NBER WORKING PAPER SERIES

PIRACY AND COPYRIGHT ENFORCEMENT MECHANISMS

Brett Danaher
Michael D. Smith
Rahul Telang

Working Paper 19150
http://www.nber.org/papers/w19150

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2013

Piracy and Copyright Enforcement Mechanisms
Brett Danaher, Michael D. Smith, and Rahul Telang
NBER Working Paper No. 19150
June 2013
JEL No. D69,L1,L11,L8,L82,M31,O30

ABSTRACT

Much debate exists around the impact that illegal file sharing may have on the creative industries. Similarly, opinions differ regarding whether the producers of artistic works should be forced to accept any weakening of intellectual property rights resulting from illegal file sharing, or if governments should intervene to protect these rights. This chapter seeks to inform these questions by outlining what we do and do not know from existing academic research.

We first discuss whether filesharing displaces sales of media goods and then discuss whether such displacement will lead to reduced incentives to produce new creative works. We continue by summarizing recent findings on what businesses can do to compete with piracy and the effectiveness of anti-piracy interventions on encouraging consumers to migrate from illegal to legal consumption channels. We conclude by demonstrating that without additional empirical evidence, it will be difficult to determine the socially optimal set of strategies and government copyright policies in the digital era.

Brett Danaher
Wellesley College
bdanaher@wellesley.edu

Rahul Telang
Carnegie Mellon University
rtelang@andrew.cmu.edu

Michael D. Smith
Carnegie Mellon University
mds@cmu.edu

1. Introduction

The goal of this chapter is to analyze what the academic literature says about piracy's impact on the media industries and on society, and to analyze the effectiveness of various industry and government efforts to respond to the threat of piracy. Our discussion proceeds in two main sections. In the first section we analyze what the academic literature says about the impact of piracy on sales and on creative incentives. In the second section we review the academic literature regarding the effectiveness of various industry and government initiatives to address piracy.

We draw several conclusions in our analysis. First, our analysis of the literature suggests that piracy represents a significant threat to media sales: The vast majority of studies in the academic literature find that piracy results in statistically and economically significant harm to media sales. Second, we conclude that, while the academic literature is only beginning to analyze the broader social implications of piracy, a strong economic argument can be made that reduced sales from piracy will, *ceteris paribus*, reduce incentives for the development of new creative works, and that reduced incentives are likely to reduce overall social welfare. Finally, we conclude that there are a variety of measures that the creative industries and governments can use to reduce the threat from piracy.

2. Piracy and Social Welfare

One can't analyze *how* governments and industries should respond to piracy without first analyzing *whether* there is a need to respond. In this section our goal is to review the academic literature on two important questions:

1. Does piracy harm media sales?

2. Does piracy impact the incentives to create media content?

Answers to these questions are critical to policy-makers and industry decision-makers because if piracy does not reduce media sales it could easily be viewed as a socially

beneficial activity that provides access to consumers who otherwise wouldn't have paid for content without cannibalizing sales from consumers who are paying for content.

If piracy does harm sales of media content, the next logical question is whether this reduction in revenue reduces the incentives to creators to produce content. In this regard, even if piracy reduces legitimate sales, if it does not harm the incentives of creators one could potentially view piracy simply as a wealth transfer from content creators to content consumers with no net reduction in overall social welfare, and possibly as socially beneficial by virtue of eliminating the deadweight loss from consumers whose willingness to pay is below the market price.

2.1 Does Piracy Harm Sales?

In this section we analyze what the academic literatures in economics, marketing, and information systems can tell us about how piracy impacts sales of media products. Within these literatures, we have chosen to focus on empirical studies of the impact of piracy because, while there are a variety of analytic models proposing theories of how piracy might impact sales,[1] we believe that the true test of these theories starts with data.

Based on our review of the empirical literature we conclude that, while some papers in the literature find no evidence of harm, the vast majority of the literature (particularly the literature published in top peer reviewed journals) finds evidence that piracy harms media sales.

