti-piracy measures impact the intensity of in- and cross-channel competitions, Geng and Lee (2013) consider a setting where legal retailers compete with each other as well as with pirate services and find that the impact of piracy could be different based on the level of in-channel competition.

The closest to our work is the branch that examines the economic impacts of anti-piracy efforts on a manufacturer’s strategy and the resulting welfare (Bae and Choi 2006, Chen and Png 2003, Lahiri and Dey 2013). However, this branch considers only demand-side enforcement. For instance, Chen and Png (2003) examine and compare three ways to curb piracy, all directed at changing the relative appeal of the pirated product to a consumer, vis-à-vis the legal one. Bae and Choi (2006) and Lahiri and Dey (2013) also assume the supply of pirated goods to be exogenous and unlimited. Our main contribution to this stream of literature is that: (i) we *model the ecosystem* of online piracy including pirate suppliers as strategic players, and (ii) we *endogenize the supply* of pirated content. In doing so, we are able to obtain a more complete picture and answer our research questions by comparing the economic impacts, especially those on innovation and public welfare, of supply- and demand-side anti-piracy measures.

Our findings relate well to the current literature and extend them logically. Results in prior research are often mixed. Some show that a lack of enforcement and higher piracy can decrease the manufacturer’s revenue, killing incentives to innovate and leading to lower quality products (Bae and Choi 2006, Jain 2008). Lahiri and Dey (2013), however, tell a different story. They argue that, in certain circumstances, less (demand-side) enforcement may surprisingly lead to higher quality products, and eventually to higher consumer and social welfare. At the core of these results lies an argument that the manufacturer can leverage a higher quality to “compete” against piracy; so, when (demand-side) enforcement is strong, the manufacturer simply responds with a lower quality, which in turn adversely impacts the legal consumers. We find that this argument continues to apply to our wider setting but only so far as demand-side enforcement is concerned—the impact of supply-side enforcement on innovation is often exactly the opposite, and its long-run impact on consumer and social welfare, strikingly different.

Finally, it is worth mentioning that, although not explicitly recognized as such, the supply side

of piracy has started garnering some attention. In particular, the role of commercial pirates, who price illegal versions to maximize profit, has been examined. For example, Jaisingh (2009) shows that the existence of commercial pirates can confound a manufacturer’s response to piracy in unpredictable ways. Tunca and Wu (2013) find that increasing enforcement against individual pirates in P2P networks might make commercial pirates more competitive, harming the manufacturer in the process. Neither work, however, explicitly models the ecosystem (comprising aggregator sites, on-line ad agencies, and pirate suppliers) that sustains online piracy, nor do they analyze the economic implications of disrupting this ecosystem.

3 Model and Analyses

We develop an economic model with three strategic players: (i) a profit-maximizing monopolist,² (ii) pirate suppliers supported by advertisements, and (iii) utility-maximizing consumers. The supplier and consumer bases are both normalized to a mass of one. Our monopolist, situated in a market with certain levels of demand- and supply-side enforcement, chooses the price and quality of its product.

The timeline is as follows: First, the manufacturer offers a product of quality $\theta > 0$ at a price $p > 0$. Each supplier of pirated content then decides whether to provide an illegal version or not—only when potential revenues from piracy can fully offset the cost imposed by supply-side enforcement, a supplier provides a pirated copy. This determines the supply level and availability of pirated content. Finally, each consumer decides whether to buy, pirate, or forgo use.

3.1 Enforcement Environment

Before describing the players and their behavior, it is important to discuss the enforcement environment in which they operate. We assume that this environment is characterized by two parameters, r and e , respectively representing the levels of demand- and supply-side enforcement activities. When demand-side enforcement is stepped up, that is, when r increases, it results in either a higher probability of getting detected when using a pirated copy or a higher penalty on detection, or both. In other words, similar to prior literature (August and Tunca 2008, Lahiri and Dey 2013), a higher r

²The assumption of a monopolist is not critical to our results. For details, please see Appendix ?? in the online supplement to the paper.

simply increases the expected legal penalty a consumer faces and is a proxy for his piracy cost. In contrast, supply-side enforcement, e , is a proxy for the entry cost faced by a pirate supplier, which includes among other things the risk of prosecution and penalty if convicted of distributing illegal copies. In addition, it captures the combined effect of all other actions that make it difficult to supply pirated content. Such actions may include blocking pirate sites, shutting down cyberlockers, and requiring search engines to filter out links to illegal content.

We argue that both r and e largely depend on the political and legal environment in which a business operates. For example, the cost of piracy in certain developing nations, both for consumption and supply, is quite low because, in those countries, either governments are remiss in enforcing intellectual property laws and international treaties, or the penalty on detection is low under their judicial systems. In contrast, there are hundreds of piracy-related criminal prosecutions in the United States every year.

Given such an enforcement environment, we first characterize the market equilibrium and then, through comparative statics, examine how this equilibrium behaves when one or the other type of enforcement is increased.

3.2 Consumer’s Preference for Quality

Our consumers are heterogeneous along two orthogonal dimensions and are indexed by $\langle v, h \rangle$ — v represents consumers’ preference for quality, that is, consumer $\langle v, \cdot \rangle$ gets a value of $v\theta$ from using a product of quality θ . On the other hand, as will be discussed in Section 3.4, h stands for their technical ability or savvy. We make the following assumption about v :

Assumption 1 *A consumer’s preference for quality, v , is uniformly distributed over $[0, 1]$.*

We now explain what quality means for a digital good and what impact piracy may have on this quality. In our context, the term *quality* is being used in the classical sense—it essentially captures those characteristics of a product that are desired by all consumers (Moorthy 1984). Quality of a digital good—also a proxy for *innovation* in our abstraction—may comprise all or some such appropriate and desirable characteristics as the ease of use, speed, functionality, flexibility, portability, resolution, fidelity and encoding bit rate, among others.

Prior research has often considered the pirated product to be of lesser quality (e.g., Duchene and Waelbroeck 2005, Sundararajan 2004). Indeed, examples abound where the physical quality of the

pirated copy is less than that of the original, as is usually the case with pirated movies (Karaganis 2011). Similarly, pirated software products do not often receive certain updates and patches (August and Tunca 2008), and may be missing important functionalities or contain embedded malicious codes (Lahiri and Dey 2013), implying that the quality of the pirated content, ϕ , is likely less than θ . It is possible to argue to the contrary as well, that is, to claim $\phi > \theta$. For example, a pirated media file may not have pesky restrictions in terms of the number of devices where it can be played (e.g., Vernik et al. 2011). Likewise, in the case of software products, such as Mathematica or SAS, a pirated copy may be viewed by some as more convenient as it may not require a periodic renewal of the license. Irrespective of whether piracy enhances or impairs quality, the pirated product is highly unlikely to have a quality level that is totally independent of the quality of the legal product. Logically, we would expect ϕ to be an increasing function of θ , satisfying $\phi|_{\theta=0} = 0$. For simplicity, we choose a linear form for ϕ , but do so without any precept about whether the pirated content is superior or inferior:

Assumption 2 *For a legal product of quality θ , the quality of its pirated version is $\phi = \beta\theta$, $\beta > 0$.*

Our own observations suggest that, in most real world situations, ϕ is likely to be less than θ , that is, $\beta < 1$. The case of $\beta \geq 1$ is included in our analysis for the sake of completeness.

3.3 Supply-Side Enforcement and Limited Supply of Pirated Content

We now consider how supply-side enforcement impacts the availability of pirated content. Because there are numerous pirate sites (e.g., cyberlockers and torrents) that can be used to easily distribute illegal content, we assume that there is a large number of identical potential pirate suppliers.³ The presence of many piracy sources is highlighted by Danaher et al. (2016) who found that, even after a total of fifty three sites were blocked in the UK, piracy was not completely eradicated there, as consumers were able to locate yet other pirate sites that remained unblocked.

³There are usually hundreds of pirate sites, serving downloaders with different types of products, such as movies, music, software, and games. Consider, for example, a partial list of top pirate sites for movies alone: 1173x.com, animesuki.com, anonamouse.com, asiatorrents.com, bitmetv.com, bitsnoop.com, blackcats.com, broadcasthe.com, bt-chat.com, btdigg.com, crystaltorrents.com, dodomovies.com, evomovies.com, extratorrents.com, eztv.com, fenopy.com, h33t.com, haloamovies.com, iptorrents.com, isohunt.com, kickasstorrents.com, mamega.com, mininova.com, mogomovies.com, movie24.cc, movieetv.com, moviewatcher.to, mydownloadtube.com, newmovies.ws, newmoviesonline.ws, nyaa.se, oldpiratebay.com, ozomovies.com, popcorn.cab, rarbg.com, revolutiontt.com, sceneaccess.com, sumotorrent.com, torrentdownloads.com, torrenthound.com, torrentleach.com, torrentproject.com, torrentR.com, torrentreactor.com, torrentus.com, torrenz.com, thepiratebay.com, transtraffic.com, what.cd, yifytorrents.com, yourbittorrent.com.

This assumption of a large number of suppliers is essentially equivalent to saying that suppliers are *atomistic*, that is, the decision of a single supplier has little impact on the overall supply level. Our atomistic suppliers are like atomistic traders in financial markets, where the decision of one trader cannot impact the market, but many traders collectively can. Indeed, there is clear empirical evidence that an individual supplier has only limited impact on the piracy market—Danaher et al. (2015) observe that blocking a single site did not have significant impact on the sales of legal copies, but blocking several at the same time did.

