Combating Online Piracy: The "Longer Arm" of Enforcement

Antino Kim (antino@uw.edu)
Debabrata Dey (ddey@uw.edu)
Atanu Lahiri (lahiria@uw.edu)

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

Last revised: February 28, 2014

Abstract

Combating online piracy, a global menace facing the manufacturers of information goods, has remained a top priority for those manufacturers, as well as for governments around the world. Despite much stricter anti-piracy measures in recent times, this menace continues unabated. Apparently, the policy debate is now shifting from the efficacy of the existing laws towards the need for enacting new ones such as the SOPA/PIPA. Our interest is at the core of this debate. We argue that, in order to develop relevant economic insights, one must understand the intricate interrelationships within a piracy ecosystem and how different enforcement approaches impact them. Based on these impacts, we propose a clear distinction between efforts that restrict online supply of pirated goods (*supply-side* enforcement) and those that penalize illegal consumption (*demand-side* enforcement). We find that, indeed, there are some fundamental differences between the two approaches 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 much more desirable impact in the long run. Our results have clear implications for manufactures and consumers, along with broader connotations for public policy and law.

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

1 Introduction

Online 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. In the software sector alone, the lost sales due to piracy, which was about US $33 billion globally in 2004, has nearly doubled in the next seven years, crossing the US $60 billion mark in 2011 (BSA 2011). Figure 1 shows how the piracy losses in the software industry have grown between 2004 and 2011; it is interesting to note that the lost sales from piracy have grown every year in this time period, except in 2009 when the global economy was experiencing a serious downturn. And, it is not just software! Given the rapid proliferation of broadband and mobile networks, essentially anything that can be digitally stored and logically rendered can also be pirated. In fact, today, digital piracy is a significant component of the global counterfeiting industry that is estimated at a whopping US $600 billion annually, accounting for 5–7% of all global trade (Bitton 2012).

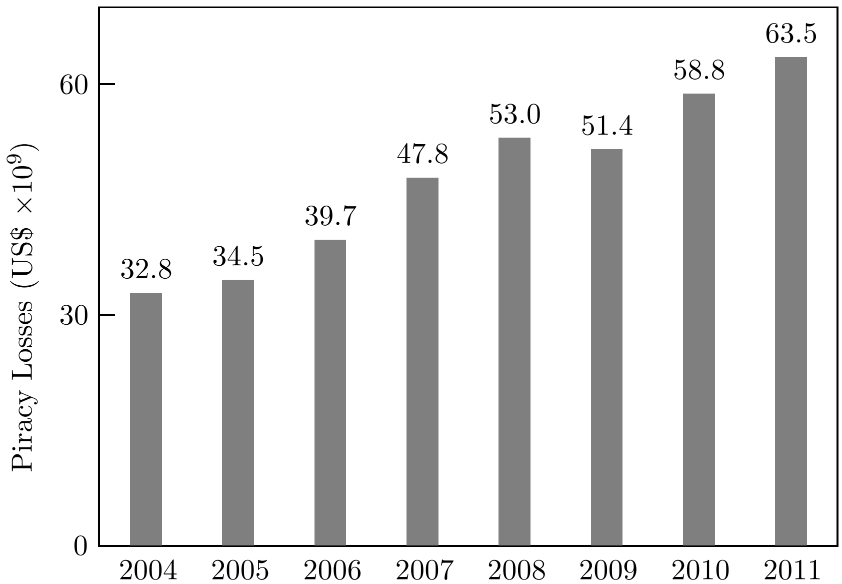

Figure 1: Total Lost Sales from Software Piracy Worldwide (Source: BSA 2011)

Naturally, to counter this growing menace, a lot of effort has been expended globally, by many governments, by manufacturers of information goods, and by different business alliances and industry lobbies. These efforts—we will call them anti-piracy measures or enforcement efforts—take on different forms. First and foremost, enforcement efforts have centered around bringing pirates to justice within the existing legal framework of a country. In 2003, for example, one of the biggest news items concerning online piracy was the thousands of lawsuits brought by the Recording Industry Association of America (RIAA) against illegal music downloaders (Bhattacharjee et al. 2006). Even though, after less than six years, the RIAA announced that it would "stop suing for illegal downloading" (Stern 2009), some manufacturers are still trying to discourage piracy by prosecuting illegal users (BBC 2012, Moon 2012). In addition, there has been a marked increase in the lawsuits brought directly by many governments. For instance, the US Copyright Group (USCG) filed two lawsuits against almost 50,000 alleged BitTorrent users in total, for downloading one of the two movies, *The Expendables* and *The Hurt Locker* (Pepitone 2011). While both cases were later dismissed, they each were a record-setting lawsuit against illegal downloaders (Burgess 2012). Similar efforts have also been observed in recent years in many other countries. For example, government agencies in Canada and S. Korea have started performing random audits of companies, universities, schools, and other government agencies, to detect software piracy and levy a fine on the offending party when a copyright violation is detected (Jackson 2009, Lee 1999).

Piracy enforcement approaches have also evolved. Intelligent monitoring of online activities and appropriate deterrent steps against offenders have been employed in Australia, New Zealand, Canada, and the US (Coutts 2013, Ducklin 2013, Farivar 2013, Hall 2013, LeMay 2013). In more recent times, governments have also started scanning for sites that distribute or aid in the distribution of pirated content. Often, governments have gone after these sites forcefully, shutting them down and prosecuting them (Epstein 2012, Perry 2012, Seidler 2013). The most prominent example, of course, is the recent shutdown of `MegaUpload.com` (Danaher et al. 2012). Even when they have been spared from being brought down completely, such pirate sites have faced significant downtime and uphill legal battles.

Despite such enforcement efforts, there is a growing concern that these efforts have largely failed to mitigate the problem of online piracy (Mick 2011, PCWorld 2011). Business Software Alliance (BSA 2011), for example, estimates that the global piracy rate for software currently stands at 42%, and the actual rate is much higher in several parts of the world. Perhaps, sensing that they might have gone amiss somewhere, governments around the world have also started considering new legislation aimed specifically at online piracy. For example, the US lawmakers have quickly come to recognize that the *Digital Millennium Copyright Act* (DMCA) of 1998 is simply not adequate in combating piracy effectively and, at different times, have considered several additional pieces of legislation: the *Protecting Intellectual Rights Against Theft and Expropriation* (PIRATE) *Act* in 2004, *Combating Online Infringements and Counterfeits Act* (COICA) in 2010, and *Stop Online Piracy Act* (SOPA) and *Protect IP Act* (PIPA) in 2012. These bills were all vigorously supported by the manufacturers and industry lobbies, such as the RIAA and MPAA (Motion Pictures Association of America). Although none of these bills could eventually be enacted into a law, primarily because of strong opposition from consumer lobbies and "netizens" (Bachman 2011), the debates have been long and are still continuing (Garrahan 2013, Gelles 2012).[1]

Our interest is at the core of this policy debate. We find that, even though there is a significant increase in legislative efforts, there has been no overall consensus about the directions for such efforts; more specifically, there has been little research that uses economic analysis as an overall guiding principle. In a broader sense, we are interested in finding answers to the following set of

[1]Of course, these are not isolated instances and are not happening in the US alone! Over the last few years, there has been a significant growth in such legislative activities around the globe, as well. One could cite the examples of New Zealand's "three strikes" law against illegal downloading (Ducklin 2013), the French HADOPI law (Danaher et al. 2012), the new anti-piracy law in Russia (Holdsworth 2013) and in Japan (Gastaldo 2012), just to name a few.

questions:

- Should manufacturers work within the existing legal framework, or should they lobby for changes?
- Should governments emphasize enforcing current laws or debate the need for new ones?
- How should consumer advocacy groups lobby to advance consumer interests in this debate?
- If new laws are indeed needed, what should be the direction or shape for them to take?
- And, what, if any, are the short- and long-run implications?

Answering these questions, we believe, would allow one to bring some economic insights to this whole debate and develop a general guiding principle for combating piracy.

In trying to answer these questions, we find it important to recognize that different types of approaches towards mitigating piracy may manifest themselves quite differently in a socioeconomic system and feel the necessity to distinguish them as such. This is in stark contrast with prior literature, where different types of enforcement are modeled uniformly—every enforcement effort is assumed to either enhance the probability of detecting illegal consumption or create the perception of a higher penalty on being detected. This way, enforcement increases the expected legal liability, or the piracy cost, associated with illegal consumption (August and Tunca 2008, Lahiri and Dey 2013). This approach, however, oversimplifies the research context, and we find it lacking in the finesse necessary to answer the above questions. A case in point is the situation described in Figure 2. This figure shows an implementation of DMCA, where a request from Viacom has forced YouTube to stop delivering the content sought by an end user. To be clear, such an effort does diminish the visibility of pirated content and abates piracy.[2] Yet, it does not translate to an increase in the piracy cost to a user—in fact, there is no piracy taking place here, and the consumer faces no legal liabilities whatsoever.

Indeed, a careful analysis reveals that enforcement efforts, both legislative and prosecutorial, are fundamentally of two types. One requires going after illegal downloaders or pirate consumers, and the other, stepping up actions against the suppliers of illegal digital goods operating through

[2]Interestingly, the primary intent of SOPA was also quite similar, leveraging the filtering capabilities of popular search engines to curb piracy. In fact, the most controversial provision in the law actually sought to make the search providers legally accountable for providing links to pirated contents.

Figure 2: DMCA-Induced Filtering of Pirated Content by Provider

cyberlockers and the like. On one hand, penalizing illegal consumption, or *demand-side enforcement*, makes piracy a more costly and less attractive proposition to potential consumers. On the other, *supply-side enforcement*, which seeks to limit the reach of illegal suppliers, makes pirated content less available.

Therefore, to fully grasp the impact of enforcement efforts, we must first classify them in a manner shown in Table 1. All the examples cited previously, along with several new ones, are classified in this table based on whether they make pirated content less attractive or less available. As evident from the examples in this table, both forms have garnered significant attention in recent years. Yet, the issue of their relative economic merits remains elusive. A major point of this paper is to compare the economic implications of these two enforcement approaches and verify if one 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. How do different types of enforcement alter the manufacturer's incentive to innovate? Recent research shows that, often, the manufacturer's reaction to demand-side enforcement is unpredictable and counterintuitive. For example, Lahiri and Dey (2013) show that the manufacturer may, in fact, respond to higher levels of demand-side enforcement by unexpectedly decreasing the quality of the product, contradicting the common argument advanced by manufacturers that stronger enforcement is necessary to foster innovation (Adobe 2013). We seek to verify whether the impact of supply-side enforcement is also qualitatively similar and, if not, why, when, and how it differs.

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

Demand Side	**Supply Side**
New Zealand's "three strikes" law against illegal downloaders (Ducklin 2013)	MegaUpload.com shutdown for online piracy (Danaher and Smith 2013)
Copyright Alert System (CAS) (also called "six strikes") in the US (Farivar 2013)	China shutting down two major sites involved in online piracy (Horwitz 2013)
Canipre, a Canadian intellectual property rights company, tracking down illegal downloaders (Coutts 2013)	French government penalizing "sites that profit from pirated material" (Datoo 2013)
Mass scanning of IP addresses in Australia (LeMay 2013) and the US (Hall 2013, Pepitone 2011)	New anti-piracy law in Russia to tackle pirate sites (Holdsworth 2013)
Enactment of the French HADOPI law (Danaher et al. 2012)	Founder of illegal movie streaming site sentenced to 4.5 years in jail, $4.7 million in fine (Epstein 2012)
Boston University graduate student fined $675,000 for illegally downloading 30 songs (Mills 2012)	Google voluntarily playing copyright cop, suppressing violators in search results (Reed 2012)
Minnesota woman fined $220,000 for illegally downloading 24 songs (Holpuch 2012)	Google dropping *Pirate Bay* from auto-complete results (Woollacott 2012)
Enactment of Japan's new law punishing illegal downloaders to a jail-term of up to 2 years (Gastaldo 2012)	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 (Jackson 2009)	Combating Online Infringements and Counterfeits Act (COICA) of 2010
Random audits in S. Korea for software piracy at companies, universities, and government agencies (Lee 1999)	"Operation In Our Sites" initative in the US resulting in the seizure of 125 websites (Gardella and Forzato 2011)
	Digital Millennium Copyright Act (DMCA) of 1998

Answering these questions is of much practical significance—for a manufacturer facing lost sales from piracy, it is essential to understand what these two approaches mean for its product-line and anti-piracy strategies.

It goes beyond, however! Since the issue of innovation is invariably intertwined with that of social welfare, the difference between welfare implications of these two types of enforcement is also unclear, and a careful analysis of consumer and social welfare is critical to assessing their relative appeal and framing public policy debates. In this paper, we build a parsimonious framework, incorporating several elements from prior research on piracy, to address these issues. Our analyses provide a clear comparison of these two enforcement types—in terms of their impacts on innovation and welfare—and our results offer important managerial implications and policy guidelines.

In studying these two different approaches, we intentionally borrow much of the model 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 extension of the standard setup used in prior literature, the new elements being the advertisement-dependent pirated content suppliers and the presence of supply-side enforcement.

Our research questions and modeling approach, therefore, put the spotlight squarely on how suppliers of pirated content actually operate, unraveling the mystery surrounding their "business" models. The underlying ecosystem shows curious interdependence among three key entities: aggregator sites, online ad agencies, and pirate suppliers. 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 to its site, a cyberlocker incentivizes individual pirate suppliers to share content that will be in high demand, as measured by the number of downloads. See, for example, the compensation rates for `FileServe.com` given in Table 2. Pirate suppliers leveraging such a site

Table 2: `FileServe.com` Rates for Each 1,000 Download of A File

	Country Groups			
Size	**A**	**B**	**C**	**D**
1 – 50 MB	\$3.00	\$2.00	\$1.00	\$0.50
50 – 100 MB	\$8.00	\$6.00	\$4.00	\$1.25
100 – 200 MB	\$15.00	\$13.00	\$10.00	\$3.00
200 – 450 MB	\$20.00	\$17.00	\$14.00	\$4.00
450 – 2048 MB	\$25.00	\$22.00	\$17.00	\$5.00

NOTE: Rates as of January 10, 2011.

typically offer free pirated content to lure hoards of downloaders to their respective pages. The number of downloaders attracted by a supplier, in turn, has a direct impact on its "revenue" funneled through the aggregator site. Indeed, cyberlockers have often been identified as an important part of this ecosystem, which the suppliers of illegal digital goods routinely use to distribute their content and generate revenue (Seidler 2011).[3] It is the explicit modeling of this ecosystem that sets our work apart from the current literature in marketing, economics, public policy, and law.

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. Although existing research does not explicitly distinguish the two, it does provide some interesting cues. For example, Danaher et al. (2012) 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. On the other hand, Danaher and Smith (2013) study the effects of shutting down `MegaUpload.com`—clearly a case of supply-side enforcement—and estimate that this shutdown resulted in a 6–10% boost in sales of digital movies. Although empirical studies can spot such immediate impact 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 efforts that curb the demand for pirated content from those that restrict its supply.