In the remainder of this section we first discuss our philosophy in evaluating the literature. We then present a high-level overview of the statistical challenges associated with

[1] See Peitz and Waelbroeck (2004) for a review of the analytic/theoretical literature on the impact of piracy. Within this literature, Chellappa and Shivendu (2005) propose a model whereby consumers may pirate products as a "sampling" strategy — using the pirated content to learn more about the true value of the content and then using that information to decide whether to buy the content. Similarly, Conner and Rumelt (1991) and Takeyama (1994) propose a theory of network effects whereby the piracy of media products by a portion of the market might increase the value to other participants in the market. Finally, some have argued that indirect appropriability, a term coined by Liebowitz (1985), might mitigate the impact of media piracy by increasing the value of the initial purchase: i.e., if I can make a copy for my friend, I might be willing to pay more for the initial product. As noted by Liebowitz (2008a), absent empirical validation serious questions can be raised regarding whether any of these theories hold in actual markets. This is why we have focused our analysis on what the data actually say.

measuring the impact of piracy, and then summarize the three main methods used in the literature for addressing these challenges. Finally we discuss the specific findings of the major papers in the literature as they pertain to the impact of piracy on music sales and on motion picture sales. We focus on music and motion picture sales because these are the media categories that are most commonly addressed in the piracy literature. Whether the results observed for music and motion picture content extend to other product categories such as books and video games, remains an open, and important, question.

Epistemology and Social Science:

One of us (Smith) once heard a physics professor say that, outside of pure mathematics, there is no such thing as a "proof" in the physical sciences. Instead, the physical sciences must deal with knowledge on the level of what is the most reasonable explanation for the observed data.

If this is true for the physical sciences, it is certainly true of the social sciences where the observed data are noisier, more incomplete, and subject to the vagaries of human behavior. As such, we believe that evaluations of the literature should start with the recognition that there is no such thing as a perfect or completely conclusive paper. Each paper has flaws, limitations, and areas that could be improved with better data or different methods. Because of this, we believe that when evaluating "what the literature says," one should start by first analyzing what each individual paper finds, but then one should take a step back and draw overall conclusions based on what the totality of the literature says. We try to follow this approach in the discussion that follows.

Methodological Approaches:

A naive approach one might take to estimating the impact of piracy is to use data on sales of individual products, say movies, and measures of piracy levels for those movies. One could then run the following regression in an attempt to measure how piracy impacts sales:

$$S_{it} = X_{it}\delta + P_{it}\beta + \varepsilon_{it} \quad (1)$$

where S_{it} represents the sales of movie i at time t, X represents a matrix of variables, possibly including fixed effects for each movie and each time period or control variables, and where P_{it} represents piracy levels on movie i at time t. In this specification, the impact of piracy on sales would be given by the β coefficient.

The problem with this approach is that there exist variables that are not included in the regression but still affect variables on both the left-hand (independent) and right-hand (dependent) sides of the regression. In this specific case, one might expect that the popularity of a movie (which is unobserved) would impact both the likelihood that it will be pirated and the likelihood that it will sell. In a situation like this, where missing variables affect both the independent and dependent variables in a regression, the affected dependent variables are referred to as being "endogenous," and it is well known that the resulting regression coefficients will be inaccurate (biased).

Indeed, in the specific case described above, if one regresses the sales of individual movies onto the piracy levels of those movies, one is very likely to see a *positive* coefficient on piracy. If the model were correct, the interpretation of this coefficient would be that increased piracy *helps* sales. However, as noted above, in this case one cannot make that interpretation because of the bias introduced by the unobserved variables.

Helberger, Huygen, and van Eijk (2009) represent a potential "real-world" example of this error. The authors survey 1,500 media customers and pirates in the Netherlands regarding their purchasing and piracy behavior. They find that media pirates purchase as many CDs as non-pirates do. In a section titled "downloading and buying as complementary activities" they conclude from this observation that "[d]ownloading need not be a threat to purchases of physical formats: it would seem that for Dutch consumers these go together" (p. 75). The problem with this conclusion is that it ignores a potential endogeneity problem: If there are unobserved characteristics of consumers (say their interest in music) that might influence both a consumer's propensity to pirate and their propensity to purchase, then one cannot conclude anything about a survey that finds that pirates purchase as much as non-pirates

do. Put another way: while pirates might purchase as much as non-pirates do, we have no way of knowing how much these pirates would have purchased if piracy weren't available.[2]

Below, we review the four major methodologies used in the academic literature to address this endogeneity problem either directly (in the case of product-level data using natural experiments and instrumental variables approaches) or indirectly (using country- or city-level data or survey data).