Despite there being a large number of pirate suppliers, not all of them will end up entering the piracy “market.” For, similar to Tunca and Wu (2013), we also take into account that each pirate supplier faces an entry cost; as explained earlier, this entry cost is denoted by e , the level of supply-side enforcement, which determines how many pirate suppliers will enter the market—when e is high, no supplier will find it profitable to enter the market, and when e is low, there will be abundant supply as almost every pirate supplier will become active. We denote the supply level of pirated content by $\eta(e)$, which is the fraction of pirate suppliers—or their number, since the mass of potential pirate suppliers has been normalized to one—active in the market in distributing illegal copies. If $\eta(e)$ were zero, there would be no supply; on the other hand, if $\eta(e)$ were one, pirated copies would be readily available.⁴ We do consider these two extreme cases in our overall analysis, but focus primarily on the case where $\eta(e) \in (0, 1)$.

3.4 Consumer Behavior in the Face of Limited Supply

Faced with a limited supply of η , consumers decide on whether to buy, pirate, or forgo use. Of course, a consumer can pirate only if he can access a pirated copy. The ease with which a consumer can eventually access an illegal copy depends not only on η but also on the consumer’s technical savvy or ability, which covers a wide spectrum of know-hows and skills—knowledge of appropriate keywords and search engines, familiarity with different types of web sites (such as aggregator sites, torrent-indexing sites, blogs, and cyberlockers), and working knowledge of torrents, protocols, and file converters (Khantwal 2016, Veneziani 2007). Literature posits that a consumer, before he can start pirating digital content, must first incur some *fixed* learning cost to develop these skills and to learn how to acquire illegal content from certain pirate sites (Clement et al. 2012, Danaher et al.

⁴Actually, when e approaches zero, $\eta(e) \rightarrow 1$ and our model reduces to that in (Lahiri and Dey 2013), which assumes that there is no supply-side enforcement, the supply level is exogenous, and pirated content, abundant.

2010, 2016, Morris and Higgins 2010).⁵

The notion of learning (or education) as an investment into the future is certainly not new, and it makes intuitive sense in our context as well. When a rational consumer decides to familiarize himself with a new pirate site and to incur an additional learning cost in turn, he knows fully well that such an investment may bear fruit not merely for the pirated product he is trying to access right now, but well beyond. The newly acquired knowledge will likely continue helping him with his future access to other pirated products. For example, when a consumer learns the intricate details of a torrent site, he can use this skill repeatedly over a long period of time to download a large variety of movies to watch, songs to listen to, or video games to play. Put differently, a consumer’s decision to “invest” in learning a new site is relevant over a much longer horizon, for many different types of products, and as such must be viewed separately from his decisions that are specific only to a single product or a single purchase. This conceptualization of learning cost as an investment towards future piracy also has solid empirical backing. Danaher et al. (2010), for example, found support for the existence of “a significant fixed cost to piracy;” they explain this fixed cost as the learning cost (or moral cost) borne by consumers.⁶

Before proceeding, we must recognize that the notion of learning cost is indeed different from what has traditionally been called the search cost in information systems and economics literature (e.g., Bakos 1997, Geng and Lee 2013, Wolinsky 1983). First, search cost is *variable*—it is incurred every time a consumer searches for a product. In contrast, learning imposes a *fixed cost* on the consumer. Once a site is learned, a consumer can continue to use it costlessly, as long as the site remains accessible (Clement et al. 2012). The fixed learning cost should not be a part of consumers’ individual rationality (IR) and incentive compatibility (IC) constraints, but the variable search cost should be. Second, in our setting, search itself is free, especially when compared to how traditional economics views search cost as the cost of physically traveling from one store to another (Wolinsky 1983). Indeed, in a world with a multitude of search engines, search is instantaneous and costless (Baye et al. 2006); a large list of possibly relevant results or links can be obtained from a search engine in a flash. Of course, whether these links are in fact the ones that a consumer

⁵We are grateful to a reviewer and the associate editor of *Management Science* for their patient guidance in our development of this part of the model using the central concept of learning cost.

⁶The seemingly different concepts of learning and moral costs are actually related in practice. Consumers with higher moral costs are likely to be more resistant towards adopting piracy. They would, therefore, be less keen on learning how to pirate, as though it is costlier for them to learn. And, at the same time, a technically challenged consumer may rationalize his lack of skills as his high moral standards.

is comfortable working with is a different proposition altogether. The latter ought to depend on the consumer’s technical savvy and his learning cost. Some consumers are technically savvy and naturally more attracted towards piracy. They learn a new site and its protocols more quickly and thus ought to incur a lower learning cost. In contrast, people lacking in technical ability are far less prone to piracy and have a higher learning cost. Thus, unlike traditional search costs, learning costs are different for different consumers. The more (less) technically savvy a consumer is, the lower (higher) should be his learning cost, and *vice versa*, implying an inverse relationship between his savvy and learning cost (Danaher et al. 2010, 2016). We make the following assumption:

Assumption 3 *A consumer’s technical ability, h , follows an exponential distribution with a mean of ω ; consumer $\langle \cdot, h \rangle$ incurs a fixed cost of $\frac{1}{h}$ for every new pirate site he learns.*

The exponential assumption for h has strong empirical backing—Table 1 in (Danaher et al. 2016), which categorizes pirate users into ten segments based on their affinity towards piracy, clearly exhibits the rapidly decaying nature that is typical of an exponential distribution.

Now, what is the maximum number of sites a consumer would learn given that he would incur a learning cost of $\frac{1}{h}$ every time he tries to learn a new site? We assume that there exists a threshold or a budget, T , for the total learning cost, beyond which the consumer simply gives up his efforts to learn about a new pirate site.⁷ The presence of T as a limit on the total learning cost immediately implies that k , the maximum number of sites a consumer would want to learn, must satisfy:

$$\frac{k}{h} \leq T < \frac{k+1}{h} \quad \Leftrightarrow \quad k \leq hT < k+1.$$

Since h is exponential with mean ω (Assumption 3), the probability mass associated with k is:

$$\Pr[k \leq hT < k+1] = \Pr\left[\frac{k}{T} \leq h < \frac{k+1}{T}\right] = \int_{\frac{k}{T}}^{\frac{k+1}{T}} \frac{e^{-\frac{h}{\omega}}}{\omega} dh = e^{-\frac{k}{\omega T}} (1 - e^{-\frac{1}{\omega T}}).$$

⁷Why should the budget T be the same for all consumers? In reality, there could indeed be heterogeneity in T as well. However, this does not impose any undue restriction in our setup because, as we will soon see, h and T always occur together as a product (hT) in our abstraction, and the overall heterogeneity in consumers’ learning behavior can be captured equally well with either one of them; we choose to capture it in h while keeping T fixed. Viewed this way, our modeling setup is exactly the same as the situation where consumers are heterogeneous along both h and T , and their product (hT) is exponentially distributed. Incidentally, if h and T are both exponentially distributed, the distribution of their product becomes a modified Bessel function of the second kind, which exhibits a similar exponential decay and is well approximated by an exponential distribution.

It is also possible that T is not independent of e , as consumers may increase their learning budget in the face of stricter enforcement. We discuss this issue in Appendix ?? in the online supplement to the paper.

Clearly, consumers with $k = 0$ are the ones who have no inclination whatsoever to violate copy-rights, so a fraction g of consumers can be viewed as inherently *ethical*, where $g = \Pr[k = 0] = \Pr[0 \leq hT < 1] = 1 - e^{-\frac{1}{\omega T}}$. It is easy to see that g is a decreasing function of ωT , and

$$\lim_{\omega T \rightarrow \infty} g = 0 \quad \text{and} \quad \lim_{\omega T \rightarrow 0} g = 1.$$

When ωT is large, consumers are either more savvy or have a lot more patience for learning, and most consumers would end up trying a hand in piracy. In contrast, when ωT is small, few consumers have the ability or the patience to consider piracy as a viable option.

Now, a consumer who has learned all k sites would not have access to the pirated version only if all of the k sources are inactive, which occurs with a probability of $(1 - \eta)^k$, because the supply level η represents the probability that any given source is active. Taking an expectation over k , we can find the fraction of consumers not having access to a pirated copy as:

$$\lambda(\eta) = \sum_{k=0}^{\infty} (1 - \eta)^k e^{-\frac{k}{\omega T}} (1 - e^{-\frac{1}{\omega T}}) = \frac{1 - e^{-\frac{1}{\omega T}}}{1 - e^{-\frac{1}{\omega T}} (1 - \eta)}. \quad (1)$$

3.5 Consumer Demand

A consumer faces an expected legal penalty of r if he ends up using an illegal copy. As explained earlier, this penalty is exogenous in our model and simply depends on the level of enforcement against the consumption of pirated goods. Clearly, a consumer can enjoy a utility of $(v\theta - p)$ from purchasing the legal version, and $(v\beta\theta - r)$ from a pirated copy; which one he chooses in the end depends not just on these utilities, but also on the availability of pirated content, as determined by supply-side enforcement activities.