There is also a vast literature on the economics of piracy using quantitative models. A branch of this literature argues that piracy may surprisingly benefit the manufacturers of digital goods in the presence of network effects (August and Tunca 2008, Conner and Rumelt 1991). Another branch examines when and how certain tools—digital rights management (DRM), nonlinear pricing, versioning, 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, Sundararajan 2004, Vernik et al. 2011, Wu and Chen 2008). The closest to our work is the branch that examines the economic impacts of anti-piracy efforts on a manufacturer's strategy and 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 reducing illegal demand. 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 we *endogenize the supply* of

[3] It is important to note that supply-side enforcement can have an impact on the entire ecosystem or some parts of it. In other words, such enforcement can reduce the availability of pirated content by making it difficult for the individual pirate suppliers, for the host pirate site, or both.

pirated content, with a clear objective of comparing economic implications of enforcement efforts on the supply side with those on the demand side.

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 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 enforcement is weak, 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. Interestingly, it does not extend to the supply side—the impact of supply-side enforcement on innovation is often exactly the opposite, and its impacts 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 that sustains online piracy, nor do they analyze the economic implications of disrupting this ecosystem.

3 Model Preliminaries

We develop an economic model with three strategic players: (i) a profit-maximizing monopolist, (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 is situated within a market with certain levels of demand- and supply-side enforcement and chooses the price and quality of an information good.

The timeline is as follows: First, the manufacturer offers a product of quality $\theta > 0$ at a price

$p > 0$. The suppliers of pirated content then decide whether to provide an illegal version or not—only when potential revenues from piracy can 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; this decision depends on the piracy cost resulting from demand-side enforcement as well as on the availability of a pirated copy. When an equilibrium is reached, the demand for pirated content matches its supply. As is customary, we traverse the timeline backwards, starting with the behavior of consumers.

3.1 A Model of Consumer Behavior

Our consumers are heterogeneous along two orthogonal dimensions and are indexed by the pair: $\langle v, k\rangle$. Consumers' preference for quality is represented by v; consumer v gets a value of $v\theta$ from using a product of quality θ. On the other hand, k stands for the maximum number of searches a consumer is willing to perform and represents his propensity to search.[4] We make the following assumption about v and k:

Assumption 1 *A consumer's preference for quality, v, is uniformly distributed over* $[0, 1]$. *A consumer's propensity to search, k, follows the geometric distribution with parameter $g \in (0, 1)$: for all $i = 0, 1, 2, \ldots$, $\Pr[k = i] = g(1 - g)^i$. A consumer knows his v and k, whereas the manufacturer only knows their distributions.*

The first part of Assumption 1—uniform distribution of v—is commonplace in economic modeling and needs little justification; the second—geometric distribution of k—is also quite reasonable. In reality, a large majority of consumers makes only a few search attempts, but a small fraction searches extensively. For instance, online "shoppers search across very few sites" (Johnson et al. 2004), and "75% of users never scroll past the first page of search results" (Siu 2012). Although the geometric distribution fits this reality quite well, this assumption is not critical; any valid distribution of k—decaying or not—provides similar results.

[4]An astute reader will see the obvious connection between the propensity to search and the well-known notion of a "search cost" (Stahl 1989, Varian 1980). In essence, one is just the dual of the other. After every failed search attempt, a consumer is likely to downgrade his chance of finding what he is looking for and will stop searching unless the chance is still sufficient to offset his search cost. Hence, a consumer with a higher search cost will stop sooner. In other words, the intuitive relationship that a higher search cost should imply a lower propensity to search is indeed preserved in our modeling setup.

Examples abound where the physical quality of the pirated copy is less than that of the original, such as in the case of 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 (Jaisingh 2009). To capture this, we follow the prior literature and assume that the quality of the pirated product is less than that of the legal version (Lahiri and Dey 2013, Sundararajan 2004):

Assumption 2 *When the quality of the legal product is θ, the quality of its pirated version is $\beta\theta$, where $\beta \in (0, 1)$.*

Consumers' decision to buy, pirate, or forgo use is determined by both the availability and attractiveness of the pirated version. If he does not have access to a pirated version, then he would buy only if his utility from buying is more than forgoing use. If, on the other hand, a pirated version is available to him—either because he was lucky in the first few searches or because he has a high propensity to search that allowed him to search extensively—he will have to compare the utility of using this pirated version with those of buying and forgoing use.

A consumer considering piracy faces an expected legal penalty of r if he ends up using an illegal copy; this penalty is exogenous in our model—it simply depends on the level of enforcement against the consumption of pirated goods (August and Tunca 2008, Lahiri and Dey 2013). This way, consistent with prior literature, our r is essentially a proxy for the level of *demand-side* enforcement. Clearly, a consumer can enjoy a net utility of $(v\theta - p)$ from purchasing the legal version and $(v\beta\theta - r)$ from a pirated copy.

Let us first consider consumers who do not have access to a pirated copy. These are the consumers either *inherently* "ethical" (i.e., $k = 0$) or those *forced* to remain "ethical" because of their inability to locate a pirated copy within their maximum number of searches. Such a consumer would buy the legal product if and only if his IR constraint is met:

$$v\theta - p \geq 0 \Rightarrow v \geq \frac{p}{\theta}. \tag{IR-L}$$

Next, we consider consumers for whom piracy is indeed an option. Such a consumer would choose to use the legal product if the following IC constraint is satisfied, in addition to (IR-L):

$$v\theta - p \geq v\beta\theta - r \Rightarrow v \geq \frac{p-r}{(1-\beta)\theta}. \tag{IC-L}$$

Finally, this consumer resorts to piracy if:

$$v\beta\theta - r \geq 0 \Rightarrow v \geq \frac{r}{\beta\theta}, \quad \text{and} \tag{IR-P}$$

$$v\beta\theta - r > v\theta - p \Rightarrow v < \frac{p-r}{(1-\beta)\theta}. \tag{IC-P}$$

The availability of pirated content, denoted $\lambda \in [0,1]$, is essentially the probability that a randomly chosen consumer will not have access to a pirated copy and will thus be "ethical." The legal and illegal demands, denoted q and $\hat{q}$, respectively, can then be written as:

$$\begin{aligned} q &= \lambda\left(1 - \min\left\{1, \frac{p}{\theta}\right\}\right) + (1-\lambda)\left(1 - \min\left\{1, \max\left\{\frac{p}{\theta}, \frac{p-r}{(1-\beta)\theta}\right\}\right\}\right), \text{ and} \\ \hat{q} &= (1-\lambda)\left(\min\left\{1, \frac{p-r}{(1-\beta)\theta}\right\} - \min\left\{1, \frac{r}{\beta\theta}, \frac{p-r}{(1-\beta)\theta}\right\}\right). \end{aligned}$$

In order to avoid trivialities, our analysis will focus on regions that satisfy $\frac{r}{\beta\theta} < \frac{p}{\theta} < 1$, ensuring that there is a market for both the legal and illegal versions. This allows us to simplify the quantity demanded (see Figure 3):

$$q = \lambda\left(1 - \frac{p}{\theta}\right) + (1-\lambda)\left(1 - \min\left\{1, \frac{p-r}{(1-\beta)\theta}\right\}\right), \text{ and} \tag{1}$$

$$\hat{q} = (1-\lambda)\left(\min\left\{1, \frac{p-r}{(1-\beta)\theta}\right\} - \frac{r}{\beta\theta}\right). \tag{2}$$

3.2 Supply-Side Enforcement and Behavior of Pirate Suppliers

We now consider how the behavior of suppliers of pirated content impacts λ. As mentioned earlier, our pirate suppliers are anyone who make use of cyberlockers or other file-sharing sites to distribute illegal content and are paid based on the number of downloads according to a menu similar to Table 2. We assume that there is a large number of identical potential pirate suppliers. However,

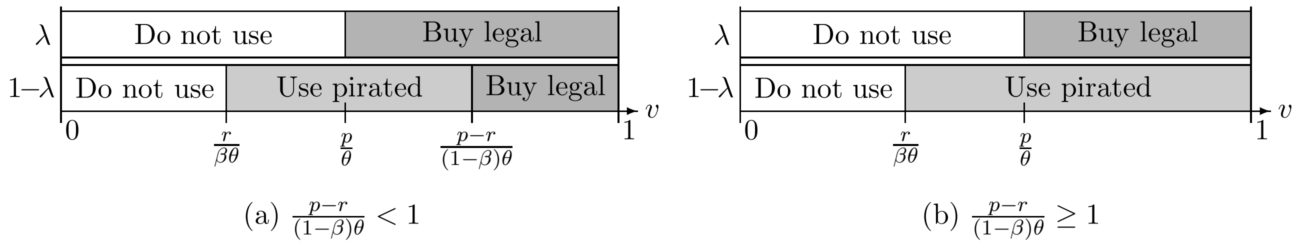

Figure 3: Consumers Self-Select Based on Their Relative Benefit and Search Result

not all of them will end up supplying an illegal copy. For, similar to Tunca and Wu (2013), we also assume that each pirate supplier faces an "entry cost," denoted e; this entry cost includes the *risk* of prosecution and penalty if convicted of distributing illegal copies, and, naturally, depends, in an aggregate sense, on all actions that amplify this risk by making it difficult to supply and profit from pirated content. Such actions may include, among other things, prosecuting the suppliers, shutting down cyberlockers, or requiring search engines to filter out links to illegal content. Therefore, e in our model essentially represents the level of *supply-side* enforcement, the extent to which the piracy ecosystem can be disrupted.

We use $\eta \in [0, 1]$ to denote the normalized supply of pirated content; this η represents the fraction of potential suppliers who end up distributing illegal copies. Because of the large base of potential suppliers, η also represents the probability that a single search by a consumer is able to locate the desired pirated content. When $\eta = 0$, there is no supply, whatsoever, of pirated copies. On the other hand, when $\eta = 1$, pirated copies are abundant and are, therefore, readily available upon a single search. Note that the assumption of a large number of suppliers is essentially equivalent to saying that the suppliers are *atomistic*, that is, the decision of a single supplier does not impact η. Our atomistic suppliers are similar to atomistic traders in financial markets (Boot and Thakor 1997), where the decision of one individual cannot impact the performance of the entire market, although the decisions of many individuals collectively can make a difference.

Since the probability of not finding the pirated content in a single search is $(1-\eta)$, the probability that a consumer searches exactly i times and still fails to find a pirated copy is simply $g(1-g)^i(1-\eta)^i$. Then, the probability that a consumer cannot find any pirated version is:

$$\lambda = \sum_{i=0}^{\infty} g(1-g)^i(1-\eta)^i = \frac{g}{1-(1-g)(1-\eta)}. \tag{3}$$

The total advertisement revenue earned by all suppliers taken together should be proportional to the number of illegal downloaders; 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 simply $\frac{\hat{q}}{\eta}$, 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 enters the market as long as the revenue is more than the cost. Thus, in a subgame equilibrium, the supply of pirated content matches its demand:

$$\frac{\hat{q}}{\eta} - e = 0 \Rightarrow \eta = \frac{\hat{q}}{e}. \tag{4}$$

Substituting (3) and (4) into (2), we can endogenize η. As shown in Figure 3, one of the two possible scenarios emerge:

- If $\frac{p-r}{(1-\beta)\theta} < 1$, then the legal demand consists of those who are unable to find a pirated version as well some of those who are able to find one. The manufacturer's price and quality decisions must consider both these groups.

- If $\frac{p-r}{(1-\beta)\theta} \geq 1$, the legal demand arises solely from those who cannot find a pirated copy, as everyone who can find one prefers it to the legal one. The manufacturer can completely ignore the latter type in its decision.

Therefore, we get:

$$\eta = \begin{cases} \frac{p-\frac{r}{\beta}}{e\theta(1-\beta)} - \frac{g}{1-g}, & \text{if } \frac{p-r}{(1-\beta)\theta} < 1, \\ \frac{\theta-\frac{r}{\beta}}{e\theta} - \frac{g}{1-g}, & \text{otherwise,} \end{cases} \quad \text{and} \quad \lambda = \begin{cases} \frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)}, & \text{if } \frac{p-r}{(1-\beta)\theta} < 1, \\ \frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)}, & \text{otherwise.} \end{cases} \tag{5}$$

These, in turn, can be substituted into (1) to obtain the legal demand, q, as a function of the context parameters e, r, β, and g, and the decision variables p and θ. Finally, if (5) does not satisfy $0 \leq \eta \leq 1$, then η is at one of its extreme values, either 0 or 1; accordingly, λ is either 1 or g.

3.3 Manufacturer's Decision and Market Equilibrium

Our manufacturer, situated in a real-world context with enforcement levels e and r, faces a demand of q and piracy loss of $\hat{q}$, both of which depend on the enforcement levels. Since the subgame equilibrium already endogenizes q, $\hat{q}$, and η, the overall equilibrium is found by simply maximizing the manufacturer's profit. Following Jones and Mendelson (2011), we assume:

Assumption 3 *The manufacturer's marginal cost of producing an additional copy is zero, and its cost of developing a product of quality level* θ *is* $c(\theta) = \frac{c\theta^2}{2}$, *where* $c > 0$.

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 shifts 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. For instance, when both e and r are low, the manufacturer sets p and θ in a way that the legal product becomes fairly unattractive and the demand for the pirated version rises sharply. In such a case, all suppliers want to provide a pirated copy because they can all make a positive profit, which drives η to 1. On the other hand, when enforcement is strong, or the legal product is priced aggressively, it is not profitable for a supplier to enter the market, and $\eta = 0$. Because of such possibilities, we need to consider three distinct cases: (i) $\eta \in (0, 1)$, (ii) $\eta = 1$, and (iii) $\eta = 0$. Depending on the context parameters, the manufacturer finds it optimal to be in exactly one of these cases; please see Figure 4.

Case 1–Limited Supply ($0 < \eta < 1$): In this case, a pirated copy has limited availability, and a search attempt reveals it with probability η. As indicated in Figure 3, the manufacturer names a price such that $\frac{p-r}{(1-\beta)\theta} < 1$ and ends up selling to both types of consumers, those who have access to a pirated copy and those who do not (Case 1A). Alternatively, it can set p so high that $\frac{p-r}{(1-\beta)\theta} \geq 1$ (Case 1B), effectively shutting out, from the legal product, all consumers who have access to a pirated version.

Case 2–Ample Supply ($\eta = 1$): In this case, $\eta = 1$, which implies that just one search attempt is sufficient to locate the pirated product with certainty. Obviously, $\lambda = g$ here, as the "ethical" consumers are now only those who never search for an illegal copy. Similar to Case 1, we again face two possibilities: the manufacturer may serve both groups, those with piracy as an option and those without (Case 2A), or consider just the "ethical" group (Case 2B).