Product-Level Analysis Using Natural Experiments: Controlled experiments are the "gold standard" of social science research. For example, a controlled experiment to analyze the impact of piracy on sales might involve obtaining a random sample of 10,000 or so consumers, randomly assigning half of them to a treatment group that must stop using Internet piracy for a period of time, and then comparing the purchase behavior of the treatment group to the control group of customers who are still able to access piracy. This approach would not suffer from the endogeneity problem described above because the decision about which consumers' behavior is left unchanged (control group) and which are no longer allowed to pirate (treatment group) is unrelated to the dependent variable (media sales).

The problem with this approach, of course, is that it is very difficult both in terms of effort and money, and thus it is not surprising that there are no papers in the literature that we are aware of that use controlled experiments to study piracy.

However, an analogous approach to the pure experiment described above involves using a "natural experiment" where a treatment is applied to one group of consumers and where the researcher can find another group of similar consumers who are unaffected by this change to serve as the control group. We have used this approach in several of our papers.

[2] Another concern with this paper is that the authors conclude that even though piracy reduces sales, increased piracy represents a net welfare gain to Dutch society. They reach this conclusion by viewing piracy as a welfare transfer from artists to consumers and noting that by reducing the cost of content to zero, piracy eliminates the deadweight loss from consumers whose utility for the music was below the market price. The problem with this conclusion is that it takes music production as a given. As we discuss in more detail below, if piracy reduces rents available to artists, and if artists were to produce less music (or lower quality music) as a result, then total social welfare could decline.

Danaher et al. (2010) use NBC's decision to remove its content from iTunes as a natural experiment and compare piracy levels and sales for NBC content (the treatment group) to ABC, CBS, and FOX content (the control group). Similarly, Danaher et al. (2013) use ABC's decision to add its content to Hulu as a natural experiment and compare piracy levels for ABC content (treatment) to NBC, CBS, and FOX content (control). Likewise, Danaher et al. (2013) use the "HADOPI" graduated response law in France as a natural experiment and compare music sales by French customers (treatment) to sales in a set of other European countries (control). Finally, Danaher and Smith (2013) use the shutdown of the site megaupload.com has an experiment of sorts, analyzing how the impact of this shutdown varies across countries with different levels of pre-shutdown usage.

The challenge with this approach is finding a suitable control group (one that has similar characteristics to the treatment group prior to the event), and finding an event that is both exogenous (i.e., is not driven by the dependent variable) and is sufficiently discrete that one can observe changes "before" and "after" it took effect.

Product-Level Analysis Using Instrumental Variables: The instrumental variables approach is similar to the "natural experiment" approach in that the researcher needs to find a variable that is correlated with the endogenous dependent variable (e.g., piracy) without being correlated with the independent variable (legal consumption).

Oberholzer-Gee and Strumpf (2007) (a paper we discuss in more detail below) apply this approach by using the number of German secondary school students who are on vacation in a particular week as an instrument for the ease of piracy among U.S. citizens. For this to work, German holidays must affect the ease of piracy in the U.S. (which the authors argue occurs because many music file sharers are German students and because German students are more likely to share files when they are not in school), and German school holidays must be otherwise uncorrelated with U.S. music sales.

In addition to the Oberholzer-Gee and Strumpf paper, Zentner (2006) and Rob and Waldfogel (2006) have used consumers' Internet sophistication and the availability of broadband as instruments for piracy. However, because of the difficulty in finding a

variable that is both correlated with the endogenous dependent variable and not directly correlated with the independent variable, beyond these papers there have been relatively few piracy-related papers that have used the Instrumental Variables approach.

City- or Country-Level Data: The third main approach used in the empirical literature is to compare sales levels across different geographical markets (typically either countries or cities). The basic idea of this approach is that after controlling for differences in the demographic characteristics of each region, changes in the dependent variable (typically broadband Internet penetration) can be treated as an experiment, and the researcher can statistically compare the change in sales resulting from this experiment for a treatment group (those regions where broadband Internet penetration increased) to a control group (regions where there was no change in broadband Internet penetration).