It is clear that the limited supply partitions the consumer base into two segments: (i) a fraction $\lambda(\eta)$ who does not have access to pirated content, and (ii) a fraction $(1 - \lambda(\eta))$ who does. Consider the first segment. A consumer who cannot locate pirated content would buy the legal product if and only if his individual rationality (IR) constraint is met: $v\theta - p \geq 0$, that is, $v \geq \frac{p}{\theta}$. On the other hand, when a consumer can locate a copy, he would choose to use the legal product if and only if, in addition to the above IR constraint, the following incentive compatibility (IC) constraint is also satisfied: $v\theta - p \geq v\beta\theta - r$. At the same time, a consumer would use the pirated version if $v\beta\theta - r > 0$ (IR) and $v\theta - p < v\beta\theta - r$ (IC). Figure 2 graphically captures these IR and IC

constraints and identifies the legal and illegal demands for various market configurations in which piracy exists.⁸ Using Figure 2, the legal and illegal demands, denoted q and \hat{q} , respectively, can

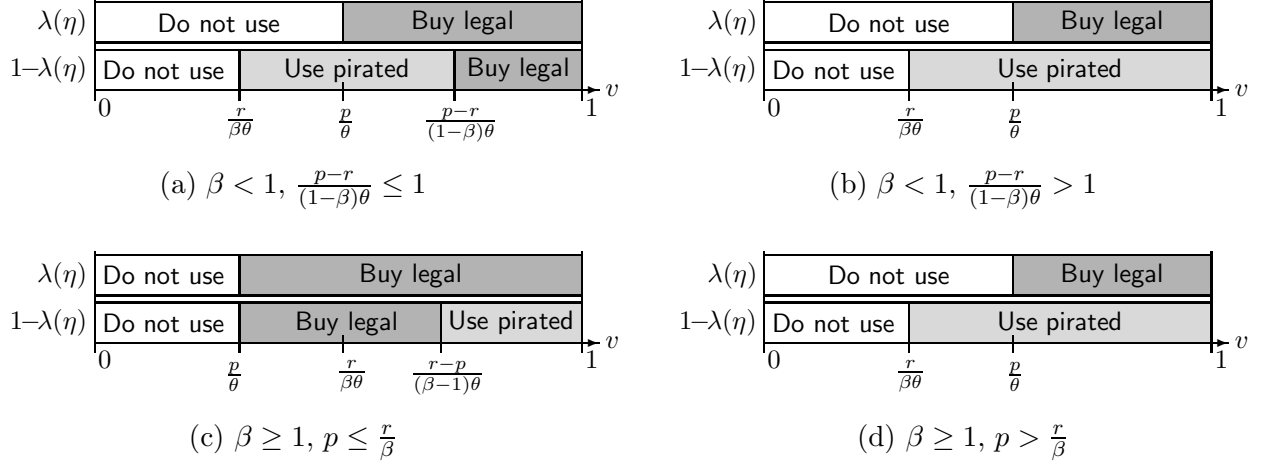


Figure 2: Consumers Self-Select Based on Their Relative Benefits

now be expressed as:

$$q = \begin{cases} \lambda(\eta) \left(1 - \frac{p}{\theta}\right) + (1 - \lambda(\eta)) \left(1 - \min\left\{1, \frac{p-r}{(1-\beta)\theta}\right\}\right), & \text{if } \beta < 1, \\ \lambda(\eta) \left(1 - \frac{p}{\theta}\right) + (1 - \lambda(\eta)) \left(\min\left\{1, \frac{r-p}{(\beta-1)\theta}\right\} - \min\left\{\frac{p}{\theta}, \frac{r-p}{(\beta-1)\theta}\right\}\right), & \text{otherwise;} \end{cases} \quad (2)$$

and

$$\hat{q} = \begin{cases} (1 - \lambda(\eta)) \left(\min\left\{1, \frac{p-r}{(1-\beta)\theta}\right\} - \frac{r}{\beta\theta}\right), & \text{if } \beta < 1, \\ (1 - \lambda(\eta)) \left(1 - \min\left\{1, \max\left\{\frac{r}{\beta\theta}, \frac{r-p}{(\beta-1)\theta}\right\}\right\}\right), & \text{otherwise.} \end{cases} \quad (3)$$

3.6 Supplier Behavior and Subgame Equilibrium

As mentioned in Section 3.3, in equilibrium, the supply level, η , ought to depend on the level of supply-side enforcement, e . Accordingly, we now complete the model specification by endogenizing η using a standard break-even analysis in which players enter the market as long as there is profit to be made (e.g., Vickers 1989). Pirate suppliers are typically paid based on the number of downloads of pirated content uploaded by them (Aguilar et al. 2015). The total advertisement revenue earned

⁸Prior research has found that every pirate demand is not necessarily a lost sale (Rob and Waldfogel 2006). Indeed, in panels (a), (b), and (d) of Figure 2, consumers with $v \in \left(\frac{r}{\beta\theta}, \frac{p}{\theta}\right)$ would not buy the legal product, even when they cannot find pirated content; they constitute what Rob and Waldfogel call “incomplete displacement.”

by all suppliers should then be proportional to \hat{q} , the total number of illegal downloads; henceforth, without loss of generality, we assume this constant of proportionality to be one. In a piracy market with identical suppliers, the revenue for each supplier is then $\frac{\hat{q}}{\eta}$, the total revenue divided by the number of suppliers participating in distributing an illegal copy. A supplier compares this revenue with his entry cost, e , and participates as long as the revenue is more than the cost (Vickers 1989). Thus, in a *subgame equilibrium*, a supplier's utility is $\frac{\hat{q}}{\eta} - e = 0$, implying

$$\hat{q} = \eta e. \quad (4)$$

Substituting (1) into (3) and equating the resulting expression to (4), we can solve for \hat{q} , η , and λ in the subgame equilibrium. We obtain:

$$\lambda = \begin{cases} \frac{e\theta\xi(1-\beta)}{p-\frac{r}{\beta}}, & \text{if } \frac{p-r}{(1-\beta)\theta} \leq 1 \text{ and } \beta < 1, \\ \frac{e\theta\xi}{\theta-\frac{r-p}{\beta-1}}, & \text{if } p \leq \frac{r}{\beta} \text{ and } \beta \geq 1, \\ \frac{e\theta\xi}{\theta-\frac{r}{\beta}}, & \text{otherwise,} \end{cases} \quad (5)$$

where $\xi = \frac{g}{1-g} = \mathbf{e}^{\frac{1}{\omega T}} - 1$.

3.7 Manufacturer's Decision and Market Equilibrium

Equation (5) allows us to estimate the legal demand in (2), in terms of the enforcement levels, r and e . The overall equilibrium is then found by simply maximizing the manufacturer's profit. We assume:

Assumption 4 *The manufacturer's marginal cost of producing an additional copy is zero, and its cost of developing a product of quality level θ is $\frac{c\theta^2}{2}$, $c > 0$.*

Thus, the manufacturer's problem is to solve: $\max_{p,\theta} \pi = pq - \frac{c\theta^2}{2}$. Although conceptually straightforward, solving this problem and analytically characterizing its solution are not simple. This is because the manufacturer's strategy may shift as we move from one point in the parameter space to another, resulting in singularities with respect to the decision variables at the boundaries of these strategies. In equilibrium, three regions/cases emerge: (i) Region/Case 1 where supply is limited ($0 < \eta < 1$), (ii) Region/Case 2 with ample supply ($\eta = 1$), and (iii) Region/Case 3 with no piracy

($\eta = 0$). A more complete description of these regions (along with all appropriate subregions and their boundaries) is provided in Appendix ?? in the online supplement to this paper. As mentioned in Section 3.3, we focus primarily on the case where $\eta \in (0, 1)$:

Definition 1 (Primary Piracy Region) *Region 1, where a limited amount of supply exists—that is, $0 < \eta < 1$, $\eta \neq 0, 1$ —is henceforth called the primary piracy region.*

Two subcases are possible in the primary piracy region: Case 1A, where the manufacturer names a price, p , in such a manner that it ends up selling to both types of consumers, with or without access to pirated content, and Case 1B, where the manufacturer names the price so high that it effectively shuts out from the legal product all consumers who have access to pirated content.

3.7.1 Short-Run Equilibrium: Exogenous Quality

We first consider the manufacturer’s decision in the short-run situation, where the quality of the information good is fixed and the manufacturer can only set the price. Here, the manufacturer’s decision problem becomes: $\max_p R = p q(\theta, p)$, for a fixed, exogenous θ .

Lemma 1 *Case 1A is not possible in equilibrium when $\beta = 1$.*

The proofs of all the results in the paper, as well as of those in the extensions in Appendix ??, are available in Appendix ??.

Proposition 1 (Short-Run Price) *For a given θ , the equilibrium price is given by:*

$$p^*(\theta) = \begin{cases} \frac{r}{2} + \frac{\theta(1-\beta)(1+e\beta\xi)}{2}, & \text{Case 1A } (\beta < 1), \\ \frac{r}{2\beta} + \frac{e\theta\xi(\beta-1)}{2\beta}, & \text{Case 1A } (\beta > 1), \\ \frac{\theta}{2}, & \text{Case 1B.} \end{cases} \quad (6)$$

From Proposition 1, we can see that $p^*(\theta)$ increases with both r and e in Case 1A—as either type of enforcement increases, piracy declines, leading to a greater pricing power for the manufacturer. In Case 1B, however, the manufacturer disregards the pirates completely while setting the price, which is why $p^*(\theta)$ no longer depends on the enforcement levels.

For certain digital goods, their manufacturers may not be in a position to respond to piracy by altering quality. In those situations, an analysis of the short-run equilibrium is not only of particular

importance, but is also sufficient. However, in many other real-world situations, a manufacturer may be able to adjust the quality level in response to piracy, by altering its investments in product development and associated R&D activities (Jain 2008, Lahiri and Dey 2013). In those cases, it is important that we consider the long-run equilibrium in which the manufacturer has control over θ as well.

To be sure, we are using the terms “short-run” and “long-run” in their traditional microeconomic sense, in which the distinction between the two terms is that (i) in the short run, firms have already incurred the (sunk) fixed costs, while there are no (sunk) fixed costs in the long run, and (ii) irrespective of the sunk cost, a firm produces in the short run as long as the market price covers the variable costs (cf. Davidson and Deneckere 1986, Mankiw 2008, Mas-Colell et al. 1995). An information good involves an up-front development cost, which depends on the quality level chosen by the firm as a part of its long-run production decision. This up-front cost can be rather large; consider, for example, the development of a star-studded movie, an enterprise-level software, or a rich and sophisticated video game. Once the product is developed and introduced to the market, however, this cost is sunk and the short-run decision of the firm is only about pricing the product.