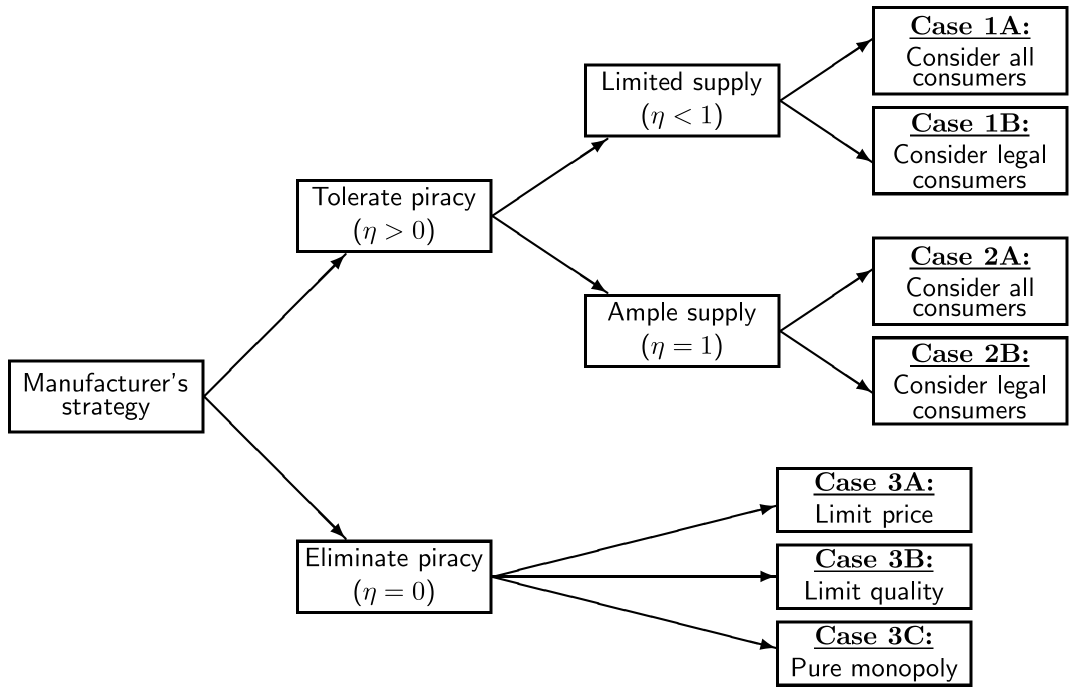

Figure 4: Manufacturer's Strategy Tree

Case 3–No Piracy ($\eta = 0$): Piracy ceases to exist when no consumer has the option to use a pirated version, or if the manufacturer chooses the price and/or quality in a way that the pirated product is rendered completely uncompetitive. In equilibrium, they are equivalent, and λ is one in both cases. There are three ways piracy may disappear: In Case 3A, the manufacturer chooses a "limit" price such that the illegal copy barely becomes unattractive compared to the legal one.[5] In Case 3B, the manufacturer chooses a "limit" quality to the same effect. In both cases, piracy ceases to exist, even though the threat of piracy still remains—unless the manufacturer holds the price or quality at the "limit" level, piracy can resurface. Finally, Case 3C happens when enforcement on either demand- or supply-side, or both, is high enough to suppress all threats completely, resulting in a pure monopoly.

[5]The concept of limit pricing here is the same as that in classical economics (Milgrom and Roberts 1982), where the limit price is used to discourage an entry of a potential competitor. Similarly, our monopolist tries to stave off the "shadow" competition from piracy by naming a price at the limit level.

4 Equilibrium Analysis

In this section, we analytically characterize the short- and long-run equilibrium for given values of e and r. Subsequently, to study the effects of changing these enforcement levels, we estimate the metrics of interest related to the level of piracy, the manufacturer's incentive to innovate, and the resulting consumer and social welfare.

4.1 Short-Run Equilibrium

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: $\max_p R = pq$, for a fixed, exogenous θ. Since the quality is fixed in the short run, Case 3B (limit quality) cannot be a part of the manufacturer's strategy space. For any β, g, and θ, the entire (r, e) space can be partitioned into several regions based on the equilibrium outcome, and their boundaries can be characterized, as shown in the appendix.

Figure 5 illustrates these partitions for $\beta = 0.75$, $g = 0.4$, and $\theta = 10$, with Region i representing the part of the parameter space where Case i occurs in equilibrium. This figure matches our

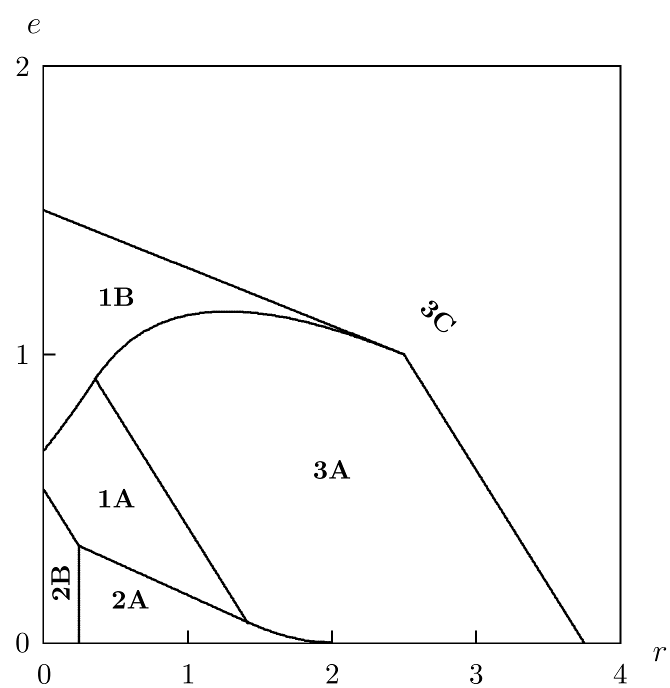

Figure 5: Short-Run: Relevant Partitions of the (r, e) Space ($\beta\!=\!0.75$, $g\!=\!0.4$, $\theta\!=\!10$)

intuition. At low values of e, we expect the supply to be abundant ($\eta = 1$), implying an equilibrium in Region 2A or 2B. As e increases, the supply of the pirated content becomes restricted ($0 < \eta < 1$), causing the equilibrium to move to either Region 1A or 1B. At an even higher e, there is no supply whatsoever ($\eta = 0$), and the equilibrium is found in Region 3C. Indeed, at very high values of e or r, both piracy and its threat should completely disappear, allowing the manufacturer to enjoy its full monopoly power. Prior research also identifies a similar possibility (Bae and Choi 2006, Lahiri and Dey 2013). However, unlike prior work, such a situation can surface in our equilibrium even when $r = 0$. This is because, in our model, a large e alone can stamp out all piracy completely. Of course, when e is low, we, too, find that achieving pure monopoly requires a very large r.

The occurrence of the equilibrium in Region 3A is somewhat more curious. For example, when $r = 1$, as e increases, the equilibrium moves from Region 1A to 3A, and then to 1B. The move from 1A to 3A is again intuitive—a higher e, resulting in a lower level of supply, allows the manufacturer to eliminate piracy by using the limit price. However, as e increases further, only a few consumers are able to locate a pirated copy, so the manufacturer finds it profitable to ignore them instead of luring them to the legal product through a low price. This moves the equilibrium from 3A to 1B.

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

$$p^*(\theta) = \begin{cases} \frac{\theta(1-\beta)(1-g(1-\beta e))}{2(1-g)} + \frac{r}{2}, & \textit{in Case 1A,} \\ \frac{\theta(1-\beta)+r(1-g)}{2(1-\beta g)}, & \textit{in Case 2A,} \\ \frac{eg\theta(1-\beta)}{1-g} + \frac{r}{\beta}, & \textit{in Case 3A,} \\ \frac{\theta}{2}, & \textit{in all other cases.} \end{cases} \tag{6}$$

It is interesting to observe that the equilibrium price in Proposition 1 depends on e only in Regions 1A and 3A; it is independent of e in all other regions. That the price does not depend on e in Region 1B is straightforward—here, the manufacturer completely ignores the fraction of consumers who have access to pirated content, and sets the price only for those who do not see piracy as an option. Similarly, in Region 3C, where the threat of piracy disappears, there is no need to consider the enforcement levels any more. Finally, in Regions 2A and 2B, where e is very small to have a material impact, the supply is abundant and the manufacturer behaves as if $e = 0$.

The connection with prior work is instructive. In Case 3A, the manufacturer sets the price at a limiting value that is low enough to eliminate piracy. This limit price of $\left(\frac{eg\theta(1-\beta)}{1-g} + \frac{r}{\beta}\right)$, however, is larger than that of $\frac{r}{\beta}$ in prior literature (Lahiri and Dey 2013). This is because our consumers do not enjoy an abundant supply of the pirated product, and the shrinkage in supply translates to a higher limit price. At the same time, though, our work is indeed consistent with prior work, since our limit price collapses to $\frac{r}{\beta}$ when $e = 0$. Another connection between our work and prior literature is through the construct of "ethical" consumers. The ethical segment in (August and Tunca 2008, Lahiri and Dey 2013), for example, actually coincides with ours when the supply of pirated content is abundant. Viewed differently, our ethical segment is repelled by pirated versions not just from ethical considerations but also out of simple economic calculations. We now estimate the metrics of interest.

Piracy Rate The *piracy rate*, μ, is defined as the number of pirated copies in use as a fraction of the total user base (BSA 2011): $\mu = \frac{\hat{q}}{q+\hat{q}}$. Substituting the equilibrium price from (6) into (1) and (2), we get:

$$\mu(\theta) = \begin{cases} \frac{r(1-g)(2-\beta)-\beta\theta(1-\beta)(1-g(1+e(2-\beta)))}{2(1-\beta)(r(1-g)-\beta\theta(1-g(1+e(1-\beta))))}, & \text{in Case 1A,} \\ \frac{2(\beta\theta-r)(r(1-g)-\beta\theta(1-g(1+e)))}{2r\beta\theta(2-g(2+e))-\beta^2\theta^2(2-g(2+e))-2r^2(1-g)}, & \text{in Case 1B,} \\ \frac{(1-g)(\beta\theta(1-\beta)-r(2-\beta(1+g)))}{(1-\beta)(\beta\theta(2-g(1+\beta))-r(1-g)(2-g\beta))}, & \text{in Case 2A,} \\ \frac{2(1-g)^2(\beta\theta-r)^2}{2r(1-g)^2(r-2\beta\theta)+\beta^2\theta^2(2-g(4-e-2g))}, & \text{in Case 2B,} \\ 0, & \text{in all other cases.} \end{cases} \tag{7}$$

Manufacturer's Profit The manufacturer's revenue—or profit since the marginal cost is zero—in equilibrium is $R^*(\theta) = p^*(\theta)q$. As before, substituting (6) into (1), we get:

$$R^*(\theta) = \begin{cases} \frac{(r(1-g)+\theta(1-\beta)(1-g(1-e\beta)))^2}{4\theta(1-\beta)(1-g)^2}, & \text{in Case 1A,} \\ \frac{eg\beta\theta^2}{4(1-g)(\beta\theta-r)}, & \text{in Case 1B,} \\ \frac{(r(1-g)+\theta(1-\beta))^2}{4\theta(1-\beta)(1-g\beta)}, & \text{in Case 2A,} \\ \frac{g\theta}{4}, & \text{in Case 2B,} \\ \frac{(\beta\theta(1-g(1+e(1-\beta)))-r(1-g))(r(1-g)+eg\beta\theta(1-\beta))}{\beta^2\theta(1-g)^2}, & \text{in Case 3A,} \\ \frac{\theta}{4}, & \text{in all other cases.} \end{cases} \tag{8}$$

Consumer Surplus We consider the surplus of only the legal consumers, which can be found from:

$$CS(\theta) = \lambda \int_{\frac{p}{\theta}}^{1} (v\theta - p)dv + (1-\lambda) \int_{\frac{p-r}{(1-\beta)\theta}}^{1} (v\theta - p)dv.$$

Substituting p from (6), and λ from (5), for different regions, we get:

$$CS(\theta) = \begin{cases} \frac{H_1+H_2}{8\theta(1-g)^2(1-\beta)^2}, & \text{in Case 1A,} \\ \frac{eg\beta\theta^2}{8(1-g)(\beta\theta-r)}, & \text{in Case 1B,} \\ \frac{g(\theta(1+\beta-2g\beta)-(1-g)r)^2+H_3}{8\theta(1-g\beta)^2}, & \text{in Case 2A,} \\ \frac{g\theta}{8}, & \text{in Case 2B,} \\ \frac{(\beta\theta(1-g(1+e(1-\beta)))-r(1-g))^2}{2\beta^2\theta(1-g)^2}, & \text{in Case 3A,} \\ \frac{\theta}{8}, & \text{in all other cases,} \end{cases} \tag{9}$$

where

$$\begin{aligned} H_1 &= (1-\beta)^2\theta^2 \left(eg\beta^2(2+g(e-2)) + 2\beta(1-g)(1-g(e+1)) + (1-g)^2\right), \\ H_2 &= 2r\theta(1-\beta)(1-g)\left(2\beta + eg\beta^2 - g\beta(e+2) - (1-g)\right) - r^2(3-2\beta)(1-g)^2, \text{ and} \\ H_3 &= \frac{(1-g)(\theta(1-\beta)(1+2\beta(1-g)) - r(3-g-2\beta))(r(1+g-2g\beta) + \theta(1-\beta)(1-2g\beta))}{(1-\beta)^2}. \end{aligned}$$

Social Welfare The total social welfare is estimated as: $SW(\theta) = R^*(\theta) + CS(\theta)$. It should be noted that this social welfare does not include the surplus generated by the illegal users. This is because we are of the opinion that an increase in welfare from an illegal activity should not influence a policymaker to embrace that activity. Consider this. The marginal value of money to the poor is more than that to the rich. However, an illegal activity such as stealing, for the sake of redistribution of wealth, cannot be supported, even though it may have some moral appeal. Indeed, Robinhoods do not have a place in the modern society (Lahiri and Dey 2013)!

4.2 Long-Run Equilibrium

We now consider the manufacturer's decision in the long run, where it has control over θ as well. We can write the manufacturer's long-run decision problem as: $\max_{\theta} \pi(\theta)$, where $\pi(\theta) = R^*(\theta) - \frac{c\theta^2}{2}$ and $R^*(\theta)$ is as given in (8). This problem can be solved for each of the regions discussed in Figure 4, in a manner analogous to the short-run one; see the appendix. Figure 6 illustrates the partitions of the (r, e) space for the same values of β and g as in Figure 5. It is not surprising that the two

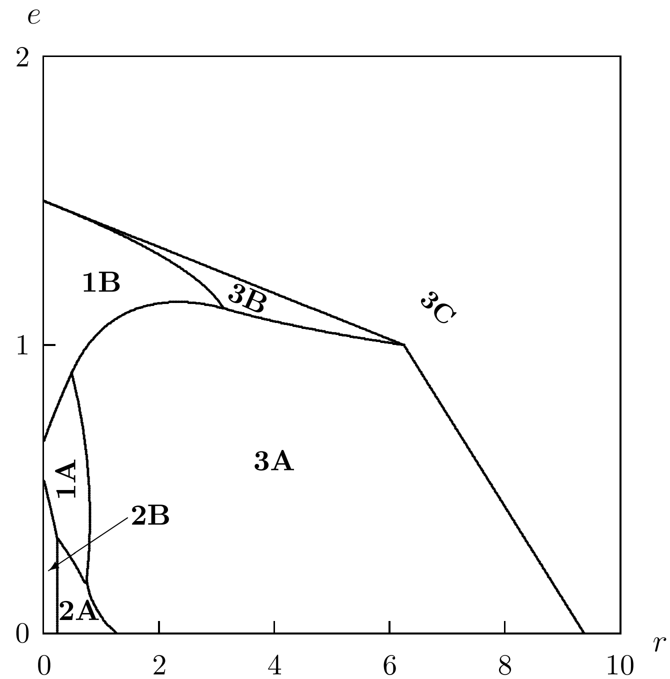

Figure 6: Long-Run: Relevant Partitions of the (r, e) Space ($\beta = 0.75$, $g = 0.4$, $c = 0.01$)

figures—Figures 5 and 6—are qualitatively similar. There are, however, two main differences. First, since the quality level is now a function of e and r, the boundaries in Figure 6 are different from the ones in Figure 5. This shifting of boundaries results in a substantial expansion of Region 3A in the long-run case. Second, since the manufacturer now has discretion over quality, Case 3B becomes a possibility, resulting in the emergence of Region 3B, where the manufacturer chooses quality as a tool to eliminate piracy, just as it uses price to do so in Region 3A. The limit quality that drives the pirated product out of market is $\frac{r(1-g)}{\beta(1-g(e+1))}$. Unlike the limit price, however, the idea that the manufacturer can use a limit quality as a means to squeeze out piracy is completely new.