Examples of this approach include Hui and Png (2003), Pietz and Waelbroeck (2004), Zentner (2009), Zentner (2012) for country-level data; and Liebowitz (2008b), Zentner (2006) and Smith and Telang (2010) for city-level data.

The main challenge of this approach is ensuring that the observed changes (for example changes in the propensity to adopt broadband Internet) are properly controlled for by the demographic characteristics of the region or by other control variables available to the researcher. For example, if there were unobserved characteristics of regions that were both driving Internet adoption and were driving media sales, and that weren't captured by observable demographic characteristics, the resulting coefficients would have similar endogeneity bias to those described above. A related challenge is that while broadband Internet access can stimulate piracy, it can also influence users in many different ways (for example provide users with other entertainment options) that can affect media sales.

Individual-Level (Survey) Data: A final category of papers use surveys of small relatively homogeneous samples of consumers, for example, college students taking an economics class (Waldfogel (2009, 2010), Rob and Waldfogel (2006, 2007), Bai and Waldfogel (2012)). In this approach, the researchers use the homogeneity of the sample, along with observed demographic and psychographic characteristics to control for unobserved correlation with

the dependent variable. The main interpretation challenges with this approach are (1) that the resulting conclusions are, of necessity, tied to the chosen sample and (2) that the stated behavior of individuals can be systematically affected by both inaccurate recall and by obfuscation. The first challenge means that, in the case of the Waldfogel papers cited above where the surveys are obtained from college students, it is difficult to generalize any results beyond college students. The second challenge may affect survey data to the extent that consumers systematically over- or under-estimate their actual purchase behavior, or intentionally obfuscate the impact of piracy on sales behavior.

The Impact of Piracy on Music Sales:

One piece of anecdotal evidence for why piracy may have harmed music sales revolves around the decline in sales in the music industry shortly after the introduction of Napster in 1999 (Figure 1). Unfortunately, using this data alone one can't tell what music sales levels would have been in the absence of Napster, or how much of the observed decline in sales can be explained by piracy versus other unrelated causes. Answering these questions requires some of the statistical techniques discussed above.

Figure 1: Global Music Sales (1990-2003)