3.7.2 Long-Run Equilibrium: Endogenous Quality

We can write the manufacturer’s long-run decision problem as: $\max_{\theta} \pi(\theta) = p^*(\theta) q(\theta, p^*(\theta)) - \frac{c\theta^2}{2}$.

Lemma 2 *In the long-run equilibrium, Case 1A is not possible when $\beta \geq 1$.*

Lemma 2 simply suggests that competing with the pirated product is futile when β is large. Rather, the manufacturer is better off just concentrating on the λ portion of the market with no access to pirated content.

Proposition 2 (Long-Run Quality and Price) *Let $\theta_{1A} > 0$ be the largest real root of:*

$$\frac{(1 - \beta)(1 + e\beta\xi)^2}{4} - \frac{r^2}{4\theta^2(1 - \beta)} - c\theta = 0.$$

The equilibrium quality is then given by:

$$\theta^* = \begin{cases} \theta_{1A}, & \text{Case 1A } (\beta < 1), \\ \theta_{1B} = \frac{r}{\beta} + \frac{e\beta\xi + \sqrt{e\beta\xi(e\beta\xi - 16cr)}}{8c\beta}, & \text{Case 1B.} \end{cases} \quad (7)$$

The equilibrium price is given by $p^*(\theta^*)$, where $p^*(\cdot)$ is as defined in (6).

4 Results

In this section, to study the effects of changing the enforcement levels, we perform comparative statics on a set of relevant metrics, with respect to r and e .



4.1 Metrics

Although our primary metric for comparing anti-piracy choices is *social welfare*—and arguably should be the only one—we consider several secondary metrics as well, mainly because of their practical relevance. *Piracy rate* is often touted as a measure of significance—a mention of this rate invariably shows up in claims made by manufacturers and their alliances to shore up support for higher levels of enforcement on both sides and, at the same time, piracy rate is much used as a valuable yardstick in policy debates and in the legal parlance (Karaganis 2011). *Quality* of a digital good is frequently used as an indicator of innovation, and a higher quality is usually viewed as socially desirable (Brynjolfsson and Zhang 2007). The manufacturer’s *profit* is also an important yardstick as it usually serves as an indicator of the health of the industry. Finally, *consumer surplus* is a useful metric to understand the basic economic rationale behind consumer activism and lobbying efforts. We now define all these metrics for the short-run case, that is, as a function of θ . The corresponding long-run versions are easily obtained by substituting θ with its case-appropriate value from (7).

Piracy Rate: The *piracy rate*, μ , is defined as the number of pirated copies in use as a fraction of the total user base: $\mu = \frac{\hat{q}}{q+\hat{q}}$. Substituting the equilibrium price from (6) into (2) and (3), we get:

$$\mu(\theta) = \begin{cases} \frac{r(2-\beta)-\beta\theta(1-\beta)(1-e\xi(2-\beta))}{2(1-\beta)(r-\beta\theta(1-e\xi(1-\beta)))}, & \text{Case 1A } (\beta < 1), \\ \frac{r(2\beta-1)-\theta(\beta-1)(e\xi+2\beta(1-e\xi))}{(\beta-1)(r+\theta(e\xi(\beta-1)-2\beta))}, & \text{Case 1A } (\beta > 1), \\ \frac{2(\beta\theta-r)(r-\beta\theta(1-e\xi))}{2r\beta\theta(2-e\xi)-\beta^2\theta^2(2-e\xi)-2r^2}, & \text{Case 1B.} \end{cases} \quad (8)$$

Manufacturer's Profit: The manufacturer's revenue—or profit since the marginal cost is zero—in equilibrium is $\pi^*(\theta) = p^*(\theta)q(\theta, p^*(\theta)) - \frac{c\theta^2}{2}$. As before, substituting (6) into (2), we get:

$$\pi^*(\theta) = -\frac{c\theta^2}{2} + \begin{cases} \frac{(r+\theta(1-\beta)(1+e\beta\xi))^2}{4\theta(1-\beta)}, & \text{Case 1A } (\beta < 1), \\ \frac{(r+e\theta\xi(\beta-1))^2}{4\beta\theta(\beta-1)}, & \text{Case 1A } (\beta > 1), \\ \frac{e\xi\beta\theta^2}{4(\beta\theta-r)}, & \text{Case 1B.} \end{cases} \quad (9)$$

Consumer Surplus: We consider the total surplus of *all* consumers, legal and illegal. Using Figure 2, this surplus can be estimated from:

$$CS(\theta) = \begin{cases} \lambda \int_{\frac{p}{\theta}}^1 (v\theta - p)dv + (1-\lambda) \int_{\frac{p}{\theta}}^1 (v\theta - p)dv + (1-\lambda) \int_{\frac{r}{\beta\theta}}^{\min\{1, \frac{p-r}{(1-\beta)\theta}\}} (v\beta\theta - r)dv, & \text{if } \beta < 1, \\ \lambda \int_{\frac{p}{\theta}}^1 (v\theta - p)dv + (1-\lambda) \int_{\min\{1, \frac{r-p}{(\beta-1)\theta}\}}^1 (v\theta - p)dv + (1-\lambda) \int_{\min\{1, \max\{\frac{r}{\beta\theta}, \frac{r-p}{(\beta-1)\theta}\}}^1 (v\beta\theta - r)dv, & \text{otherwise.} \end{cases}$$

Substituting the equilibrium price from (6), we get:

$$CS(\theta) = \begin{cases} \frac{r^2(4-3\beta)}{8\beta\theta(1-\beta)} - r \left(\frac{3}{4} - \frac{e\xi(1-\beta)}{2} \right) - \frac{e^2\beta^2\theta\xi^2(1-\beta)}{8} - \frac{e\beta\theta\xi(1-\beta)}{2} + \frac{\theta(1+3\beta)}{8}, & \text{Case 1A } (\beta < 1), \\ \frac{r^2(4\beta-3)}{8\beta\theta(\beta-1)} + r \left(\frac{e\xi(\beta-1)}{2\beta} - 1 \right) - \frac{e^2\theta\xi^2(\beta-1)}{8\beta} - \frac{e\theta\xi(\beta-1)}{2} + \frac{\beta\theta}{2}, & \text{Case 1A } (\beta > 1), \\ \frac{e\beta\theta^2\xi}{8(\beta\theta-r)} + \frac{(\beta\theta-r)(\beta\theta(1-e\xi)-r)}{2\beta\theta}, & \text{Case 1B.} \end{cases} \quad (10)$$

Social Welfare: For a good with zero marginal cost, the total social welfare is estimated from the consumption benefits of *all* consumers, legal and illegal, taken together, minus the cost of

development:⁹

$$SW(\theta) = -\frac{c\theta^2}{2} + \begin{cases} \lambda \int_{\frac{p}{\theta}}^1 v\theta dv + (1-\lambda) \int_{\frac{p-r}{(1-\beta)\theta}}^1 v\theta dv + (1-\lambda) \int_{\frac{r}{\beta\theta}}^{\min\{1, \frac{p-r}{(1-\beta)\theta}\}} v\beta\theta dv, & \text{if } \beta < 1, \\ \lambda \int_{\frac{p}{\theta}}^1 v\theta dv + (1-\lambda) \int_{\min\{1, \frac{r-p}{(\beta-1)\theta}\}}^1 v\theta dv + (1-\lambda) \int_{\min\{1, \max\{\frac{r}{\beta\theta}, \frac{r-p}{(\beta-1)\theta}\}}^1 v\beta\theta dv, & \text{otherwise.} \end{cases}$$

Once again, substituting the equilibrium price from (6), after some algebra, we obtain:

$$SW(\theta) = -\frac{c\theta^2}{2} + \begin{cases} -\frac{r^2(4-3\beta)}{8\beta\theta(1-\beta)} + r\left(\frac{1}{4} - \frac{e\xi(1-\beta)}{2}\right) + \frac{e^2\beta^2\theta\xi^2(1-\beta)}{8} + \frac{\theta(3+\beta)}{8}, & \text{Case 1A } (\beta < 1), \\ -\frac{r^2(4\beta-3)}{8\beta\theta(\beta-1)} - \frac{er\xi(\beta-1)}{2\beta} + \frac{e^2\theta\xi^2(\beta-1)}{8\beta} - \frac{e\theta\xi(\beta-1)}{2} + \frac{\beta\theta}{2}, & \text{Case 1A } (\beta > 1), \\ \frac{3e\beta\theta^2\xi}{8(\beta\theta-r)} + \frac{\beta^2\theta^2-r^2}{2\beta\theta} - \frac{e\xi(\beta\theta+r)}{2}, & \text{Case 1B.} \end{cases} \quad (11)$$

For completeness, we also perform comparative statics on the *legal social welfare* (SW_L), which excludes the surplus generated by illegal use, mostly because certain governments or policymakers may be interested in this surplus as a secondary metric. Considering the illegal surplus when choosing policy directions may, at times, be politically inconvenient. Furthermore, just as manufacturers do not like lost sales due to piracy, governments may also not like losses in tax revenues. Therefore, it is only practical that anti-piracy efforts are often moderated by political and pecuniary calculations.