Lemma 1 *The following equations represent the first order conditions for Cases 1A, 2A, and 3A,* *respectively:*

$$\frac{(1-\beta)(1-g(1-e\beta))^2}{4(1-g)^2} - \frac{r^2}{4\theta^2(1-\beta)} - c\theta = 0, \tag{FOC1A}$$

$$\frac{\theta^2(1-\beta)^2 - r^2(1-g)^2}{4\theta^2(1-\beta)(1-\beta g)} - c\theta = 0, \ \ and \tag{FOC2A}$$

$$\frac{eg(1-\beta)(1-g(1+e(1-\beta)))}{(1-g)^2} + \frac{r^2}{\beta^2\theta^2} - c\theta = 0. \tag{FOC3A}$$

Further, within the region of validity for each case, there exists a unique real positive root of the first order condition that maximizes its corresponding profit.

We denote these roots as $\theta_{1\mathrm{A}}$, $\theta_{2\mathrm{A}}$, and $\theta_{3\mathrm{A}}$ for (FOC1A), (FOC2A), and (FOC3A), respectively.

Proposition 2 (Quality) *In the long run, the equilibrium quality is given by:*

$$\theta^* = \begin{cases} \theta_{1\mathrm{A}}, & \textit{in Case 1A}, \\ \frac{eg\beta + \sqrt{eg\beta(eg\beta - 16cr(1-g))}}{8c\beta(1-g)} + \frac{r}{\beta} & \textit{in Case 1B}, \\ \theta_{2\mathrm{A}}, & \textit{in Case 2A}, \\ \frac{g}{4c}, & \textit{in Case 2B}, \\ \theta_{3\mathrm{A}}, & \textit{in Case 3A}, \\ \frac{r(1-g)}{\beta(1-g(e+1))}, & \textit{in Case 3B}, \\ \frac{1}{4c}, & \textit{in Case 3C}, \end{cases} \tag{10}$$

where $\theta_{1\mathrm{A}}$, $\theta_{2\mathrm{A}}$, *and* $\theta_{3\mathrm{A}}$ *are as defined above.*

The result in Proposition 2 can be better visualized in Figure 7, which shows the manufacturer's quality decisions in equilibrium over the entire (r, e) space for $\beta = 0.75$, $g = 0.4$, $c = 0.01$. As can be seen from this figure, the equilibrium quality attains its highest value of $\frac{1}{4c}$ in Region 3C, where the manufacturer enjoys complete monopoly power and the threat of piracy does not exist. In all other regions, quality is below that level and is also a function of both e and r. The implication is clear. When operating in a market fraught with piracy or its threat, the manufacturer must take into consideration the enforcement efforts from both sides, supply and demand. Curiously, in Regions 1A and 1B, where piracy is present in equilibrium, θ^* seems to decrease in r but increase in e, suggesting that the manufacturer's response to e and r are indeed very different.

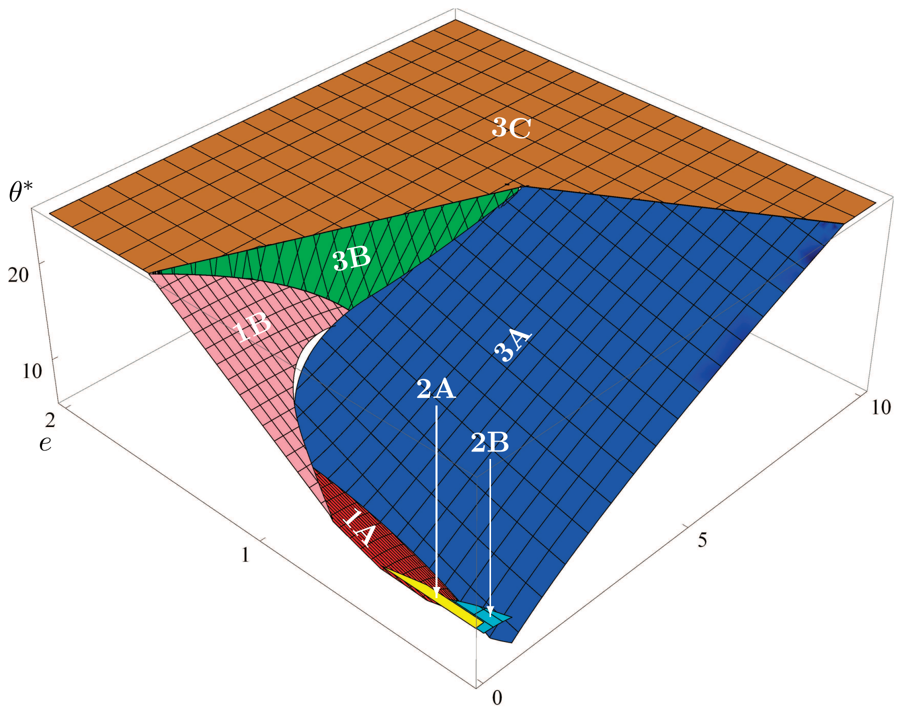

Figure 7: Equilibrium Quality as a Function of e and r ($\beta = 0.75$, $g = 0.4$, $c = 0.01$)

The equilibrium price (in the long run) can also be obtained by substituting θ^* from (10), on a case-by-case basis, into (6). It can be verified that the equilibrium price shows a trend that is quite similar to that of the equilibrium quality. Viewed differently, the manufacturer decides to invest in quality only when it can recover the additional investment through a higher price. In order to see this relationship more clearly, we plot the equilibrium price-to-quality ratio in Figure 8; this ratio is essentially a metric for the relative competitiveness of the legal product against its pirated version—the higher the ratio, the lower is the competitiveness. There are several interesting observations that can be made from this figure. First, in equilibrium, the price-to-quality ratio is a constant in Regions 1B, 2B, 3B, and 3C. Of course, in Regions 1B, 2B, and 3C, the manufacturer cares little about piracy, and the ratio—the relative competitiveness—is not impacted by either demand- or supply-side enforcement. What is surprising is that it is also the same in Region 3B, where the threat of piracy is neither absent nor has the manufacturer decided to ignore this threat in its quality decision.

The behavior of the price-to-quality ratio in the other three regions (1A, 2A, and 3A) is also curious. In these regions, the ratio increases with r, but remains quite flat with changes in e. As r increases, pirated content becomes less attractive, easing the competitive pressure on the

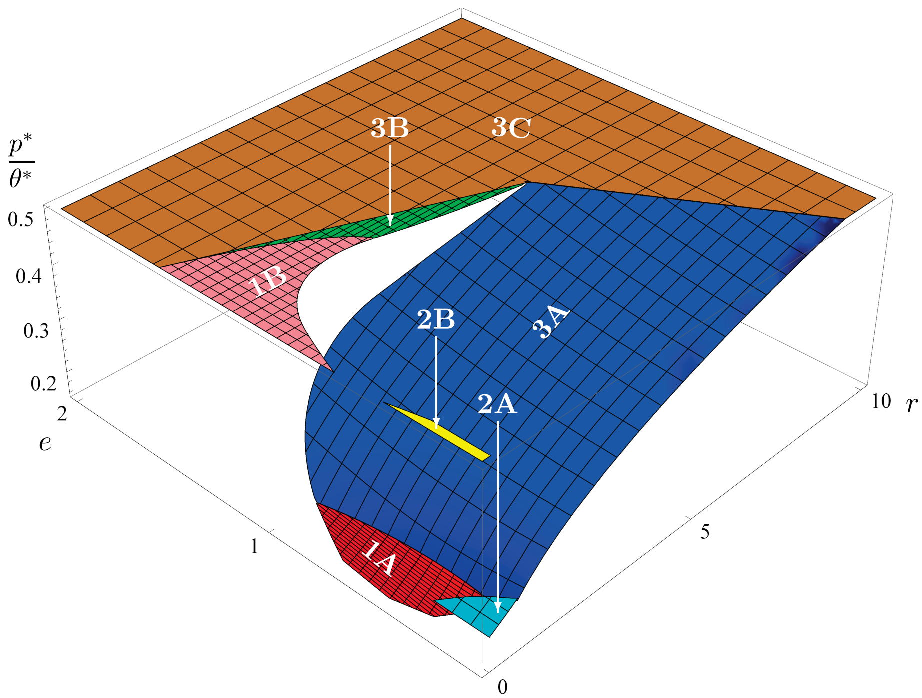

Figure 8: Equilibrium Price-to-Quality Ratio as a Function of e and r ($\beta = 0.75$, $g = 0.4$, $c = 0.01$)

manufacturer and allowing it to command a relatively higher price. On the other hand, e has little impact on the relative competitiveness, since the pirated copy remains as attractive as before when e changes. An increasing e only restricts the supply of pirated content, but the manufacturer has to be competitive for the fraction of the consumer base that still has access to the pirated version.

With the equilibrium quality and price known, it is conceptually easy to estimate the metrics of interest, now for the long-run equilibrium. First, the *piracy rate* can be estimated from (7), again by substituting θ^* from (10) for the appropriate case. Likewise, the *consumer surplus* can be obtained from (9) after suitable substitutions. Finally, we can use (8) to calculate the manufacturer's *profit*: $\pi^* = R^*(\theta^*) - \frac{c\theta^{*2}}{2}$.

A word on the veracity of our analysis is now in order. Given the level of complexity, it may be difficult for a reader to accept that our analysis is fully correct—there are no unintentional algebraic errors, mistranscriptions, or hidden assumptions. We must confess that we ourselves have been a bit skeptical at different points in this project. However, to allay those fears, as well as to gain valuable insights into the problem structure, we have also run hundreds of purely numerical experiments in Mathematica with a wide range of parameter values. We have found that all our numerical results fully mimic the respective analytical ones. This way, we have been able to ascertain that, even if errors or assumptions are still unfortunately hidden, their impact on the conclusions is minimal.

5 Comparative Statics

Now, we proceed to answer our research questions regarding the economic impacts of e and r. To that end, we conduct comparative statics with respect to e and r. Given the overwhelming presence of piracy around the globe (Bitton 2012)—as well as the fact that it is extremely difficult to eradicate piracy completely (BSA 2011)—we are naturally interested in only those cases where piracy exists: Cases 1A, 1B, 2A, and 2B. Of these four, 2A and 2B are not that interesting. As evident from Propositions 1 and 2, e in these regions is too small to have any effect on the equilibrium outcome. Together, they represent the case of ample supply of pirated content ($\eta = 1$), a case that has already attracted due attention in the literature (Bae and Choi 2006, Lahiri and Dey 2013). Hence, for the rest of the analysis, we limit our attention only to Regions 1A and 1B; together, we call them the *primary piracy region.*

Theorem 1 (Impact on Piracy Rate) *The piracy rate in the primary piracy region is monotonically decreasing in both e and r in the short run. In the long run, however, it is monotonically decreasing only in r; its relationship with e is ambiguous.*

The short-run result in Theorem 1 is quite intuitive. In the short run, both e and r have a direct impact on μ—as enforcement increases, irrespective of their type, fewer consumers use (or get to use) the pirated version, resulting in a lower piracy rate. However, the long-run case, where the manufacturer can change the quality level, is not that obvious, as it must also account for the indirect impact from this change in quality. It turns out that the indirect impact of r matches the trend of the direct one. On the other hand, the indirect impact due to e could have an exactly opposite tendency, often strong enough to overcome the direct impact altogether.

It is this last part of Theorem 1 that is most surprising. Intuitively, a higher level of an enforcement activity, regardless of its type, ought to lead to 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 e or r, or both; other popular search engines provide similar results as well. Furthermore, a mention of the piracy rate invariably shows up in lobbying efforts from manufacturers and their alliances to shore up support for higher levels of enforcement on both sides (Adobe 2013, BSA 2011). And, the 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 simply curbing the long-run piracy rate, Theorem 1 suggests that enforcement efforts are perhaps better directed at the demand side.

We now discuss the impact of enforcement on innovation. More specifically, we are interested in studying how the equilibrium quality, θ^*, changes with e and r. This is critical to understanding whether the two types of enforcement have the same impact on the manufacturer's long-run *incentive to innovate*, another potential metric for comparing the two approaches. Of course, it cannot be used as a metric for the short-run equilibrium whereby the quality is fixed.

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

Evidently, in the primary piracy region, the two enforcement approaches have exactly the opposite effect on the manufacturer's long-run quality decision. The equilibrium quality increases in e throughout the primary piracy region. On the other hand, in this region, it is decreasing in r. In Region 1A in particular, a higher r makes the pirated product less competitive, thereby reducing the manufacturer's incentive to differentiate the legal product in terms of quality. Since the quality difference between the legal and pirated products is $(1-\beta)\theta$, the manufacturer has a tendency to respond to a higher r with a lower θ. Now, in Region 1B, the pirated product, already weakened by a higher r, can be further impaired by a lower $\beta\theta$, which again provides the manufacturer with the strange incentive to cut θ.

Interestingly, though, e does not have any impact on the appeal of the pirated product. Increasing e directly limits the reach of the pirated product, which allows the manufacturer to easily recoup its investments in quality through higher prices. Hence, the manufacturer finds it profitable to increase θ in response to an increase in e. This increase in θ, however, could result in a higher piracy rate. An increasing θ, along with a steady price-to-quality ratio, could shift the boundary between legal and illegal users, $v = \frac{p-r}{(1-\beta)\theta}$, to the right (see Figure 3), resulting in a higher rate of piracy in the end (the last part of Theorem 1).

These contrasting impacts of e and r on quality provide important insights. First, a manufacturer must consider the two types of enforcement differently when forming its own strategy. Second, our finding highlights the fact that these two types of enforcement indeed result in very different incentives for innovation. More specifically, supply-side enforcement has a much more desirable effect on the incentive to innovate, and it may even lead to a better quality product. At the same time, demand-side enforcement has a markedly dissimilar—in fact, adverse—effect on the product quality. This is instructive. A major reason why policymakers are so concerned about digital piracy is that they believe piracy leads to reduced incentives for innovation (GAO 2010, OECD 2008). Theorem 2 suggests that policymakers interested in controlling piracy without sacrificing innovation should direct resources towards limiting the supply of pirated content instead of expending them on making piracy more costly and less attractive to consumers.