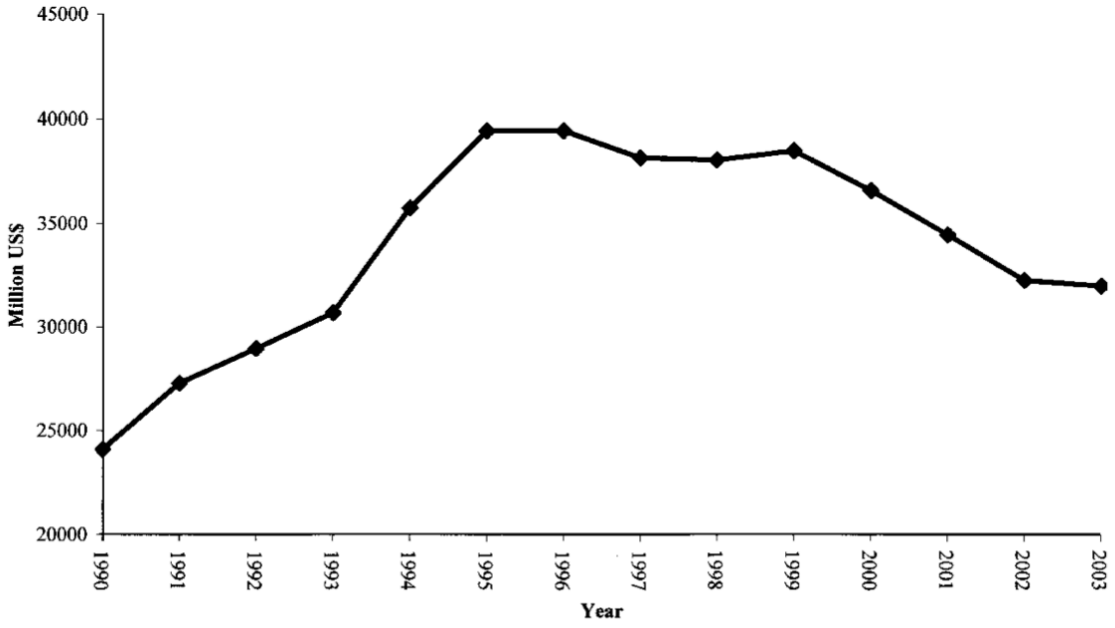

Source: Zentner (2006), p. 64

Below we review the major academic papers that have looked at the impact of piracy on music sales. While the vast majority of these studies have found some harm from piracy, one of the earliest and most prominent studies found no evidence of harm from piracy, and we start our discussion with that paper.

Oberholzer-Gee and Strumpf: The Oberholzer-Gee and Strumpf (2007) paper is not only one of the first papers in the economics literature to look at the impact of piracy on sales, it was also published as the "lead article" in the *Journal of Political Economy*, one of the most well respected journals in the economics literature. Because of this, it is one of the most cited papers on the impact of piracy on sales.

In the paper, the authors analyze data from the latter part of 2002. Their data include U.S. piracy downloads by album (collected from two large OpenNap servers) and U.S. album sales. To address the endogeneity problem described above, Oberholzer-Gee and Strumpf use the number of German secondary school students who are on vacation in a particular week. The authors argue that German school vacations are an appropriate instrument for the ease of piracy in the U.S. because German users provide one out of every six files that are downloaded by U.S. users, and because during school vacations there are more files shared by German users (students have more time to spend online sharing files). After using this instrument, Oberholzer-Gee and Strumpf find that file sharing has a statistically insignificant impact on music purchases.

On one hand, we believe the authors deserve "full marks" for being the first to address a very interesting question, for using a very creative combination of datasets to answer the question, for being the first to think carefully about how to measure piracy levels using data from pirate networks, and finally for developing a creative instrument to break the endogeneity problem in the data.

However, there have also been some significant concerns raised about their methods. The most forceful objections have been raised by Stan Liebowitz (e.g., Liebowitz 2007, Liebowitz 2010). Liebowitz's main critique of the Oberholzer-Gee and Strumpf's work focuses on the appropriateness of their chosen instrumental variable. Specifically, as noted

above, for this instrument to do its job it must be correlated with the ease of file sharing in the U.S., and uncorrelated with U.S. sales. Liebowitz (2010) argues that it fails in both respects. Specifically, Liebowitz (2010, p. 3) argues that German school holidays are negatively correlated with U.S. sales, primarily through a correlation with the Christmas holiday season.

Liebowitz also argues that the Oberholzer-Gee and Strumpf data likely overstate the impact of German users on U.S.-based downloads. Specifically, Liebowitz observes that Oberholzer-Gee and Strumpf's estimate indicate that a one standard deviation increase in German students on school vacation (an increase of 3.6 million students) would predict a 50% increase in American file sharing (an increase of about 2.2 million downloads). Liebowitz argues that this predicted increase is unreasonably large, so large that in any weeks where no German schools kids were on vacation (which occurs in 7 out of 17 weeks in the Oberholzer-Gee and Strumpf data), that all U.S. file sharing should fall to zero.

Liebowitz is certainly the most forceful, but is not the only academic to express reservations about Oberholzer-Gee and Strumpf's work. Rob and Waldfogel (2006) also critique Oberholzer-Gee and Strumpf's for relying on the contemporaneous relationship between piracy and sales (does piracy in a particular week reduce sales in that same week?) as opposed to focusing on the impact of piracy over a longer time horizon, and also expresses concern about the use of piracy and sales data whose correlation likely biases the results toward finding no effect.

In short, while we emphasize that Oberholzer-Gee and Strumpf deserve credit for writing the first major paper on this question and for doing so using an innovative dataset and innovative methods, one may wish to view their result relative to the large number of more recent papers, and more recent datasets, finding a strong and significant impact of piracy on music sales. We review this literature below.