4.2 Comparative Statics

Now, we proceed to answer our research questions regarding the economic impacts of the two potential anti-piracy instruments, namely, r and e . In general, we are interested in identifying whether the impact of changing one is more desirable than that of changing the other from the perspectives of the manufacturer, consumers, and policymakers. The manufacturer would, of course, want a higher profit; likewise, consumers would prefer a higher consumer surplus, and policymakers,

⁹In estimating the social surplus, we do not consider the cost of either type of piracy enforcement. This is because both r and e are exogenous parameters in the model. Thus, our modeling experiment does not aim to provide justification for additional investment in either type of enforcement activities. Our purpose is simply to find an answer to this question: if the government has already decided to invest in piracy enforcement, where is that marginal dollar better spent, on the supply or the demand side?

a higher social surplus. Further, it is likely that the manufacturer, and sometimes the policymakers, would also prefer to see a decrease in the piracy rate. To see if these expectations can be met by an anti-piracy instrument, we are particularly interested in whether the impact of changing r or e is uniformly favorable throughout the primary piracy region. Why do we seek such uniformity? After all, every product or its market is represented by a single point in our parameter space and, to understand how its online piracy is impacted by changes in enforcement, a simple point-specific comparative statics should have sufficed. However, the purpose of this research is to identify possible unintended consequences of different enforcement activities and to seek guidance accordingly. A recommendation can be useful only if the nature of its impact is not tied to a few specific products. Put differently, there is little comfort in knowing that a prescribed approach would work only for some products facing piracy, and not all, as having different anti-piracy policies in place at the same time could be quite impractical.

Theorem 1 (Short-Run Impact on Social Welfare) *In the short run, if $\beta < 1$, social welfare is not monotonic in the primary piracy region with respect to either r or e . On the other hand, if $\beta \geq 1$, social welfare is monotonically decreasing in both r and e throughout the primary piracy region.*

As either type of enforcement increases, piracy declines and, with it declines the consumer surplus. When $\beta \geq 1$, the negative impact of enforcement on the consumer surplus dominates, and social welfare decreases with both r and e . Interestingly, though, this relationship need not hold when $\beta < 1$. This may seem counterintuitive at first—for a product with zero marginal cost, the social surplus ought to increase with the overall consumption level, but stronger enforcement means only a reduction in consumption. Then, why should higher enforcement not always result in a smaller social surplus, just as it does for $\beta \geq 1$? Actually, when β is small, every conversion of an illegal consumption to a legal one generates a large upsurge (specifically, by a factor of $\frac{1}{\beta}$) in the surplus, which can offset the reduction in consumption.

Theorem 1 also shows that, for $\beta < 1$, neither enforcement policy uniformly dominates the other one. Each instrument has some positive impact, but only in parts of the primary piracy region. As mentioned above, the situation with $\beta \geq 1$ is even bleaker. There, both types of enforcement, supply- and demand-side, have a uniformly negative impact on social welfare. The implication is clear. When it is difficult to change the quality level of the product—for example, if the current state

of the technology imposes a quality ceiling—very strict piracy restrictions may have an unintended, and perhaps even undesirable, impact on social welfare. Enforcement authorities must consider such possibilities in formulating their directions. Now, what impact does piracy enforcement have on the secondary metrics? Our next set of results answers that question:

Proposition 3 (Short-Run Impact on Secondary Metrics) *In the short run, over the entire primary piracy region, the following relationships are observed:*

- i) The piracy rate is monotonically decreasing in both r and e .*
- ii) The manufacturer's profit is monotonically increasing in both r and e .*
- iii) The consumer surplus is monotonically decreasing in both r and e .*
- iv) The legal social welfare is monotonically increasing in both r and e .*

Proposition 3 has several important insights. First, the two types of enforcement are still indistinguishable in terms of their impacts. In other words, the secondary metrics, too, do not provide any clear anti-piracy directions for products with exogenous quality. Second, this indistinguishability aside, piracy enforcement seems to have a more desirable impact on the secondary metrics, with the only exception of consumer surplus. Third, Proposition 3 states that the consumer surplus decreases with enforcement even as the manufacturer's profit increases, suggesting that the extra legal surplus generated by additional enforcement activities accrues mostly to the monopolist manufacturer and is not duly shared with its consumers. This imbalance, sometimes coupled with additional inconveniences faced by consumers, perhaps explains why there have been so many instances of consumer discontent towards enforcement legislation and activities (e.g., Guarini 2013, Krebs 2005). Finally, all the findings in Proposition 3 consistently support the basic message in Theorem 1 that, in the short run, anti-piracy efforts, despite the best of intentions, can sometimes have undesirable consequences. Moderation and careful considerations are needed in formulating enforcement strategies towards online piracy.

We now turn our attention to the long-run situation, where we endogenize quality. Quality is often an important strategic decision for many manufacturers of digital products, and endogenizing it is necessary for obtaining insights applicable to these products. There is a second motivation as well—some policymakers may be specifically interested in learning the overall long-term impact of an enforcement choice, and, often times, long-term decisions center around the notion of innovation, which plays a crucial role in the development and design of a large majority of information

goods (Brynjolfsson and Zhang 2007). Economic history of nations provides sufficient testimony that the concept of innovation is important for public policy, as the former is often linked with an economy's rapid growth (cf. Grossman and Helpman 1993). Therefore, we start our analysis with the impact on innovation:

Theorem 2 (Impact on Innovation) *In the primary piracy region, the long-run equilibrium quality is always increasing in e but decreasing in r .*

In other words, we find that supply-side enforcement provides, to the manufacturers engaged in research, development, and design of information goods, added incentives to innovate and invest in quality. This result is quite surprising, especially in light of exactly the opposite impact from demand-side enforcement. Ultimately, any enforcement effort, supply- or demand-side, essentially makes it more difficult to pirate, one way or the other, resulting in a shrinkage of illegal use. Why then should they impact the equilibrium quality so differently?

The answer lies in how different enforcement types reveal themselves in the piracy ecosystem. Indeed, the shrinkage of illegal use is accompanied by an expansion of legal use, which affords the manufacturer a bigger base from which it can now recoup its investment in quality, providing incentives to invest more. Therefore, both r and e ought to have a positive impact on quality. And, they do. However, r also has a negative impact that is not shared by e . Curiously, for r , the negative impact dominates the positive, resulting in a quality level decreasing in r .

This negative impact is, in essence, rooted in the shadow competition from piracy within the $(1 - \lambda)$ segment, the intensity of which is measured by the incentive-compatibility (IC) threshold $\frac{p-r}{(1-\beta)\theta}$, a quantity that is decreasing in r but independent of e . Therefore, when there is an increase in demand-side enforcement, the pirated version appears less attractive to a potential consumer. A lessening shadow competition from the pirated copy is met with diminishing aggression from the manufacturer in the form of a lower quality, at either the same or a higher price-to-quality ratio. On the other hand, when supply-side enforcement is stepped up, the relative attractiveness of pirated content does not suffer; only its supply reduces, making it less available to a potential copyright violator. A portion of the consumer base, however, can still access the pirated copy. If the manufacturer wishes to remain competitive for this consumer segment, it can ill afford to drastically reduce the quality or hike the price-to-quality ratio. As mentioned above, irrespective of whether the manufacturer wants to cater to this segment or not, the presence of a bigger λ segment

now provides the manufacturer with a marginally better return from investing in quality and allows the manufacturer to leverage their “loyalty” by increasing the quality level. This is essentially at the crux of the model and its results.

Although the result in Theorem 2 is quite edifying in itself, we must still consider social welfare to be our primary metric:

Theorem 3 (Long-Run Impact on Social Welfare) *In the long run, if $\beta < 1$, social welfare is increasing in e over the entire primary piracy region but is not monotonic in r . For $\beta \geq 1$, social welfare is decreasing in r throughout the primary piracy region but is not monotonic in e .*

Evidently, in the primary piracy region, the two enforcement approaches have contrasting effects on social welfare. Theorem 3 states that, when $\beta < 1$, increasing supply-side enforcement has an overall positive impact on social welfare, but the impact from the demand side is not uniformly positive. In fact, our extensive numerical experiments suggest that, unless r or β is very small, this impact is largely negative and, as β increases beyond a threshold, social welfare actually becomes monotonically decreasing in r . Theorem 3 also states that, for $\beta \geq 1$, the impact of demand-side enforcement becomes uniformly negative. In contrast, for $\beta \geq 1$, the impact of e on social welfare is not monotonic—it is mostly positive, except when e is excessively large. Taken together, it appears that supply-side enforcement, overall, has a more desirable impact on social welfare when compared to its demand-side counterpart.

A point is in order. Although, for the sake of completeness, we provide analyses for both cases, $\beta < 1$ and $\beta \geq 1$, prior research has mostly considered the former (e.g., August and Tunca 2008, Bae and Choi 2006, Chellappa and Shivendu 2005, Chen and Png 2003, Jaisingh 2009, Lahiri and Dey 2013, Sundararajan 2004). The literature correctly recognizes that pirated products are usually inferior. Indeed, it is quite common for stolen product keys to not work, illegal copies of software to lack manufacturers’ support, or illegally downloaded movies or music to not have the same resolution. Given this reality, anti-piracy discussions are expected to focus primarily on this class of information goods for which $\beta < 1$. What we find here is that, for this predominant class, supply-side enforcement is not only better than its demand-side counterpart but is also uniformly desirable across the entire primary piracy region. Moreover, as we will show shortly, whether or not $\beta < 1$, supply-side enforcement increases the social welfare generated from legal sales (that is, welfare excluding pirates) throughout the primary piracy region, whereas demand-side enforcement again

falls short, further strengthening the appeal of supply-side measures from a long-run perspective.

The reason supply-side enforcement outperforms demand-side is also telling. As we have seen already, from a short-run perspective, they are similar. In other words, if we hold quality fixed, they would have similar impacts on both consumption and welfare. However, in the long run, quality is not fixed. In particular, as Theorem 2 establishes, θ is increasing in e but decreasing in r . Thus, every dollar spent on raising e instead of r actually generates more innovation. It is this innovation that enhances the value from consumption, eventually translating to a higher social welfare, legal or otherwise, making the two types of enforcement fundamentally different from a long-run perspective.