Proposition 3 (Impact on Profit) *Both e and r have a favorable impact on the manufacturer's profit—in the primary piracy region, the manufacturer's profit, both in the short and long run, is increasing in e as well as r.*

Proposition 3 shows that our model is consistent with real-world observations that an increase in any type of enforcement is indeed a good news for the manufacturer (Danaher et al. 2012, Danaher and Smith 2013). We now turn our attention to consumer surplus.

Theorem 3 (Impact on Consumer Surplus) *In the short run, the consumer surplus of legal users in the primary piracy region is not monotonic in either e or r; in the long run, it is increasing in e, but its relationship with r is ambiguous.*

Theorem 3 is indeed an important finding as it provides consumer advocates with a new perspective. Thus far, the battle against piracy has mainly taken place on the manufacturer's turf, with little support coming from consumer circles. In fact, the conventional belief is 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, some restrictions (e.g., DRM) designed to make procurement of pirated content more difficult can also reduce the utility of the legal product (Guarini 2013, Krebs 2005, Vernik et al. 2011). Naturally, consumers have often been against anti-piracy measures. Our result shows that such an attitude can be economically justified only in the short run, or only as far as demand-side enforcement is concerned. The story quickly changes, though, when we turn to supply-side enforcement—an increase in e not only results in a better quality product but can also improve consumer welfare in the long run.

We now turn our attention to social welfare, an issue of paramount importance; after all, it is the most natural yardstick for measuring effectiveness of a public policy.

Theorem 4 (Impact on Social Welfare) *Social welfare in the primary piracy region is increasing in both e and r in the short run; in the long run, it is increasing in e, but its relationship with r is ambiguous.*

The intuition is straightforward for the long-run case. Since the profit increases in both e and r, the social welfare can decrease only when the consumer surplus decreases sharply, offsetting the increase in the profit. This is exactly what happens when r increases in the primary piracy region. For example, in Region 1A, an increasing price-to-quality ratio has a negative impact on the consumer welfare, strong enough to nullify the higher profit, in the long run. The higher quality and the insignificant change in the price-to-quality ratio associated with a higher e, however, enhances both consumer and social welfare in the long run.

The short-run implications are somewhat more curious. From Theorem 3, we know that the short-run impact of enforcement on consumer surplus is ambiguous. Therefore, the short-run increase in social welfare in Theorem 4 must primarily be driven by an associated increase in the manufacturer's profit. Viewed differently, even though there is an increase in social welfare in the short run, its fruits are primarily enjoyed by the manufacturer.

6 Discussion

Our comparative statics have important practical implications for manufacturers, consumers, and the piracy ecosystem, along with broader connotations for public policy and law. In trying to answer the set of questions raised in the introduction, 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, both in the short and long run. To see this more clearly, we now summarize our results in Table 3.

Table 3: Impacts of Demand- and Supply-Side Enforcements

Enforcement ⟶ *Metric*	Demand-Side ($r \uparrow$)		Supply-Side ($e \uparrow$)	
	Short-Run	**Long-Run**	**Short-Run**	**Long-Run**
Piracy Rate	$\downarrow^*$	$\downarrow^*$	$\downarrow^*$	$\uparrow\downarrow$
Product Quality	—	$\downarrow^{**}$	—	$\uparrow^*$
Manufacturer's Profit	$\uparrow^*$	$\uparrow^*$	$\uparrow^*$	$\uparrow^*$
Consumer Surplus	$\uparrow\downarrow$	$\uparrow\downarrow$	$\uparrow\downarrow$	$\uparrow^*$
Social Welfare	$\uparrow^*$	$\uparrow\downarrow$	$\uparrow^*$	$\uparrow^*$
Overall Score	**3**	**1**	**3**	**4**

NOTE: *direction desirable; **direction undesirable

Each cell in this table shows the direction of the impact of an anti-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 in Table 3. These can be further consolidated into *overall scores*, counted as follows: Each cell with a desirable impact contributes a +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 score, one may be tempted to conclude that, even though demand-side enforcement may appear equally effective in the short run, supply-side enforcement is a far superior alternative in the long run (4 to 1). Since public policies should be made with long-run consequences in mind, one may then decide to argue in favor of supply-side enforcement, whether it is in the form of new legislation or implementation of the existing laws. And, such a conclusion may indeed be the right one, but the arguments ought to be somewhat more nuanced!

First and foremost, consider the plight of the piracy rate as a social metric. Accepting the above conclusion—supply-side enforcement is better in the long run—would immediately cast doubts about the piracy rate as a metric for formulating public policies on online piracy. Although touted as a measure of significance and widely viewed as such, the piracy rate could actually be a rather myopic metric, pointing towards the wrong policy direction in the long run. Its use for other purposes may be fine, but an increasing piracy rate need not be that alarming altogether, at least from a policy perspective. This is an important point because, often, "you get what you measure!" 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, an important aspect of supply-side enforcement is the positive impact it leaves on innovation. Often times, policy debates center around the issue of innovation. Indeed, the notion of innovation plays a crucial role in the development and design of a large majority of information goods (Brynjolfsson and Zhang 2007). We find that supply-side enforcement provides, to manufacturers of such products, 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. After all, any enforcement effort, supply- or demand-side, essentially makes it more difficult to pirate, one way or the other. Why then should they impact the equilibrium quality differently?

The answer lies in how different enforcement types reveal themselves in the market and the piracy ecosystem. 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 terms of a lower quality, at either the same or 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. Because the manufacturer must remain competitive for this consumer segment, it cannot afford to drastically reduce the quality or jack up the price-to-quality ratio. In fact, the presence of a bigger "ethical" 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.

Third, manufacturers always cheer an increased level of enforcement activities, whether they are supply- or demand-side. The increasing profit (in Table 3) with either type of enforcement, both in the short and long run, must surely provide ample justification, for thinking this way, to any manufacturer struggling with lost sales due to piracy. Or, perhaps not! A manufacturer completely oblivious of the surplus of its consumer base is likely to see a higher attrition rate, and a lower profit, in the long run. Indeed, examples abound where manufacturers have been soundly criticized in consumer circles for their overzealous anti-piracy actions (e.g., BBC 2008, Guarini 2013, Krebs 2005). And, such uproar has often resulted in the manufacturer scaling back its efforts. A 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).

Fourth, our results clearly show that, as far as piracy is concerned, manufacturers and consumers can finally agree on something—supply-side enforcement presents a "win-win" situation for both. Although there may be some ambivalence, or even some resistance, from consumer advocates in the short run, manufacturers would be better off backing supply-side efforts and backing off from the demand-side ones. At the same time, though, consumer groups should try to understand that not all enforcement activities have the same economic impact. If properly designed and executed, a supply-side action can, in fact, enhance overall consumer welfare. Opposition to anti-piracy measures should be moderated!

Fifth, often, the most important metric for policy debates is social welfare. In the short run, both types of enforcement seem to do well in this regard. The actual comparison, however, is a bit more subtle. Here, an increase in social welfare is primarily driven by a larger profit. In other words, it is possible to argue that, in the short run, most of the surplus generated through enforcement activities of either type would mainly go to the manufacturer, with very little coming the consumers' way. This imbalance, sometimes coupled with additional inconveniences faced by consumers, perhaps explains why there have been so many instances of consumer uproar against enforcement legislation and activities (Bachman 2011, BBC 2008, Guarini 2013, Krebs 2005).

It turns out that the long-run impact of demand-side enforcement is quite dubious. In contrast, stepping up enforcement on the supply side actually contributes to a higher social welfare in the long run. The key here is, again, innovation! It is the increasing quality level that translates to a higher consumer surplus, eventually leading to an increase in social welfare. Absent such innovation and investment in quality, as is the case in the short run, a large part of the social surplus generated through a higher enforcement level is extracted by the manufacturer, sometimes leading to a loss in consumer welfare. All in all, it seems reasonable to feel that supply-side enforcement has a "longer arm"—a much more desirable economic impact in the long run—even though its short-run implications could sometimes be a little bitter for consumers to swallow.

Finally, a word of caution! Our findings should not be automatically extrapolated to justify all supply-side legislation efforts, such as the SOPA, the PIPA, the PIRATE Act, or any other specific piece of legislation. When it comes to the letter of the law, the devil always ought to be in the detail. It is important to remember that, indeed, the original intent of the SOPA and PIPA was to limit the supply of pirated content. The massive campaign that followed against them, however, was not because they intended to limit the supply of pirated content; instead, that opposition was firmly rooted in the vagueness and unbounded scope of these bills (Fagin 2011, Heald 2012). Our work simply provides researchers and practitioners with a new lens—a clear distinction between the policies that diminish attractiveness of pirated content versus those that limit its supply—through which they can view and evaluate policy proposals.

7 Conclusion

In recent times, the policy debates about curbing online piracy has seen a major shift in their focus from the demand side 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 seeks to push the supply curve down to achieve the same goal of reducing piracy. It essentially involves combating piracy at its source, for example, by limiting the reach of pirate suppliers through shutting down their websites, filtering them out from search-engine results, and by penalizing 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 develop a parsimonious model capturing both sides of piracy. On the supply side of our model, there are online suppliers of pirated content who profit from advertisements. Their decision to provide illegal content depends on how their expected revenue compares with the cost imposed by the level of supply-side enforcement. On the demand side, when consumers use illegal content, they face a penalty, the expected value of which depends on the level of demand-side enforcement. Situated in this context, our manufacturer decides on the price and quality of its product in a way that maximizes its profit. The manufacturer's strategic decisions, coupled with the responses from consumers and pirate suppliers, determine the equilibrium outcome.

We find that, in fact, there are some fundamental differences between the economic impacts of these two types of enforcement. 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 consumer and social welfare as well, the effect of supply-side enforcement turns out to be quite desirable in the long run, although that of demand-side enforcement is not necessarily so.

It is true that we do not explicitly consider the costs incurred by policymakers in implementing these two types of enforcement. However, we feel that the inclusion of such costs would only bolster our case for supply-side enforcement. Often, supply-side enforcement is relatively inexpensive—shutting down one cyberlocker is likely much more economical than filing individual lawsuits against a multitude of illegal users. Individual lawsuits also lead to varying amounts of penalties depending on the specific instance of piracy, the legal jurisdiction, the jury handling the cases, and the contextual interpretation of ambiguous statutes in the law. For example, the set of 261 individual lawsuits filed by the RIAA in 2003 were mostly settled outside courtrooms for undisclosed amounts or dismissed altogether; the two conspicuous ones that went to trial resulted in vastly different amounts of penalty per song, \$80,000 versus \$9,250 (Lahiri and Dey 2013). In the end, the RIAA recovered only a tiny fraction of its legal costs but generated ample negative publicity; "the individual lawsuits were unbelievably counterproductive" (Holpuch 2012).

On the other hand, the closure of `MegaUpload.com`—one of the major cyberlockers accused of abetting online piracy—clearly illustrates how the beneficial effect of shutting down a major hosting site can swiftly propagate through the entire piracy ecosystem. When `MegaUpload.com` was taken down, "many linking websites, custom search engines, and custom streaming scripts that relied on the hosted content became inoperable" (Masnick 2012); as a direct consequence, "some websites were abandoned by their operators, others lost traffic, while still others shifted their business model." A recent empirical work by Danaher and Smith (2013) lends strong support to this assessment—the closure of `MegaUpload.com` resulted in a significant increase (6–10%) in digital movie sales over the following eighteen weeks. In light of such effects, some are calling the takedown of `MegaUpload.com` a "massive success," despite the site's relaunch a year later under a different domain name (Parfeni 2013). Admittedly, we do not model such issues, but incorporating them can only strengthen our argument, by tipping the balance further towards the "longer arm" of supply-side enforcement.

Our work is not without limitations; no work is! In reality, the ecosystem that supports online piracy is vastly more complex. This ecosystem often involves 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 policy implications. Despite such limitations, this work would have achieved its goal if it has succeeded in providing a new economic lens through which policy decisions can be viewed and evaluated.

References

Adobe. 2013. Whom does piracy hurt? URL `http://www.adobe.com/uk/aboutadobe/antipiracy/hurt.html`. Accessed on June 12, 2013.

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.

Bachman, K. 2011. House online piracy bill faces uproar from internet community: Lobbying against proposal got a head start. URL `http://www.adweek.com/news/technology/house-online-piracy-bill-faces-uproar-internet-community-136099`. Accessed on February 9, 2014.

Bae, S.-H., J.P. Choi. 2006. A model of piracy. *Information Economics and Policy* **18**(3) 303–320.

BBC. 2008. Copyright row dogs Spore release. URL `http://news.bbc.co.uk/2/hi/technology/7604405.stm`. Accessed on February 1, 2010.

BBC. 2012. US court reimposes $220,000 music piracy fine. URL `http://www.bbc.co.uk/news/technology-19572817`. Accessed on February 9, 2013.

Bhattacharjee, S., R.D. Gopal, K. Lertwachara, J.R. Marsden. 2006. Impact of legal threats on online music sharing activity: An analysis of music industry legal actions. *Journal of Law and Economics* **49**(1) 91–114.

Bitton, M. 2012. Rethinking the anti-counterfeiting trade agreement's criminal copyright enforcement measures. *The Journal of Criminal Law & Criminology* **102**(1) 67–117.

Boot, A.W.A., A.V. Thakor. 1997. Financial system architecture. *The Review of Financial Studies* **10**(3) 693–733.

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.

BSA. 2011. Eighth annual BSA global software 2010 piracy study. Business Software Alliance. URL `http://globalstudy.bsa.org/2011/downloads/study_pdf/2011_BSA_Piracy_Study-Standard.pdf`. Accessed on February 5, 2012.

Burgess, R. 2012. Hurt locker lawsuit targeting pirates comes to a halt. URL `http://www.techspot.com/news/46861-hurt-locker-lawsuit-targeting-pirates-comes-to-a-halt.html`. Accessed on May 21, 2012.

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.-N., 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.

Conner, K.R., R.O. Rumelt. 1991. Software piracy: An analysis of protection of strategies. *Management Science* **37**(2) 125–139.

Coutts, M. 2013. Is an all-out battle against internet piracy coming to Canada? URL `http://ca.news.yahoo.com/blogs/right-click/battle-against-internet-piracy-coming-canada-172625021.html`. Accessed on July 14, 2013.

Danaher, B., M.D. Smith. 2013. Gone in 60 seconds: The impact of the MegaUpload shutdown on movie sales. *Working Paper* URL `http://ssrn.com/abstract=2229349`.

Danaher, B., M.D. Smith, R. Telang, S. Chen. 2012. The effect of graduated response anti-piracy laws on music sales: Evidence from an event study in France. *Working Paper* URL `http://ssrn.com/abstract=1989240`.

Datoo, S. 2013. France drops controversial 'Hadopi law' after spending millions. URL `http://www.guardian.co.uk/technology/2013/jul/09/france-hadopi-law-anti-piracy`. Accessed on August 21, 2013.

Ducklin, P. 2013. First pirate prosecution in New Zealand under "three strikes" law. URL `http://nakedsecurity.sophos.com/2013/01/31/first-pirate-prosecution-in-new-zealand/`. Accessed on February 9, 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.

Fagin, B. 2011. SOPA, another flawed effort to make the Internet behave. URL `http://gazette.com/article/130569`. 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.

GAO. 2010. Intellectual property: Observations on efforts to quantify the economic effects of counterfeit and pirated goods. Government Accountability Office (Report# GAO-10-423).