Other Papers: In addition to the Oberholzer-Gee and Strumpf paper, there are three other academic papers we are aware of that find no evidence that piracy harms sales. First, Boorstin (2004), in his undergraduate thesis at Princeton University, used census data on

the number of individuals in each city with and without broadband Internet access in 1998, 2000, and 2002 and combined this data with CD sales data for the same metro areas. However, while Boorstin finds no negative effect of broadband Internet penetration on music sales, Liebowitz (2005) uses the same dataset used by Boorstin and argues that after controlling for demographic characteristics at a city level (which might influence Internet adoption), and after adding year-level fixed effects, the results show that piracy harms sales.

Andersen and Frenz (2010) also find no evidence of harm from file sharing after analyzing the results of survey responses from Canadian file sharers taken in 2006. However, in a recent paper, Barker and Maloney (2012) find the opposite result in this dataset after correcting for two "fundamental errors" in the econometric analysis in Andersen and Frenz's. The first major change made by Barker and Maloney is restoring over 400 respondents to the analysis. Andersen and Frenz exclude over 400 individuals from their analysis (about 20% of their sample) because they did not purchase CDs in 2005, arguing that these consumers "may never have been active in CD purchasing." Barker and Maloney note that since the survey data include both CDs sales in 2005 and 2004, one can partially test this assumption (a test that Andersen and Frenz surprisingly did not conduct). And in fact that nearly a third of the excluded respondents had purchased CDs in 2004 even though they did not purchase in 2005. Barker and Maloney suggest that many of these consumers who stopped purchasing CDs may be exactly the sort of customers whose purchases were most affected by the availability of piracy. Second, Barker and Maloney specify a system of both CD and P2P demand, an econometric change that partially controls for potential endogeneity problems from unobservable consumer-level characteristics that may drive both CD and P2P demand. After making these changes Barker and Maloney find that the Canadian data actually reflects a strong and consistent negative impact of piracy on sales.

Hammond (2012) is the third paper we are aware of that finds no harm from file sharing on music sales, albeit in a nuanced way. Hammond's analysis occurs in the context of pre-release leaks of CDs, with data obtained from a private tracker site specializing in pre-

release file sharing. Hammond finds that additional downloads of an album already available on pre-release piracy sites causes no statistical change in post-release sales. However, Hammond notes that this result doesn't contradict the "well-documented fact that file sharing is harmful to the music industry (p. 1)." Rather it suggests that given the presence of rampant piracy, additional piracy is not harmful to an individual artist.

With the exception of these papers, all of the other papers we are aware of in the literature find that music piracy has harmed sales. We review these papers briefly below.

In the context of survey data, Zentner (2006) uses a sample of 15,000 people in 2001 and 2002 and finds that, after controlling for Internet sophistication and broadband speed, peer-to-peer usage reduces by about 30% the probability that an individual will purchase music, and overall that piracy reduced music sales by about 7.8% in 2002. Rob and Waldfogel (2006) survey the piracy and music purchasing behavior of 412 college students at 4 colleges in 2003 and find that each pirated download displaces about 0.2 album sales, and that overall piracy reduced per capita expenditures on music by about 20%. Waldfogel (2010) uses a survey of University of Pennsylvania undergraduates in January 2009 analyzing piracy and purchase behavior for music, and finds that each pirated download displaces between 0.15 and 0.3 album sales. Finally, in the context of Census data, Michel (2006) and Hong (2004) use purchase data from the Consumer Expenditure Survey combined with census data on broadband usage and find that file sharing causes a 5 to 7.6% reduction in sales.

In the context of country-level cross sectional data, Hui and Png (2003) use country-level data for 28 countries from 1994-1998 and find that physical piracy reduces sales by about 42%. Pietz and Waelbroeck (2004) use CD sales for 16 countries from 2000-2001 and find that piracy explains about 25% of the decline in music sales observed over that time frame. Finally, Zentner (2009) uses country-level music sales and broadband penetration for 49 countries from 1997-2008 and finds that file-sharing may explain up to 50% of the decline in music sales observed during that period.

As noted above, it is also possible to conduct similar analyses using city-level (MSA-level) data. This approach has been used by Zentner (2006) and Liebowitz (2008b) for broadband penetration and music sales major cities in the U.S. from 1998-2003, with both papers finding that piracy (via broadband availability) caused a significant portion of the observed decline in CD sales during this period.