In order to complete the comparison between the two approaches, we also look at the impacts on the secondary metrics:

Proposition 4 (Long-Run Impact on Secondary Metrics) *In the long run, the following relationships are observed across the entire primary piracy region:*

- i) The piracy rate is monotonically decreasing in r , but it is not monotonic in e .*
- ii) The manufacturer's profit is monotonically increasing in both r and e .*
- iii) The consumer surplus is monotonically decreasing in r , but it is not monotonic in e .*
- iv) The legal social welfare is monotonically increasing in e , but it is not monotonic in r .*

There are several interesting observations that can be made from Proposition 4. First, even though the piracy rate decreases as r increases, the impact of e is not uniform, that is, in certain portions of the piracy region, the piracy rate could actually increase with e . This may seem surprising—the first order effect of a higher level of enforcement, regardless of its type, ought to be a reduction in the piracy rate. At least, that is what is widely believed. A quick Google search, for example, on “how to reduce the piracy rate” returns links to a large set of articles and news items, all of which point to higher levels of enforcement—a higher r or e , or both. To see the intuition behind this result, however, we must also consider the second order effect—the increase in quality resulting from a higher e has a positive impact on the illegal demand, which works in a direction contrary to the first order effect. As mentioned earlier, piracy rate is often used as a valuable yardstick in policy debates and in the legal parlance (Karaganis 2011). We find that this conventional wisdom is not always correct. It makes economic sense only in the short run, or only for demand-side enforcement, but a higher supply-side enforcement, in certain real-world

contexts, may actually result in a higher piracy rate in the long run. From another angle, if one is interested in merely curbing the long-run piracy rate, according to Proposition 4, enforcement efforts are perhaps better directed at the demand side. Thus, an enforcement activity geared towards reducing the long-run piracy rate may end up doing exactly that, but it might not be the socially desirable outcome after all.

Second, Proposition 4 echoes the sentiment shared by most manufacturers of digital goods that ramping either type of enforcement up is good for the industry and is corroborated by empirical findings in prior research (Danaher and Smith 2014, Danaher et al. 2014). For consumer advocates, however, it provides a brand new perspective. The long-held belief has been that piracy is beneficial to consumers in several ways. Piracy creates an opportunity to use the product free of charge, which, in turn, puts a pressure on the manufacturer to decrease the price (Lahiri and Dey 2013). Besides, certain restrictions (e.g., DRM) designed to deter piracy also reduce the utility of the legal product (Guarini 2013, Krebs 2005, Vernik et al. 2011). Naturally, consumers have often rallied against overzealous anti-piracy measures.¹⁰ Our result shows that such an attitude can be economically justified only in the short run, or only so far as demand-side enforcement is concerned. The story may change, though, when we turn to supply-side enforcement—an increase in e results in better quality and, in certain cases (specifically, if the current enforcement level is not excessively high), could actually improve consumer welfare in the long run. Viewed differently, consumer groups should try to understand that not all enforcement activities have the same economic impact and, if properly designed and executed, supply-side interventions can, in fact, enhance overall consumer welfare. Opposition to anti-piracy measures should be moderated.

Third, Proposition 4 also tells us that, if a policymaker is concerned about the legal social welfare in the context of piracy, supply-side enforcement is likely to provide a uniformly positive impact, whereas such impact from the demand side is often ambiguous. Thus, our results, taken together, find supply-side enforcement in a comparatively better light.

¹⁰A case in point is the DRM code used by music giant Sony-BMG on its music CDs; a massive consumer uproar was triggered by this action and, within days, Sony caved in by providing a patch to remove the DRM code (Krebs 2005).

4.3 Summary and Comparison

Clearly then, our comparative statics have important practical implications for manufacturers and consumers, along with some broader connotations for public policy. In trying to answer our research questions, we find that supply- and demand-side enforcement activities not only manifest themselves distinctly in the piracy ecosystem but also end up impacting different economic metrics very differently, especially in the long run. A side-by-side comparison of the results from Theorem 1 and Proposition 3 to those from Theorem 3 and Proposition 4 allows us to see the stark contrast between the short- and long-run impacts. To see this contrast more clearly, we now summarize our results as a scorecard in Table 2. Each cell in this table shows the direction of the impact of an anti-

Table 2: Economic Impacts of Demand- and Supply-Side Enforcement

Enforcement \longrightarrow Metric \downarrow	Demand-Side ($r \uparrow$)				Supply-Side ($e \uparrow$)			
	Short-Run		Long-Run		Short-Run		Long-Run	
	$\beta < 1$	$\beta \geq 1$	$\beta < 1$	$\beta \geq 1$	$\beta < 1$	$\beta \geq 1$	$\beta < 1$	$\beta \geq 1$
<i>Social Welfare</i>	$\uparrow\downarrow$	\downarrow^{**}	$\uparrow\downarrow$	\downarrow^{**}	$\uparrow\downarrow$	\downarrow^{**}	\uparrow^*	$\uparrow\downarrow$
<i>Piracy Rate</i>	\downarrow^*	\downarrow^*	\downarrow^*	\downarrow^*	\downarrow^*	\downarrow^*	$\uparrow\downarrow$	$\uparrow\downarrow$
<i>Product Quality</i>	—	—	\downarrow^{**}	\downarrow^{**}	—	—	\uparrow^*	\uparrow^*
<i>Manufacturer Profit</i>	\uparrow^*	\uparrow^*	\uparrow^*	\uparrow^*	\uparrow^*	\uparrow^*	\uparrow^*	\uparrow^*
<i>Consumer Surplus</i>	\downarrow^{**}	\downarrow^{**}	\downarrow^{**}	\downarrow^{**}	\downarrow^{**}	\downarrow^{**}	$\uparrow\downarrow$	$\uparrow\downarrow$
<i>Legal Social Welfare</i>	\uparrow^*	\uparrow^*	$\uparrow\downarrow$	$\uparrow\downarrow$	\uparrow^*	\uparrow^*	\uparrow^*	\uparrow^*
Overall Score	2	1	0	-1	2	1	4	3

NOTE: *direction desirable; **direction undesirable

piracy measure, either in the short or long run, for each of the economic metrics we study; whether such a direction is socially desirable or not is also indicated. These can be further consolidated into *overall scores* counted as follows: Each cell with a desirable impact contributes +1 to the overall score, but -1 if the impact is undesirable. When the directionality is unclear, no contribution is made. Just looking at this scorecard, it becomes apparent that supply-side enforcement may indeed be a superior alternative in the long run.

Several broad insights can be gleaned from the above results. First, although there is little to choose between the two enforcement types in terms of their impacts in the short run, their long-run impacts are widely different, with supply-side enforcement emerging as the more desirable one. The implication is clear from a perspective of anti-piracy strategy. Unless the long-run consequences

are properly taken into considerations, one may end up with a wrong policy choice. It is a relevant point, as it throws a bit of caution in the wind by pointing out that, while deploying anti-piracy efforts, failing to consider the long-run impacts and focusing only on the short-run ones could have an undesirable consequence. This dilemma is perhaps best illustrated by the French HADOPI Law. Enacted in 2009 to pursue illegal downloaders, this law has seen only a handful of cases taken to trial, with offenders being allowed to escape with a minimal fine, even though the law allows for significantly higher fines, service disconnections, and jail terms. Perhaps, sensing its undesirable impacts, authorities have dragged their feet in enforcing this short-sighted law, and indications are that, instead of pursuing illegal downloaders, French authorities are now targeting sites that profit from pirated material (Dato 2013).

Second, even though supply-side enforcement seems better in the long run, its short-term consequence might not be all that great. When adopting such a policy, a policymaker should be ready to see some temporary opposition due to welfare losses in the short run. A policymaker should stay the course till the long-run benefits start showing up. At the same time, consumer protection agencies and other lobbies should be aware of this distinction between the two types of enforcement and their long- and short-run impacts. They should recognize that supply-side enforcement may foster better innovation, lead to better quality products, and may actually enhance consumer welfare in the long run, despite some short-term pains. Such insights may help them bring moderation to their opposition to anti-piracy measures.

5 Costs of Enforcement

In reality, enforcement is not free. Designing and legislating copyright laws, monitoring network traffic, scanning potential sites for pirated content, finding, blocking, or shutting down illegal sites, apprehending and prosecuting offenders—all enforcement activities, supply- or demand-side—require substantial investments. Earlier, we ignored these enforcement costs and were able to conclude that supply-side enforcement has a more positive impact in an overall sense; such a conclusion was entirely based on the results from our comparative statics. In this section, we would like to study how robust that conclusion would be if these costs were included in our analysis and the enforcement levels were endogenized. In order to do so in a rigorous manner, we must include the policymaker as a fourth player who, at the very beginning of the game, takes the enforcement

costs into consideration and chooses r and e in such a way that maximizes the total social welfare net of the enforcement costs. In other words, conceptually, the policymaker solves the following problem:

$$\max_{r,e} SW(\theta^*) - (K_r r^2 + K_e e^2), \quad (12)$$

where, θ^* and $SW(\cdot)$ are as given in (7) and (11), respectively. In this formulation, we assume quadratic cost functions for both demand- and supply-side enforcement, with $K_r > 0$ and $K_e > 0$ being the respective cost parameters. Solving this optimization problem would essentially transform the equilibrium regions from the (r, e) space to the (K_r, K_e) space.