Gardella, R., J. Forzato. 2011. US goes on offense against digital piracy. URL `http://investigations.nbcnews.com/_news/2011/05/28/6732624-us-goes-on-offense-against-digital-piracy`. Accessed on February 9, 2014.

Garrahan, M. 2013. Google urged to step up fight against film piracy. The Financial Times. URL `http://www.ft.com/intl/cms/s/0/d088e55e-46c8-11e3-bdd2-00144feabdc0.html#axzz2tbverbUh`. Accessed on Februray 17, 2014.

Gastaldo, E. 2012. Japan's new fine for illegal downloads: 2 years in jail. URL `http://www.newser.com/story/155016/japans-new-fine-for-illegal-downloads-2-years-in-jail.html`. Accessed on February 9, 2014.

Gelles, D. 2012. Debate over online piracy rages in us. The Financial Times. URL `http://www.ft.com/cms/s/0/6162237a-402b-11e1-82f6-00144feab49a.html`. Accessed on Februray 17, 2014.

Gopal, R.D., A. Gupta. 2010. Trading higher software piracy for higher profits: The case of phantom piracy. *Management Science* **56**(11) 1946–1962.

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.

Hall, A. 2013. Copyright infringement tips for illegal movie & music download cases. URL `http://thompsonhall.com/copyright-infringement-tips-for-illegal-movie-music-download-cases/`. Accessed on February 9, 2014.

Heald, C. 2012. Why SOPA is dangerous. URL http://mashable.com/2012/01/17/sopa-dangerous-opinion/. Accessed on May 8, 2013.

Holdsworth, N. 2013. New anti-piracy law allows Russian rights owners to tackle illegal internet download sites. URL http://www.hollywoodreporter.com/news/new-anti-piracy-law-allows-597962. Accessed on September 8, 2013.

Holpuch, A. 2012. Minnesota woman to pay $220,000 fine for 24 illegally downloaded songs. URL http://www.guardian.co.uk/technology/2012/sep/11/minnesota-woman-songs-illegally-downloaded. Accessed on May 8, 2013.

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.

Jackson, B. 2009. Small Canadian firms, school division fork out $270,000 in "piracy" damages. URL http://www.itbusiness.ca/news/small-canadian-firms-school-division-fork-out-270000-in-piracy-damages/13208. Accessed on August 8, 2010.

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.

Johnson, E.J., W.W. Moe, P.S. Fader, S. Bellman, G.L. Lohse. 2004. On the depth and dynamics of online search behavior. *Management Science* **50**(3) 299–308.

Jones, R., H. Mendelson. 2011. Information goods vs. industrial goods: Cost structure and competition. *Management Science* **57**(1) 164–176.

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.

Krebs, B. 2005. Study of sony anti-piracy software triggers uproar. The Washington Post. URL http://www.huffingtonpost.com/2013/06/19/xbox-one-always-online_n_3468098.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.

Lee, H. 1999. Illegal software crackdown begins. URL http://english.chosun.com/site/data/html_dir/1999/05/11/1999051161402.html. Accessed on March 11, 2010.

LeMay, R. 2013. Mass piracy lawsuits are back in Australia: Law firm targets end users' details. URL http://delimiter.com.au/2013/05/23/mass-piracy-lawsuits-are-back-in-australia-law-firm-targets-end-users-details/. Accessed on August 2, 2013.

Masnick, 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.

Mick, J. 2011. Hurt Locker lawsuit fragged in court. URL http://www.dailytech.com/Hurt+Locker+Lawsuit+Fragged+in+Court/article23608.htm. Accessed on May 8, 2013.

Milgrom, P., J. Roberts. 1982. Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica* **50**(2) 443–460.

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/. Accessed on May 8, 2013.

Moon, B. 2012. Entertainment industry to launch new anti-piracy campaign. URL http://www.marketplace.org/topics/tech/entertainment-industry-launch-new-anti-piracy-campaign. Accessed on May 8, 2013.

Moorthy, K. S. 1984. Market segmentation, self-selection, and product line design. *Marketing Science* **3**(4) 288–307.

Mussa, M., S. Rosen. 1978. Monopoly and product quality. *Journal of Economic Theory* **18**(2) 301–317.

OECD. 2008. The economic impact of counterfeiting and piracy. Organization for Economic Co-operation and Development.

Parfeni, L. 2013. The Pirate Bay again the biggest pirate site in the world, after MegaUpload's destruction. URL `http://news.softpedia.com/news/The-Pirate-Bay-Again-the-Biggest-Pirate-Site-in-the-World-After-MegaUpload-s-Destruction-342041.shtml`. Accessed on May 8, 2013.

PCWorld. 2011. So, you're being sued for piracy. URL `http://www.pcworld.com/article/230515/So_Youre_Being_Sued_for_Piracy.html`. Accessed on May 8, 2013.

Pepitone, J. 2011. 50,000 BitTorrent users sued for alleged illegal downloads. URL `http://money.cnn.com/2011/06/10/technology/bittorrent_lawsuits/index.htm`. Accessed on May 8, 2013.

Perry, N. 2012. Popular file-sharing website MegaUpload shut down. URL `http://usatoday30.usatoday.com/tech/news/story/2012-01-19/megaupload-feds-shutdown/52678528/1`. Accessed on May 8, 2013.

Reed, B. 2012. Google voluntarily plays copyright cop, punishes violators in search results. URL `http://news.yahoo.com/google-voluntarily-plays-copyright-cop-punishes-violators-search-204016430.html`. Accessed on May 8, 2013.

Seidler, E. 2011. Cyberlockers: Explaining piracy's profit pyramid. URL `http://popuppirates.com/?p=1249`. Accessed on May 8, 2013.

Seidler, E. 2013. Movie 2k blocked? URL `http://popuppirates.com/?p=2187`. Accessed on August 2, 2013.

Siu, E. 2012. 24 eye-popping SEO statistics. URL `http://www.searchenginejournal.com/24-eye-popping-seo-statistics/42665/`. Accessed on May 8, 2013.

Stahl, D. O. 1989. Oligopolistic pricing with sequential consumer search. *The American Economic Review* **79**(4) 700–712.

Stern, D. 2009. RIAA to stop suing for illegal downloading. URL `http://www.tuftsdaily.com/riaa-to-stop-suing-for-illegal-downloading`. Accessed on May 18, 2010.

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.

Varian, H. R. 1980. A model of sales. *The American Economic Review* **70**(4) 651–659.

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.

Woollacott, E. 2012. Google drops Pirate Bay from autocomplete results. URL `http://www.tgdaily.com/games-and-entertainment-features/66104-google-drops-pirate-bay-from-autocomplete-results`. Accessed on May 8, 2013.

Wu, S.-Y., P.-Y. Chen. 2008. Versioning and piracy control for digital information goods. *Operations Research* **56**(1) 157–172.

Appendix A: Proofs

Proof of Proposition 1

Case 3B is not possible here, as θ is exogenous in the short run. In each of the remaining cases (except Case 3A), we maximize the total revenue earned by the manufacturer, for a given θ, to find the its optimal price from the first order condition for that case. We have also verified that the second order condition is satisfied in each case. Case 3A is different. In this case, the manufacturer uses a limit price to weed out piracy; please see Figure 4. The boundaries separating all the different regions in Figure 5 are discussed in a separate technical appendix.

Case 1A: From the legal demand in (1), the revenue in this case is:

$$R = p\left(\lambda\left(1-\frac{p}{\theta}\right)+(1-\lambda)\left(1-\frac{p-r}{(1-\beta)\theta}\right)\right).$$

We substitute $\lambda = \frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)}$ into this to obtain $R = p+\frac{egp\beta}{1-g}-\frac{p(p-r)}{(1-\beta)\theta}$. The first order condition is then given by:

$$\frac{\partial R}{\partial p} = 1+\frac{eg\beta}{1-g}-\frac{2p-r}{(1-\beta)\theta} = 0,$$

which can be easily solved to obtain $p^*(\theta)$.

Case 2A: Since $\lambda = g$ in this case, from (1), the revenue for this case is given by:

$$R = p\left(g\left(1-\frac{p}{\theta}\right)+(1-g)\left(1-\frac{p-r}{(1-\beta)\theta}\right)\right) = p+\frac{p(r(1-g)-p(1-g\beta))}{(1-\beta)\theta}.$$

Solving $\frac{\partial R}{\partial p} = 1+\frac{r(1-g)-2p(1-g\beta)}{(1-\beta)\theta} = 0$, we get the desired result.

Case 3A: The limit price is obtained by equating $\lambda = \frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)}$ to 1.

Cases 1B, 2B, and 3C: In all the remaining cases, λ is independent of p: $\lambda = \frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)}$ in Case 1B, $\lambda = g$ in 2B, and $\lambda = 1$ in 3C. Therefore, the revenue in each case, $R = p\lambda\left(1-\frac{p}{\theta}\right)$, is clearly maximized at $p^*(\theta) = \frac{\theta}{2}$. ∎

Proof of Lemma 1

We consider each case separately.

Case 1A: From the legal demand in (1), the manufacturer's profit in this case, as a function of θ, is given by:

$$\pi = p^*(\theta)\left(\lambda\left(1-\frac{p^*(\theta)}{\theta}\right)+(1-\lambda)\left(1-\frac{p^*(\theta)-r}{(1-\beta)\theta}\right)\right)-\frac{c\theta^2}{2},$$

where $\lambda = \frac{eg\theta(1-\beta)}{(1-g)\left(p^*(\theta)-\frac{r}{\beta}\right)}$ and $p^*(\theta) = \frac{\theta(1-\beta)(1-g(1-\beta e))}{2(1-g)}+\frac{r}{2}$. Substituting these, we get:

$$\pi = \frac{(r(1-g)+(1-\beta)(1-g(1-e\beta))\theta)^2}{4\theta(1-g)^2(1-\beta)}-\frac{c\theta^2}{2}. \tag{A1}$$

The first order condition can then be derived as:

$$\frac{\partial\pi}{\partial\theta} = \frac{(1-\beta)(1-g(1-e\beta))^2}{4(1-g)^2}-\frac{r^2}{4\theta^2(1-\beta)}-c\theta = 0,$$

which is the same as (FOC1A).

We now show that (FOC1A) has a unique real positive root that maximizes the profit. First we observe that (FOC1A) is essentially a cubic equation, and the signs of its coefficients indicate that the product of the three roots is negative but their sum is positive. This immediately implies that there are exactly one negative and two positive roots. Of these two positive roots, one is a minimum, and the other a maximum, of the underlying profit. Since $\frac{\partial\pi}{\partial\theta}$ approaches $-\infty$ as θ becomes large, the higher of the two positive roots must be the maximum.

Case 2A: The profit for this case is given by:

$$\pi = p^*(\theta)\left(g\left(1-\frac{p^*(\theta)}{\theta}\right)+(1-g)\left(1-\frac{p^*(\theta)-r}{(1-\beta)\theta}\right)\right)-\frac{c\theta^2}{2},$$

where $p^*(\theta)=\frac{\theta(1-\beta)+r(1-g)}{2(1-\beta g)}$. Substituting this $p^*(\theta)$ and taking the derivative of π we get (FOC2A):

$$\frac{\partial\pi}{\partial\theta}=\frac{\theta^2(1-\beta)^2-r^2(1-g)^2}{4\theta^2(1-\beta)(1-\beta g)}-c\theta=0.$$

The rest of the proof is similar to the one for Case 1A.

Case 3A: The profit for this case is given by: $\pi=p^*(\theta)\left(1-\frac{p^*(\theta)}{\theta}\right)-\frac{c\theta^2}{2}$. Substituting $p^*(\theta)=\frac{eg\theta(1-\beta)}{1-g}+\frac{r}{\beta}$, we can derive the first order condition in (FOC3A):

$$\frac{\partial\pi}{\partial\theta}=\frac{eg(1-\beta)(1-g(1+e(1-\beta)))}{(1-g)^2}+\frac{r^2}{\beta^2\theta^2}-c\theta=0.$$

Since $\frac{\partial^2\pi}{\partial\theta^2}=-\frac{2r^2}{\beta^2\theta^3}-c<0$, we have a unique maximum in this case, as well. ■

Proof of Proposition 2

In each case (except Case 3B), we maximize the total profit, $\pi(\theta)=p^*(\theta)q-\frac{c\theta^2}{2}$, to find the optimal quality from the first order condition for that case. Verifying the second order condition is easy in Cases 1B, 2B, and 3C. Furthermore, Lemma 1 ensures the existence of a unique maximum for Cases 1A, 2A, and 3A, and we omit these cases in this proof. Of course, Case 3B is different in the long run problem. Here, the manufacturer uses a limit quality to eliminate piracy; please see Figure 4. The boundaries separating all the different regions in Figure 6 are similar to the short run case and are discussed in a separate technical appendix.

Cases 1A, 2A, and 3A: See Lemma 1.

Case 1B: In this case, the profit is: $\pi=p^*(\theta)\lambda\left(1-\frac{p^*(\theta)}{\theta}\right)-\frac{c\theta^2}{2}$, where $p^*(\theta)=\frac{\theta}{2}$ and $\lambda=\frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)}$. Substituting these, we get :

$$\pi=\frac{eg\beta\theta^2}{4(1-g)(\beta\theta-r)}-\frac{c\theta^2}{2}, \tag{A2}$$

which is concave in θ and can be maximized by solving for the first order condition:

$$\theta^*=\frac{eg\beta+\sqrt{eg\beta(eg\beta-16cr(1-g))}}{8c\beta(1-g)}+\frac{r}{\beta}.$$

Case 2B: Here, the profit is: $\pi=p^*(\theta)g\left(1-\frac{p^*(\theta)}{\theta}\right)-\frac{c\theta^2}{2}$, where $p^*(\theta)=\frac{\theta}{2}$. The resulting first order condition can be solved to get $\theta^*=\frac{g}{4c}$.

Case 3B: The limit quality is obtained by equating $\lambda=\frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)}$ to 1.

Case 3C: Here, the profit is: $\pi = p^*(\theta)\left(1 - \frac{p^*(\theta)}{\theta}\right) - \frac{c\theta^2}{2}$, where $p^*(\theta) = \frac{\theta}{2}$. The first order condition then results in $\theta^* = \frac{1}{4c}$. ■

Proof of Theorem 1

We consider the piracy rates of Cases 1A and 1B for both the short- and long-run situations.