In the context of the impact of piracy on the distribution of revenue between CD sales and live performances, Mortimer, Nosko, and Sorensen (2012) find that while piracy displaces CD sales, it increases concert revenue for less well-known artists.

A final set of papers analyze how piracy impacts the composition of "best of" and bestselling lists of music, with Bhattacharjee et al. (2007) finding that piracy reduces the duration albums spend on bestselling charts, and with Waldfogel (2011) finding that file-sharing does not seem to impact the quality of music appearing on "best of" lists.

The Impact of Piracy on Motion Picture Sales:

Given that the academic literature seems to show that piracy hurts music sales, why might motion picture piracy have a different effect? Two obvious differences between motion picture and music piracy are first that motion picture files are typically much larger than music files, and second that motion picture piracy developed into a significant phenomenon with the development of the BitTorrent protocol (in 2003), well after the mainstream development of music piracy (with Napster in 1999). Based on these differences, we might expect that the impact of motion picture piracy would be seen later than the impact of music piracy was, and that motion picture piracy might be more strongly tied to the presence of broadband Internet connections. And this is exactly what we see in the literature.

This effect can be seen quite clearly when analyzing the differences between three papers in the academic literature. First, Liebowitz (2008b), discussed above, used broadband Internet penetration and music sales at an MSA level for the 99 largest MSAs from 1998-2003 to show that increased broadband penetration led to a sharp decline in music sales

during this timeframe. Note that this time period was after the introduction of Napster in 1999 and before the introduction of BitTorrent in 2003.

In contrast to these results, Smith and Telang (2010) uses essentially the same methods as Liebowitz used for music, but use DVD sales data and broadband Internet penetration for 2000-2003. Their results also differ from Liebowitz's results: increased broadband penetration caused about 9.3% of the $14.1 billion *increase* in DVD sales from 2000-2003.

Their interpretation of the difference between their results those in Liebowitz (2008b) is that, while music piracy was prevalent from 2000-2003, movie piracy was much less developed during that period. As such increased broadband penetration from 2000-2003 may have exposed DVD consumers to the beneficial aspects of the Internet (increased information about movies, increased product selection through online retailers, and lower prices) without exposing these customers to the potentially harmful aspects of the Internet (increased availability of pirated content).

Based on this we would expect that increased broadband Internet penetration would harm DVD sales after 2003. And indeed evidence from descriptive statistics on DVD sales seems to bear this out. Figure 2 is taken from Zentner (2012) and shows that DVD sales flattened in 2004, and that combined VHS and DVD sales dropped by 27% from 2004 to 2008 when measured in constant 2008 dollars.

Was this decline causally related to an increase in broadband-enabled Internet-based piracy from 2003-2008? Zentner (2012) finds that it was. Specifically, Zentner uses country-level panel data from 2001 through 2008 for 36 European countries and the U.S. and Japan. These data document theatrical revenue, video rental, and video sales data for movies, along with broadband Internet penetration for the time period. The use of panel data allows Zentner to control for factors unrelated to broadband penetration in a way that is not possible if one were to only observe time series or cross sectional data. Zentner attempts to isolate the impact of broadband-enabled movie piracy by comparing the impact of increased broadband penetration before and after the widespread adoption of BitTorrent in 2003. Zentner finds that prior to the introduction of BitTorrent, increased

broadband penetration had a positive impact on motion picture sales (consistent with Smith and Telang 2010), but after 2003, increased broadband penetration had a negative impact on sales.[3]

Figure 2: U.S. Video Rental and Sales Revenue, 1994-2008
(Measured in Constant 2008 Dollars)