Before we proceed, a word of caution is in order. We note that the purpose of the optimization problem in (12) is not *normative*—we are not seeking to derive the optimal enforcement level (r^*, e^*) that a policymaker might want to actually implement in a real-world situation. The reason is simple. Both r and e are abstract concepts in our model, with unspecified scales or units. Furthermore, it is quite likely that the optimal values themselves are different for different information goods. It seems impractical to have widely varying enforcement strategies for different products in a market. Viewed this way, our model should not be used to derive how much investment should be made in enforcement activities; rather, it should only be used to see the impact of shifting resources from one side to the other. The purpose of this part of the analysis is, therefore, still *positive*—we would like to examine if the desirability of supply-side enforcement as an anti-piracy instrument continues to hold in a qualitative sense.

Although conceptually straightforward, the solution to the optimization problem in (12) is not fully tractable. We say “not fully” because, in fact, it is partially tractable. In three of the equilibrium regions, namely Regions 1A, 2A, and 3A—see Appendix ?? for details about these regions—the optimal quality, θ^* , is obtained by solving appropriate cubic equations. In each case, we can and do identify the correct root (from the three possible ones) and theoretically prove that it indeed exists. However, the sheer size of the expressions for those roots (in Mathematica) would immediately preclude them from an academic paper. That kind of size also makes it impossible to dissect it any further to identify useful analytical properties—a natural eventuality in any modeling experiment, a point is reached where it simply becomes impossible to go forward. To be sure, certain

limited-impact analytical observations can still be made,¹¹ but we do not consider them significant.

Not all is lost, however, because (12) can be solved numerically and, although a tedious undertaking, the solutions for different combinations of parameter values can be carefully checked and compared to test the veracity of our earlier results. To that end, we perform extensive numerical analyses with a wide variety of parameter value combinations and find that equilibrium regions themselves depend heavily on K_e . When K_e is small, the policymaker finds it optimal to completely eradicate piracy, leading to $\eta = 0$; the resulting equilibrium happens in Regions 3A and 3C. However, as K_e increases, the policymaker can no longer eliminate piracy fully and is forced to tolerate it to an extent. The equilibrium now shifts to the primary piracy region—Regions 1A or 1B—where $\eta > 0$. As K_e increases further and supply-side enforcement becomes prohibitively costly, the policymaker is forced to stop supply-side enforcement altogether, resulting in an abundant supply ($\eta = 1$) of pirated content in Regions 2A and 2B. Figure 3 illustrates this behavior in the (K_r, K_e) space for $\beta = \frac{3}{4}$, $\xi = \frac{2}{3}$, and $c = 0.01$.

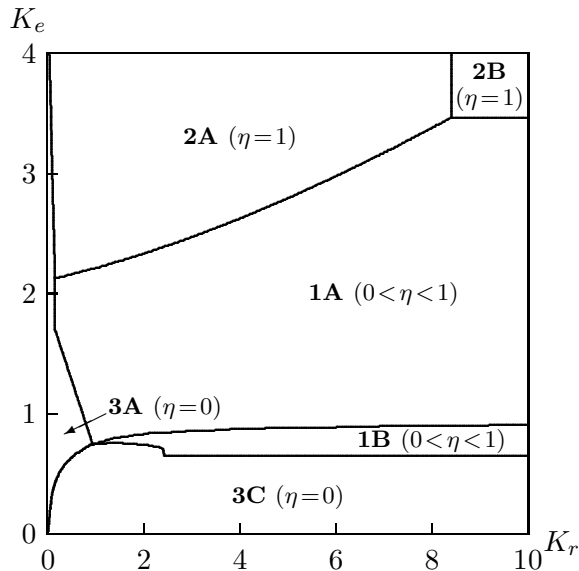


Figure 3: Relevant Partitions of the (K_r, K_e) Space; $\beta = \frac{3}{4}$, $\xi = \frac{2}{3}$, $c = 0.01$

Interestingly, even though the intuitive trend of η increasing with enforcement cost is observed for K_r as well, the influence of K_r on the equilibrium regions is a lot more subdued. In contrast,

¹¹For example, if $\beta > \frac{3}{4}$, throughout Region 1B—and if $\beta \geq 1$, throughout the entire primary piracy region— r^* is zero, implying a complete substitution of demand-side enforcement by its supply-side counterpart in those cases. Additionally, for $\beta < 1$, throughout Region 1A, r^* is capped at a theoretical ceiling, given by σ_{1A} in Technical Lemma ?? in Appendix ??.

K_e plays a much bigger role in defining the boundary of the primary piracy region. Region 1B is primarily defined by two thresholds of K_e and seems independent of K_r , except for low values of K_r . The boundaries of Region 1A are also determined largely by K_e thresholds, beyond low values of K_r . By carefully comparing the results from the solutions to (12) obtained with a wide set of parameter values, we are able to make the following observations:

- Although it now has the option to completely eradicate piracy, in a significant portion of the relevant parameter space, the policymaker is still willing to tolerate piracy to an extent ($0 < \eta < 1$), that is, the primary piracy region can still occur. This is consistent with the realities of today.
- Unless supply-side enforcement becomes prohibitively costly, the policymaker always employs it. More specifically, throughout the entire primary piracy region, e^* is strictly greater than zero. Clearly, when operating in the primary piracy region, supply-side enforcement is an indispensable instrument for the policymaker, even after accounting for its cost.
- In contrast, a policymaker may choose to ignore demand-side enforcement altogether in several parts of the parameter space, even when K_r is fairly low. More importantly, in a significant portion of the region of interest—the primary piracy region—we find that $r^* \approx 0$, suggesting that it is optimal for the policymaker to fully shift resources from the demand to the supply side of enforcement activities. This is not only true for all of Region 1B, but also for all but a small portion of Region 1A.

We construe that supply-side enforcement is more likely to be preferred by a policymaker, further corroborating our earlier findings.

6 Conclusion

In recent times, anti-piracy efforts aimed at curbing online piracy has seen a major shift in their focus from the demand to the supply side. The objective of demand-side enforcement, which has long been in use, is to reduce piracy by depressing the demand for pirated content. Such enforcement primarily involves making the pirated product less attractive by imposing penalties for illegal use. On the other hand, supply-side enforcement, which has lately started gaining popularity, does not

aim to shift the demand curve—it simply pushes the supply curve down, making pirated content less available. It involves combating piracy essentially at its source, for example, by limiting the reach of pirate suppliers through shutting down websites, filtering them out from search-engine results, and by prosecuting illegal content distribution. In this work, we compare these two types of enforcement in terms of their impacts on innovation and welfare, to answer whether a shift to the supply side indeed has merits.

We find that, although the two types of enforcement are somewhat indistinguishable with respect to their short-run impacts, there are some fundamental differences between their economic impacts in the long run. In situations where piracy exists, making the pirated product less available leads to an increase in the equilibrium quality of the legal product, whereas making piracy less attractive decreases this quality. This contrast is indeed fascinating, because both types of enforcement actually have similar effects of protecting the manufacturer's profit. More interestingly, in terms of social welfare as well, the effect of supply-side enforcement turns out to be more desirable in the long run, when compared to its demand-side counterpart.

Our basic model does not explicitly consider the costs incurred by a policymaker in implementing the two types of enforcement. For, in this research, we are only interested in exploring whether the policymaker, situated in a real-world context, might find it favorable to shift resources from the demand to the supply side. Comparative statics from our model indicate that, given a total budget for enforcement activities, the relative allocation of that budget should perhaps favor the supply side. In order to further test the veracity of this result, we later consider enforcement costs and endogenize the enforcement levels. The results from our original model are further corroborated by this exercise.

Our work is not without limitations. In reality, the ecosystem that supports online piracy is vastly more complex, often involving many legitimate business entities, such as online payment systems, search engine providers, and advertisement services; they all benefit from piracy (Seidler 2011). Thus, it is difficult to attribute legal liabilities to any specific entity. It is this difficulty that essentially makes the supply side of piracy a rather nebulous concept around which a definitive boundary is often difficult to draw. Moreover, many blog owners involved in piracy only provide links to pirated content but do not actually host any. Such instances create tensions between stopping online piracy and protecting freedom of speech. Clearly, all these issues need to be

examined carefully in order to better understand the supply side of piracy and related economic implications. Furthermore, our manufacturer—either individually on its own or collectively through an industry group—could attempt to influence enforcement activities. The extent of such an action, its direct impact on the enforcement levels, and the consequent spillover effects on the price and quality decisions are also among the issues that we do not explicitly address. Despite such limitations, this work would have achieved its goal if it has succeeded in providing a new economic lens through which enforcement activities can be viewed and evaluated.