Case 1A, Short Run: From (7), the piracy rate in this case is given by:

$$\mu = \frac{r(1-g)(2-\beta) - \beta\theta(1-\beta)(1-g(1+e(2-\beta)))}{2(1-\beta)(r(1-g) - \beta\theta(1-g(1+e(1-\beta))))}.$$

Taking partial derivative with respect to e and r, we get:

$$\begin{aligned}
\frac{\partial\mu}{\partial e} &= -\frac{g\beta^2\theta^2(1-g)}{2\left(\beta\theta(1-g(1+e(1-\beta))) - r(1-g)\right)^2} < 0, \text{ and} \\
\frac{\partial\mu}{\partial r} &= -\frac{\beta\theta(1-g)^2}{2(1-\beta)\left(\beta\theta(1-g(1+e(1-\beta))) - r(1-g)\right)^2} < 0.
\end{aligned}$$

Case 1A, Long Run: To see how μ changes with r in the long run, we use the chain rule:

$$\left.\frac{d\mu}{dr}\right|_{\theta=\theta^*} = \left.\frac{\partial\mu}{\partial r}\right|_{\theta=\theta^*} + \left.\frac{\partial\mu}{\partial\theta}\right|_{\theta=\theta^*}\frac{d\theta^*}{dr}.$$

We know from the short-run situation that $\left.\frac{\partial\mu}{\partial r}\right|_{\theta=\theta^*} < 0$. Furthermore, from Theorem 2, we know that $\frac{d\theta^*}{dr} < 0$. Finally,

$$\left.\frac{\partial\mu}{\partial\theta}\right|_{\theta=\theta^*} = \frac{r\beta(1-g)^2}{2(1-\beta)((1-g(1+e(1-\beta)))\beta\theta^* - (1-g)r)^2} > 0.$$

Combining everything, we get $\left.\frac{d\mu}{dr}\right|_{\theta=\theta^*} < 0$.

We now consider how μ changes with e in the long run; once again, we use the chain rule:

$$\left.\frac{d\mu}{de}\right|_{\theta=\theta^*} = \left.\frac{\partial\mu}{\partial e}\right|_{\theta=\theta^*} + \left.\frac{\partial\mu}{\partial\theta}\right|_{\theta=\theta^*}\frac{d\theta^*}{de}.$$

However, the first term on the right hand side is negative, whereas the second is not. This is because $\frac{d\theta^*}{de} > 0$ from Theorem 2. As a result, μ is not a monotonic function of e in the long run. To prove non-monotonicity, an example is sufficient: when $g = 0.1$, $\beta = 0.75$, $r = 0.5$, and $c = 0.01$, in the Region 1A, μ is increasing for $e \in [0.22, 0.31]$ but decreasing for $e \geq 0.32$.

Case 1B, Short Run: From (7), the piracy rate in this case is given by:

$$\mu = \frac{2(\beta\theta - r)(r(1-g) - \beta\theta(1-g(1+e)))}{2r\beta\theta(2-g(2+e)) - \beta^2\theta^2(2-g(2+e)) - 2r^2(1-g)}.$$

Taking partial derivative with respect to e and r, we get:

$$\begin{aligned}
\frac{\partial\mu}{\partial e} &= -\frac{2g\beta^2\theta^2(1-g)(\beta\theta - r)^2}{\left(2r^2(1-g) + \beta\theta(\beta\theta - 2r)(2-g(2+e))\right)^2} < 0, \text{ and} \\
\frac{\partial\mu}{\partial r} &= \frac{2eg\beta^2\theta^2(2r(1-g) - \beta\theta(2-g(2+e)))}{\left(2r^2(1-g) + \beta\theta(\beta\theta - 2r)(2-g(2+e))\right)^2}.
\end{aligned}$$

Now, to show that $\frac{\partial \mu}{\partial r}$ is also negative, we need to prove that $\beta\theta(2 - g(2+e)) > 2r(1-g)$, which follows directly from the fact that $\eta > 0$ in this region.

Case 1B, Long Run: This is quite similar to Case 1A in the long run. As before, using the chain rule, we can easily show that μ is again monotonic in r. However, as before, the directionality of μ with respect to e is ambiguous. ■

Proof of Theorem 2

Case 1A: Taking a partial derivative of (A1) with respect to θ, we get:

$$\pi' = \frac{\partial \pi}{\partial \theta} = \frac{(1-\beta)(1-g(1-e\beta))^2}{4(1-g)^2} - \frac{r^2}{4\theta^2(1-\beta)} - c\theta.$$

Clearly,

$$\frac{\partial \pi'}{\partial e} = \frac{g(1-\beta)\beta(1-g(1-e\beta))}{2(1-g)^2} > 0,$$

because $1 - g(1-e\beta) > 0$. On the other hand,

$$\frac{\partial \pi'}{\partial r} = -\frac{r}{2(1-\beta)\theta^2} < 0.$$

Finally, the second order condition must be satisfied in optimality, that is, $\left.\frac{\partial \pi'}{\partial \theta}\right|_{\theta=\theta^*} < 0$. It now follows from the Implicit Function Theorem that:

$$\frac{d\theta^*}{de} = -\frac{\left.\frac{\partial \pi'}{\partial e}\right|_{\theta=\theta^*}}{\left.\frac{\partial \pi'}{\partial \theta}\right|_{\theta=\theta^*}} > 0 \quad \text{and} \quad \frac{d\theta^*}{dr} = -\frac{\left.\frac{\partial \pi'}{\partial r}\right|_{\theta=\theta^*}}{\left.\frac{\partial \pi'}{\partial \theta}\right|_{\theta=\theta^*}} < 0.$$

Case 1B: Taking a partial derivative of (A2) with respect to θ, we get:

$$\pi' = \frac{eg\beta\theta(\beta\theta - 2r)}{4(1-g)(\beta\theta - r)^2} - c\theta,$$

differentiating which, we get:

$$\frac{\partial \pi'}{\partial e} = \frac{g\beta\theta(\beta\theta - 2r)}{4(1-g)(\beta\theta - r)^2} > 0 \quad \text{and} \quad \frac{\partial \pi'}{\partial r} = -\frac{egr\beta\theta}{2(1-g)(\beta\theta - r)^3} < 0.$$

The rest follows from the Implicit Function Theorem, in a manner similar to the proof for 1A above. ■

Proof of Proposition 3

Case 1A, Short Run: From (8), the manufacturer's profit in this case is:

$$R^* = \frac{(r(1-g) + \theta(1-\beta)(1-g(1-e\beta)))^2}{4\theta(1-\beta)(1-g)^2}.$$

Therefore, we can differentiate R^* with respect to e and r to get:

$$\frac{\partial R^*}{\partial e} = \frac{g\beta(r(1-g) + \theta(1-\beta)(1-g(1-e\beta)))}{2(1-g)^2} > 0 \quad \text{and} \quad \frac{\partial R^*}{\partial r} = \frac{(1-g) + eg\beta}{2(1-g)} + \frac{r}{2\theta(1-\beta)} > 0.$$

Case 1A, Long Run: Since $\pi(\theta) = R^*(\theta) - \frac{c\theta^2}{2}$, from the Envelope Theorem, we get:

$$\left.\frac{d\pi}{de}\right|_{\theta=\theta^*} = \left.\frac{\partial\pi}{\partial e}\right|_{\theta=\theta^*} = \left.\frac{\partial R^*}{\partial e}\right|_{\theta=\theta^*} > 0 \quad \text{and} \quad \left.\frac{d\pi}{dr}\right|_{\theta=\theta^*} = \left.\frac{\partial\pi}{\partial r}\right|_{\theta=\theta^*} = \left.\frac{\partial R^*}{\partial r}\right|_{\theta=\theta^*} > 0.$$

Case 1B, Short Run: From (8), the manufacturer's profit in this case is:

$$R^* = \frac{eg\beta\theta^2}{4(1-g)(\beta\theta - r)}.$$

Therefore, we can differentiate R^* with respect to e and r to get:

$$\frac{\partial R^*}{\partial e} = \frac{g\beta\theta^2}{4(1-g)(\beta\theta - r)} > 0 \quad \text{and} \quad \frac{\partial R^*}{\partial r} = \frac{eg\beta\theta^2}{4(1-g)(\beta\theta - r)^2} > 0.$$

Case 1B, Long Run: Once again, $\pi(\theta) = R^*(\theta) - \frac{c\theta^2}{2}$; and we can apply the Envelope Theorem to get the desired result. ■

Proof of Theorem 3

Case 1A, Short Run: From (9), the consumer surplus in this case is:

$$CS = \frac{H_1 + H_2}{8\theta(1-g)^2(1-\beta)^2}, \tag{A3}$$

where

$$\begin{aligned} H_1 &= (1-\beta)^2\theta^2\left(eg\beta^2(2+g(e-2)) + 2\beta(1-g)(1-g(e+1)) + (1-g)^2\right), \text{ and} \\ H_2 &= 2r\theta(1-\beta)(1-g)\left(2\beta + eg\beta^2 - g\beta(e+2) - (1-g)\right) - r^2(3-2\beta)(1-g)^2. \end{aligned}$$

Therefore, we can differentiate CS with respect to e and r to get:

$$\begin{aligned} \frac{\partial(CS)}{\partial e} &= -\frac{g\beta(r(1-g) + \theta(1-g-\beta-g\beta(e-1)))}{4(1-g)^2} \quad \text{and} \\ \frac{\partial(CS)}{\partial r} &= -\frac{r(1-g)(3-2\beta) + \theta(1-\beta)\left(1-g-2\beta+g\beta(e+2) - eg\beta^2\right)}{4\theta(1-g)(1-\beta)^2}. \end{aligned}$$

The sign of neither derivative can be ascertained one way or the other and depends on the parameter values. For example, let $g = 0.4$, $\beta = 0.75$, and $\theta = 10$. Then, the above derivatives are both negative when $(r, e) = (1, 0.25)$ but both positive at $(0.4, 0.75)$. Since both these points belong to Region 1A in the (r, e)-space (see Figure 5), clearly, the consumer surplus is not monotonic in either e or r.

Case 1A, Long Run: To see the long-run impact of e on the consumer surplus, we use the chain rule:

$$\left.\frac{d(CS)}{de}\right|_{\theta=\theta^*} = \left.\frac{\partial(CS)}{\partial e}\right|_{\theta=\theta^*} + \left.\frac{\partial(CS)}{\partial\theta}\right|_{\theta=\theta^*}\frac{d\theta^*}{de}.$$

Now, differentiating (A3) with respect to θ, we get:

$$\frac{\partial(CS)}{\partial\theta} = \frac{1}{8}\left(1 + \frac{\beta(2(1-g)(1-g(e+1)) + eg\beta(2+g(e-2)))}{(1-g)^2} + \frac{r^2(3-2\beta)}{\theta^2(1-\beta)^2}\right).$$

To find $\frac{d\theta^*}{de}$, we use the Implicit Function Theorem:

$$\frac{d\theta^*}{de} = -\frac{\left.\frac{\partial \pi'}{\partial e}\right|_{\theta=\theta^*}}{\left.\frac{\partial \pi'}{\partial \theta}\right|_{\theta=\theta^*}} = -\frac{g\beta(1-\beta)^2(1-g(1-e\beta))\theta^{*3}}{(1-g)^2\left(r^2 - 2c\theta^{*3}(1-\beta)\right)}.$$

Putting all the pieces together, we get:

$$\left.\frac{d(CS)}{de}\right|_{\theta=\theta^*} = \frac{g\beta}{4(1-g)^2}\left(\frac{(1-\beta)(1-g+eg\beta)X}{Z} - (1-g)(1-\beta)Y + eg\beta\theta^*\right),$$

where X, Y, and Z are given by:

$$X = \left(1 + \frac{\beta(2(1-g)(1-g-eg) + eg(2-(2-e)g)\beta)}{(1-g)^2} + \frac{r^2(3-2\beta)}{(1-\beta)^2\theta^{*2}}\right),$$

$$Y = \theta^*\left(\frac{r}{(1-\beta)\theta^*} + 1\right), \quad \text{and} \quad Z = 4c - \frac{2r^2}{(1-\beta)\theta^{*3}}.$$

Therefore,

$$\begin{aligned}
\left.\frac{d(CS)}{de}\right|_{\theta=\theta^*} &> \frac{g\beta}{4(1-g)^2}\left(\frac{(1-\beta)(1-g+eg\beta)X}{Z} - (1-g)(1-\beta)Y\right) \\
&> \frac{g\beta}{4(1-g)^2}\left(\frac{(1-\beta)(1-g)X}{Z} - (1-g)(1-\beta)Y\right) \\
&= \frac{g\beta(1-\beta)}{4(1-g)Z}\left(X - YZ\right).
\end{aligned}$$

Evidently, then, the sign of $\left.\frac{d(CS)}{de}\right|_{\theta=\theta^*}$ is the same as that of $(X-YZ)$. After substitution and rearrangement, we can show:

$$\begin{aligned}
X - YZ &= \left(1 + \frac{\beta(2(1-g)(1-g(e+1)) + eg\beta(2+g(e-2)))}{(1-g)^2} + \frac{2r^3}{\theta^{*3}(1-\beta)^2} - \frac{4c\theta^*}{1-\beta}\right) \\
&\quad + \left(\frac{2r^2}{\theta^{*2}(1-\beta)} + \frac{r^2(1+2(1-\beta))}{\theta^{*2}(1-\beta)^2} + \frac{4c(\beta\theta^* - r)}{1-\beta}\right).
\end{aligned}$$

Now, since $\frac{r}{\beta\theta} < 1$, the term on the second line of the above expression is clearly positive, implying:

$$X - YZ > 1 + \frac{\beta(2(1-g)(1-g-eg) + eg(2+(e-2)g)\beta)}{(1-g)^2} + \frac{2r^3}{\theta^{*3}(1-\beta)^2} - \frac{4c\theta^*}{1-\beta}.$$

Substituting $c\theta^*$ from (FOC1A) into the above expression leads to:

$$X - YZ > \frac{r^2(2r+\theta^*)}{\theta^{*3}(1-\beta)^2} + \frac{2\beta(1-g(1+e(2-\beta)))}{1-g}.$$

which would be positive as long as $(1-g(1+e(2-\beta)) > 0$. Fortunately, this last inequality is indeed true in this region—it follows from the condition $\eta > 0$.

We now show that the long-term impact of r on the consumer surplus is ambiguous. As before, we can use the chain rule to estimate:

$$\left.\frac{d(CS)}{dr}\right|_{\theta=\theta^*} = \left.\frac{\partial(CS)}{\partial r}\right|_{\theta=\theta^*} + \left.\frac{\partial(CS)}{\partial \theta}\right|_{\theta=\theta^*}\frac{d\theta^*}{dr}.$$

It turns out that this derivative does not have the same sign throughout Region 1A, and its sign depends on the parameter values. For example, let $g = 0.4$, $\beta = 0.75$, and $c = 0.01$. Then, the derivative is negative when $(r, e) = (0.6, 0.3)$ but positive at $(0.4, 0.8)$. It can be easily verified that both these points belong to Region 1A in the (r, e)-space (see Figure 6). Clearly, the consumer surplus is not monotonic in r.

Case 1B, Short Run: From (9), the consumer surplus is $CS = \frac{eg\beta\theta^2}{8(1-g)(\beta\theta-r)}$. It is quite easy to see that it is increasing in both e and r. Still, since there is no monotonicity in 1A, overall, the directions are ambiguous in the primary piracy region.

Case 1B, Long Run: Finding the long-run impact of e on the consumer surplus is significantly easier in this case. First, from the chain rule, we get:

$$\left.\frac{d(CS)}{de}\right|_{\theta=\theta^*} = \left.\frac{\partial(CS)}{\partial e}\right|_{\theta=\theta^*} + \left.\frac{\partial(CS)}{\partial\theta}\right|_{\theta=\theta^*}\frac{d\theta^*}{de}.$$

We already know from the short-run case that $\frac{\partial(CS)}{\partial e}$ is positive. And, $\frac{d\theta^*}{de}$ is positive, too, from Theorem 2. Finally, $\frac{\partial(CS)}{\partial\theta} = \frac{eg\beta\theta(\beta\theta-2r)}{8(1-g)(\beta\theta-r)^2} > 0$. Clearly, then, $\left.\frac{d(CS)}{de}\right|_{\theta=\theta^*} > 0$, implying a monotonic relationship between the consumer surplus and e.