References

- Aguiar, L., J. Claussen, C. Peukert. 2015. Online copyright enforcement, consumer behavior, and market structure. Tech. Rep. JRC93492, Institute for Prospective Technological Studies. URL https://ec.europa.eu/jrc/sites/jrcsh/files/JRC93492_Online_Copyright.pdf. Accessed on August 1, 2016.
- August, T., T.I. Tunca. 2008. Let the pirates patch? An economic analysis of network software security patch restrictions. *Information Systems Research* **19**(1) 48–70.
- Bae, S.-H., J.P. Choi. 2006. A model of piracy. *Information Economics and Policy* **18**(3) 303–320.
- Baird, A., C.J. Miller, T.S. Raghu, R.K. Sinha. 2016. Product line extension in consumer software markets in the presence of free alternatives. *Information Systems Research* **27**(2) 282–301.
- Bakos, J.Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science* **43**(12) 1676–1692.
- Baye, M.R., J. Morgan, P. Scholten. 2006. Information, search and price dispersion. T. Hendershott, A.B. Whinston, eds., *Handbooks in Information Systems*, vol. 1, chap. 6. Emerald, Bingley, UK, 323–376.
- Brynjolfsson, E., X. Zhang. 2007. Innovation incentives for information goods. Adam B. Jaffe, Josh Lerner, Scott Stern, eds., *Innovation Policy and the Economy*, vol. 7. National Bureau of Economic Research, 99–123.
- Chellappa, R.K., S. Shivendu. 2005. Managing digital piracy: Pricing and sampling strategies for digital experience goods in vertically segmented markets. *Information Systems Research* **16**(4) 400–417.
- Chen, Y., I. Png. 2003. Information goods pricing and copyrights enforcement: Welfare analysis. *Information Systems Research* **14**(1) 107–123.
- Cho, W.-Y., B.-H. Ahn. 2010. Versioning of information goods under the threat of piracy. *Information Economics and Policy* **22**(4) 332–340.
- Clement, M., A. Rangaswamy, S. Vadali. 2012. Consumer responses to legal music download services that compete with illegal alternatives. *Service Science* **4**(1) 4–23.
- Conner, K.R., R.O. Rumelt. 1991. Software piracy: An analysis of protection of strategies. *Management Science* **37**(2) 125–139.
- Danaher, B., S. Dhanasobhon, M.D. Smith, R. Telang. 2010. Converting pirates without cannibalizing purchasers: The impact of digital distribution on physical sales and internet piracy. *Marketing Science* **29**(6) 1138–1151.
- Danaher, B., M.D. Smith. 2014. Gone in 60 seconds: The impact of the MegaUpload shutdown on movie sales. *International Journal of Industrial Organization* **33** 1–8.
- Danaher, B., M.D. Smith, R. Telang. 2015. The effect of piracy website blocking on consumer behavior. *Social Science Research Network* URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2612063. Accessed on January 17, 2016.
- Danaher, B., M.D. Smith, R. Telang. 2016. Website blocking revisited: The effect of the UK November 2014 blocks on consumer behavior. *Social Science Research Network* URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2766795. Accessed on August 30, 2016.

- Danaher, B., M.D. Smith, R. Telang, S. Chen. 2014. The effect of graduated response anti-piracy laws on music sales: Evidence from an event study in France. *The Journal of Industrial Economics* **62**(3) 541–553.
- Dattoo, S. 2013. France drops controversial ‘Hadopi law’ after spending millions. *The Guardian* URL <http://www.guardian.co.uk/technology/2013/jul/09/france-hadopi-law-anti-piracy>. Accessed on August 21, 2013.
- Davidson, C., R. Deneckere. 1986. Long-run competition in capacity, short-run competition in price, and the Cournot model. *The RAND Journal of Economics* **17**(3) 404–415.
- de Beer, J., C.D. Clemmer. 2009. Global trends in online copyright enforcement: A non-neutral role for network intermediaries? *Jurimetrics Journal* **49**(4) 375–409.
- Diola, C. 2013. 3 PC shops busted for loading pirated software. *The Philippine Star* URL <http://www.philstar.com:8080/nation/2013/10/14/1245204/3-pc-shops-busted-loading-pirated-software>. Accessed on September 23, 2014.
- Duchene, A., P. Waelbroeck. 2005. Peer-to-peer, piracy and the copyright law: Implications for consumers and artists. L.N. Takeyama, W.J. Gordon, R. Towse, eds., *Developments in the Economics of Copyright*, chap. 4. Edward Elgar Publishing, Northampton, MA, 60–79.
- Duhigg, C. 2006. Is copying a crime? Well... *Los Angeles Times* URL <http://www.latimes.com/news/la-fi-pollmusic9aug09-story.html#page=1>. Accessed on September 23, 2014.
- Epstein, Z. 2012. Founder of illegal movie streaming site sentenced to 4.5 years in jail, \$4.7 million fine. URL <http://news.yahoo.com/founder-illegal-movie-streaming-sentenced-4-5-years-202537619.html>. Accessed on February 9, 2014.
- Farivar, C. 2013. Here’s what an actual “six strikes” copyright alert looks like. URL <http://arstechnica.com/tech-policy/2013/02/heres-what-an-actual-six-strikes-copyright-alert-looks-like/>. Accessed on February 9, 2014.
- Geng, X., Y-J. Lee. 2013. Competing with piracy: A multichannel sequential search approach. *Journal of Management Information Systems* **30**(2) 159–184.
- Gopal, R.D., A. Gupta. 2010. Trading higher software piracy for higher profits: The case of phantom piracy. *Management Science* **56**(11) 1946–1962.
- Grossman, G.M., E. Helpman. 1993. *Innovation and Growth in the Global Economy*. MIT press, Cambridge, MA.
- Guarini, D. 2013. Microsoft drops internet, used game restriction on Xbox one after huge gamer backlash. *The Huffington Post* URL http://www.huffingtonpost.com/2013/06/19/xbox-one-always-online_n_3468098.html. Accessed on February 9, 2014.
- Hashim, M.J., K.N. Kannan, S. Maximiano, J.R. Ulmer. 2014. Digital piracy, teens, and the source of advice: An experimental study. *Journal of Management Information Systems* **31**(2) 211–244.
- Horwitz, J. 2013. China cracks down on online piracy and removes two major sites for illegal downloads. URL <http://thenextweb.com/asia/2013/04/26/chinas-top-two-movie-piracy-sites-removed-by-authorities/>. Accessed on May 8, 2013.
- Jain, S. 2008. Digital piracy: A competitive analysis. *Marketing Science* **27**(4) 610–626.
- Jaisingh, J. 2009. Impact of piracy on innovation at software firms and implications for piracy policy. *Decision Support Systems* **46**(4) 763–773.
- Johar, M., N. Kumar, V. Mookerjee. 2012. Content provision strategies in the presence of content piracy. *Information Systems Research* **23**(3) 960–975.
- Kan, M. 2012. Microsoft: Most PCs running pirated windows in China have security issues. *Computerworld* URL <http://www.computerworld.com/article/2493787/malware-vulnerabilities/microsoft--most-pcs-running-pirated-windows-in-china-have-security-issues.html>. Accessed on May 3, 2014.
- Kannan, K., M.S. Rahman, M. Tawarmalani. 2016. Economic and policy implications of restricted patch distribution. *Management Science* **62**(11) 3161–3182.
- Karaganis, J. 2011. Rethinking piracy. J. Karaganis, ed., *Media Piracy in Emerging Economies*. Social Science Research Council, 1–73. URL <http://piracy.ssrc.org/the-report/>. Accessed on May 8, 2013.

- Khantwal, S. 2016. Best top 10+ websites to download free movies. URL <https://www.infogeekers.com/websites-to-download-free-movies-no-registration/>. Accessed on March 13, 2016.
- Krebs, B. 2005. Study of Sony anti-piracy software triggers uproar. *The Washington Post* URL <http://www.washingtonpost.com/wp-dyn/content/article/2005/11/02/AR2005110202362.html>. Accessed on May 8, 2013.
- Lahiri, A., D. Dey. 2013. Effects of piracy on quality of information goods. *Management Science* **59**(1) 245–264.
- Liebelson, D. 2014. Why it's getting harder to sue illegal movie downloaders. *Mother Jones* URL <http://www.motherjones.com/politics/2014/02/bittorrent-illegal-downloads-ip-address-lawsuit>. Accessed on January 10, 2014.
- Mankiw, N.G. 2008. *Principles of Economics*. 5th ed. South-Western, Mason, OH.
- Mas-Colell, A., M.D. Whinston, J.R. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York, NY.
- Masnack, M. 2012. MPAA To USTR: More shutdowns like MegaUpload, please. URL <https://www.techdirt.com/articles/20121205/16451021247/mpaa-to-ustr-more-shutdowns-like-megaupload-please.shtml>. Accessed on May 8, 2013.
- Mills, E. 2012. Court affirms \$675,000 penalty in music-downloading case. URL [http://news.cnet.com/8301-13578_3-57499519-38/court-affirms-\\$675000-penalty-in-music-downloading-case/](http://news.cnet.com/8301-13578_3-57499519-38/court-affirms-$675000-penalty-in-music-downloading-case/). Accessed on May 8, 2013.
- Moorthy, K. S. 1984. Market segmentation, self-selection, and product line design. *Marketing Science* **3**(4) 288–307.
- Morris, R.G., G.E. Higgins. 2010. Criminological theory in the digital age: The case of social learning theory and digital piracy. *Journal of Criminal Justice* **38**(4) 470–480.
- Mussa, M., S. Rosen. 1978. Monopoly and product quality. *Journal of Economic Theory* **18**(2) 301–317.
- Rob, R., J. Waldfogel. 2006. Piracy on the high Cs: Music downloading, sales displacement, and social welfare in a sample of college students. *The Journal of Law & Economics* **49**(1) 29–62.
- Seidler, E. 2011. Cyberlockers: Explaining piracy's profit pyramid. URL <http://popuppirates.com/?p=1249>. Accessed on May 8, 2013.
- Sundararajan, A. 2004. Managing digital piracy: Pricing and protection. *Information Systems Research* **15**(3) 287–308.
- Tunca, T.I., Q. Wu. 2013. Fighting fire with fire: Commercial piracy and the role of file sharing on copyright protection policy for digital goods. *Information Systems Research* **24**(2) 436–453.
- Veneziani, V. 2007. Help key: The essential guide to piracy. URL <http://techcrunch.com/2007/05/14/help-key-the-essential-guide-to-piracy/>. Accessed on November 17, 2015.
- Vernik, D.A., D. Purohit, P.S. Desai. 2011. Music downloads and the flip side of digital rights management. *Marketing Science* **30**(6) 1011–1027.
- Vickers, J. 1989. The nature of costs and the number of firms at Cournot equilibrium. *International Journal of Industrial Organization* **7**(4) 503–509.
- Wolinsky, A. 1983. Retail trade concentration due to consumers' imperfect information. *The Bell Journal of Economics* **14**(1) 275–282.
- Wu, S.-Y., P.-Y. Chen. 2008. Versioning and piracy control for digital information goods. *Operations Research* **56**(1) 157–172.