Whether the consumer surplus in this region is monotonic in r or not, it is certainly not so in Region 1A (see above), implying that the relationship is indeed ambiguous in an overall sense. ■

Proof of Theorem 4

Case 1A, Short Run: Since social welfare is the sum total of the consumer surplus and the manufacturer's profit, it can be easily shown that:

$$\frac{\partial(SW)}{\partial e} = \frac{g\beta(r(1-g) + \theta(1-g)(1-\beta) + eg\beta\theta(3-2\beta))}{4(1-g)^2} \quad \text{and} \quad \frac{\partial(SW)}{\partial r} = \frac{eg\beta}{4(1-g)} + \frac{\theta(1-\beta)-r}{4\theta(1-\beta)^2}.$$

The first derivative is clearly positive. The second derivative would also be positive if $\theta(1-\beta) > r$, which is indeed true and follows from the fact that $\eta > 0$ in this region.

Case 1A, Long Run: Since the consumer surplus and the manufacturer's profit are both increasing in e, the long-run impact of e on social welfare is positive in this case. The marginal impact of r, however, does not have the same sign throughout Region 1A, as its sign depends on the parameter values. For example, let $g = 0.4$, $\beta = 0.75$, and $c = 0.01$. Then, the marginal impact is negative when $(r, e) = (0.6, 0.3)$ but positive at $(0.4, 0.8)$. Both these points belong to Region 1A in the (r, e)-space (see Figure 6). Therefore, social welfare is not monotonic in r.

Case 1B, Short Run: Since the consumer surplus and the manufacturer's profit are increasing in both e and r, the result follows.

Case 1B, Long Run: Since the consumer surplus and the manufacturer's profit are increasing in both e and r and, therefore, so is social welfare. ■

Appendix B: Characterization of Different Equilibrium Regions

In this appendix, we discuss both the short- and long-run equilibrium. For the equilibrium to occur within a specific region, it must first provide a solution that is valid in that region, that is, the manufacturer's decisions about p and θ must abide by the following restrictions derived from (5):

Region 1A: $g < \frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)} < 1$ and $\frac{p-r}{(1-\beta)\theta} < 1$.

Region 1B: $g < \frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)} < 1$ and $\frac{p-r}{(1-\beta)\theta} \geq 1$.

Region 2A: $\frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)} = g$ and $\frac{p-r}{(1-\beta)\theta} < 1$.

Region 2B: $\frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)} = g$ and $\frac{p-r}{(1-\beta)\theta} \geq 1$.

Region 3A: $\frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)} = 1$ and $\frac{p-r}{(1-\beta)\theta} < 1$.

Region 3B: $\frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)} = 1$ and $\frac{p-r}{(1-\beta)\theta} \geq 1$.

Region 3C: $\frac{eg\theta(1-\beta)}{(1-g)\left(p-\frac{r}{\beta}\right)} > 1$ or $\frac{eg\theta}{(1-g)\left(\theta-\frac{r}{\beta}\right)} > 1$.

In addition, for the equilibrium to occur in a region, the manufacturer's profit there must also be higher than its profit in all other regions that provide a valid solution.

B.1 Short-Run Equilibrium

In the short run θ is exogenous, and the manufacturer only chooses the price for a given θ. We consider each region separately.

Region 1A: Equating the profits for Regions 1A and 1B, we get:

$$e = h_1(r) = \frac{(1-g)(r+(1-\beta)\theta)^2}{g\beta\theta(1-\beta)(\beta\theta-r)}. \tag{B1}$$

We find that 1A dominates if $e < h_1(r)$, 1B dominates if $e > h(r)$, and both provide the same profit at $e = h_1(r)$.

Similarly, comparing the profits from 1A and 3A, we get the following boundary:

$$e = h_2(r) = \frac{1-g}{g}\left(\frac{1}{2-\beta} - \frac{r}{\beta(1-\beta)\theta}\right), \tag{B2}$$

and 1A dominates when $e < h_2(r)$.

Now, we compare the profit in 1A with that from 2A and 2B to obtain:

$$e = h_3(r) = \frac{(1-g)\left(\sqrt{(1-g\beta)(r(1-g)+\theta(1-\beta))^2} - (1-g\beta)(r+\theta(1-\beta))\right)}{g\beta\theta(1-\beta)(1-g\beta)}, \text{ and} \tag{B3}$$

$$e = h_4(r) = \frac{(1-g)\left(\theta\left(\sqrt{g(1-\beta)} - (1-\beta)\right) - r\right)}{g\beta\theta(1-\beta)}. \tag{B4}$$

1A dominates to the right of these boundaries, that is, when e is greater then both $h_3(r)$ and $h_4(r)$. Region 1A can now be fully characterized as follows:

$$\{(r,e)|\max\{h_3(r),h_4(r)\} \leq e < \min\{h_1(r),h_2(r)\}; e,r \geq 0\}. \tag{RGN1A}$$

Point to note here is that, in obtaining the boundaries of Region 1A, its profit need not be compared with the profit from either 3B or 3C. This is because, considering the validity conditions, we see that 1A can never have a valid solution when 3B or 3C does.

Region 1B: The result of profit comparison with 1A is already captured in (B1). Next, we compare its profit with that of 3A to obtain the following boundary:

$$\begin{aligned} e = h_5(r) \quad = \quad & \frac{1-g}{8g\beta\theta(1-\beta)^2(\beta\theta - r)}\Big(8r^2(1-\beta) - 12r\beta\theta(1-\beta) + \beta^2\theta^2(3-4\beta) + \\ & \beta\theta\sqrt{\beta\left(\beta\theta^2(3-4\beta)^2 - 16r^2(1-\beta) - 8r\theta(1-\beta)(1-4\beta)\right)}\Big), \end{aligned} \tag{B5}$$

such that Region 1B occurs only above this boundary. Combining the boundaries in (B1) and (B5), we get:

$$e = h_6(r) = \begin{cases} h_1(r), & \text{if } 0 \le r < \rho_1 = \frac{\theta(1-\beta)(3\beta-2)}{4-3\beta}, \\ h_5(r), & \text{otherwise,} \end{cases} \tag{B6}$$

where ρ_1 is the solution of $h_1(r) = h_5(r)$. Clearly, (B6) provides a combined lower boundary for 1B. To find the only other possible lower boundary for this region, we now compare its profit with that of Region 2B to obtain:

$$e = h_7(r) = \frac{(1-g)(\beta\theta - r)}{\beta\theta}. \tag{B7}$$

To find the upper boundary for Region 1B, we equate its profit with that from Region 3C to obtain:

$$e = \frac{(1-g)(\beta\theta - r)}{g\beta\theta} = \frac{h_7(r)}{g}. \tag{B8}$$

Region 1B can now be fully characterized as follows:

$$\left\{(r,e)|\max\{h_6(r), h_7(r)\} \le e < \tfrac{h_7(r)}{g}; r \le \rho_2 = \tfrac{\theta(2\beta-1)}{2}; e, r \ge 0\right\}, \tag{RGN1B}$$

where ρ_2 is the solution of $h_5(r) = \frac{h_7(r)}{g}$.

Region 2A: Comparing the profits from Regions 2A and 2B, we get:

$$r = \rho_3 = \frac{\theta\left(\sqrt{g(1-\beta)(1-g\beta)} - (1-\beta)\right)}{1-g}. \tag{B9}$$

Region 2A occurs to the right of this boundary and 2B, to the left. The boundary with Region 1A is already derived in (B3). We now find the boundary between Regions 2A and 3A by comparing their profits:

$$\begin{aligned} e = h_8(r) = & \frac{1-g}{2g\beta\theta(1-\beta)^2(1-g\beta)}\Big((1-\beta)(1-g\beta)(\beta\theta - 2x) - \\ & \beta\sqrt{(1-g)(1-\beta)(1-g\beta)\left(\beta\theta^2(1-\beta) - x^2(1-g) - 2x\theta(1-\beta)\right)}\Big). \end{aligned} \tag{B10}$$

Then, the combined upper boundary can be found as:

$$e = h_9(r) = \begin{cases} h_3(r), & \text{if } 0 \le r < \rho_4 = \frac{\theta(1-\beta)\left(2\sqrt{(1-g\beta)} - (2-\beta)\right)}{(1-g)(2-\beta)}, \\ h_8(r), & \text{if } \rho_4 \le r < \rho_5 = \frac{\beta\theta(1-\beta)}{2-\beta(1+g)}, \\ 0, & \text{otherwise,} \end{cases} \tag{B11}$$

where ρ_4 and ρ_5 are solutions of $h_3(r) = h_8(r)$ and $h_8(r) = 0$, respectively. Region 2A can now be expressed as:

$$\{(r,e)|e < h_9(r); \rho_3 \leq r < \rho_5; e, r \geq 0\}, \tag{RGN2A}$$

Region 2B: The possible boundaries with Regions 1A, 1B and 2A have already been found in (B4), (B7), and (B9), respectively. Of these, (B4) is a valid boundary for Region 2B only if Region 1A occurs for the given set of parameter values. Hence, we modify it as:

$$e = h_{10}(r) = \begin{cases} h_4(r), & \text{if } h_4(r) < h_2(r), \\ \infty, & \text{otherwise,} \end{cases} \tag{B12}$$

We now find the possible boundary with Region 3A by comparing the profits:

$$e = h_{11}(r) = \frac{(1-g)\left(\beta\theta\left(1-\sqrt{1-g}\right)-2r\right)}{2g\beta\theta(1-\beta)}. \tag{B13}$$

The complete characterization of Region 2B is, therefore, given by:

$$\{(r,e)|e < \min\{h_7(r), h_{10}(r), h_{11}(r)\}; r < \rho_3; e, r \geq 0\}, \tag{RGN2B}$$

Region 3A: The only remaining profit comparison is between Regions 3A and 3C. We do so now to obtain the boundary between them:

$$e = h_{12}(r) = \frac{(1-g)(\beta\theta - 2r)}{2g\beta\theta(1-\beta)}. \tag{B14}$$

It is now possible to express this region as:

$$\{(r,e)|\max\{h_2(r), h_9(r), h_{11}(r)\} \leq e < \min\{h_5(r), h_{12}(r)\}; e, r \geq 0\}, \tag{RGN3A}$$

Region 3B: As mentioned in the paper, Region 3B cannot occur in the short-run equilibrium.

Region 3C: When enforcement is very high, on either side, the threat of piracy disappears completely, and we enter this region of pure monopoly. Therefore, there is no upper boundary for this region. It only has a lower boundary, shared with Regions 1B and 3A. Both these boundaries have been found above in (B8) and (B14). Therefore, we can characterize this region as:

$$\left\{(r,e)|e \geq \min\left\{\tfrac{h_7(r)}{g}, h_{12}(r)\right\}; e, r \geq 0\right\}, \tag{RGN3C}$$

Finally, we note that, even though the equilibrium occurs in exactly one of the six possible regions (see Figure 5), not all regions may exist always. Depending on the parameter values, certain regions may actually disappear, implying that the manufacturer would never choose the associated strategy (see Figure 4). For example, one can show that, if $g \leq \frac{1-\beta}{\beta}$, Region 2B can never occur.

B.2 Long-Run Equilibrium

Finding the boundaries of different regions for the long-run equilibrium is conceptually a straightforward extension of the short-run exercise. Here, when comparing the profits of two regions, we must replace the exogenous θ in the short run by its long-run equilibrium value of θ^*. As shown in Proposition 2, θ^* is analytically characterizable, even though its expression is rather cumbersome. We leave out the tedious details, and only provide an example of how a region can be characterized. We pick Region 3B for this illustration.

Region 3B: Region 3B can occur only in the long-run equilibrium. This region has boundaries with Region 3C above and Regions 1B and 3A below. From Propositions 1 and 2, the equilibrium quality, θ^*, and price, $p^*(\theta^*)$, in each case can be found:

$$(\theta^*, p^*(\theta^*)) = \begin{cases} \left(\frac{eg\beta+\sqrt{eg\beta(eg\beta-16cr(1-g))}}{8c\beta(1-g)} + \frac{r}{\beta}, \frac{eg\beta+\sqrt{eg\beta(eg\beta-16cr(1-g))}}{16c\beta(1-g)} + \frac{r}{2\beta}\right) & \text{in Case 1B,} \\ \left(\theta_{3\mathrm{A}}, \frac{\theta_{3\mathrm{A}}}{2}\right), & \text{in Case 3A,} \\ \left(\frac{r(1-g)}{\beta(1-g(e+1))}, \frac{r(1-g)}{2\beta(1-g(e+1))}\right), & \text{in Case 3B,} \\ \left(\frac{1}{4c}, \frac{1}{8c}\right), & \text{in Case 3C,} \end{cases}$$

where $\theta_{3\mathrm{A}}$ is the unique real solution of (FOC3A) that maximizes the profit in Case 3A. It is now possible to calculate the equilibrium profit in each case as:

$$\pi^* = \begin{cases} \pi_{1\mathrm{B}} = \frac{\left(3eg\beta-\sqrt{eg\beta(eg\beta-16cr(1-g))}\right)\left(8cr(1-g)+eg\beta+\sqrt{eg\beta(eg\beta-16cr(1-g))}\right)^2}{128c\beta^2(1-g)^2\left(eg\beta+\sqrt{eg\beta(eg\beta-16cr(1-g))}\right)}, & \text{in Case 1B,} \\ \pi_{3\mathrm{A}} = \frac{(\beta\theta_{3\mathrm{A}}(1-g(1+e(1-\beta)))-r(1-g))(r(1-g)+eg\beta\theta_{3\mathrm{A}}(1-\beta))}{\beta^2\theta_{3\mathrm{A}}(1-g)^2} - \frac{c\theta_{3\mathrm{A}}^2}{2}, & \text{in Case 3A,} \\ \pi_{3\mathrm{B}} = \frac{r(1-g)(\beta(1-g(e+1))-2cr(1-g))}{4\beta^2(1-g(e+1))^2}, & \text{in Case 3B,} \\ \pi_{3\mathrm{C}} = \frac{1}{32c}, & \text{in Case 3C,} \end{cases}$$

Now, let $H_1(r)$, $H_2(r)$, and $H_3(r)$ be functions of r, such that $e = H_1(r)$ is equivalent to $\pi_{3\mathrm{B}} = \pi_{1\mathrm{B}}$, $e = H_2(r)$ to $\pi_{3\mathrm{B}} = \pi_{3\mathrm{A}}$, and $e = H_3(r)$ to $\pi_{3\mathrm{B}} = \pi_{3\mathrm{C}}$. Then, Region 3B can be characterized as:

$$\{(r, e)|\max\{H_1(r), H_2(r)\} \le e < H_3(r); e, r \ge 0\}.$$

Although, these boundaries are not always closed form, for a few of them, closed-form expressions exist. For example, after some algebraic manipulations, $\pi_{3\mathrm{B}} = \pi_{3\mathrm{C}}$ can be simplified to:

$$e = H_3(r) = \frac{(1-g)(\beta - 4cr)}{g\beta}.$$