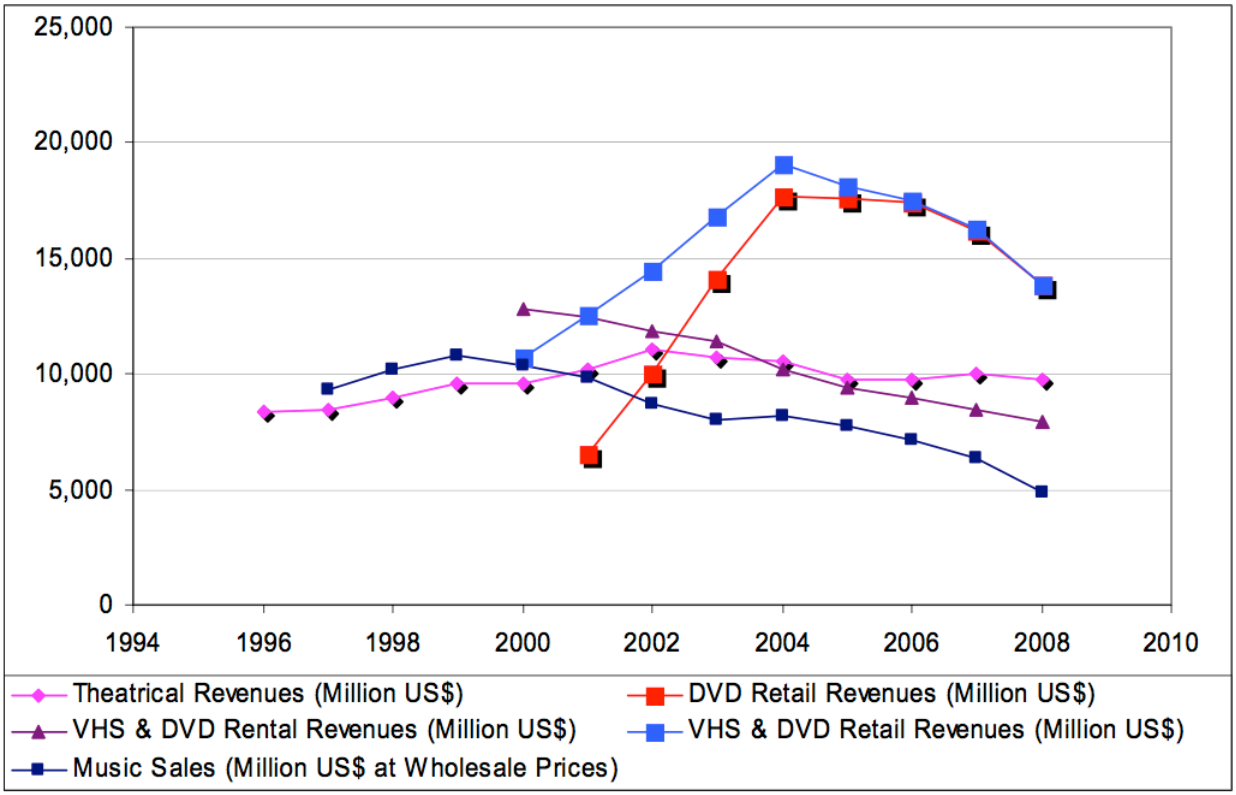

Source: Zentner (2012), p. 2

This result — that piracy harms motion picture sales — is consistent with all but two of the academic papers we are aware of that have looked at the impact of Internet piracy on movie sales, a set of papers that span a variety of datasets, settings, and statistical methodologies.

In the context of surveys, Bounie, Bourreau, and Waelbroeck (2006) use data on purchase and file sharing behavior for a sample of 620 French individuals in 2005, and find that,

[3] Zentner also observes (p. 25) that legal online video sales and rentals were a very small part of the market (approximately 1% of DVD sales and rentals) even as late as 2008.

while file sharing has no statistical impact on theatrical attendance, file sharing leads to a large decline in both video sales and video rentals. Similarly, Rob and Waldfogel (2007) conducted a survey of 500 University of Pennsylvania undergraduates in the fall 2005 semester. These students were given a list of the top 50 movies from each of the 3 previous years (150 movies in all) and asked whether they saw the movie, and what channel (piracy, theater, television, rental, purchase) they used to view the movie. The authors used this data to find that unpaid consumption of movies reduces paid consumption on nearly a one-for-one basis.

Likewise, Hennig-Thurau, Henning, and Sattler (2007) use a combination of customer's stated intentions and reported behavior to analyze the impact of file sharing on movie sales and find that file-sharing results in significant cannibalization of theater visits, DVD rentals, and DVD purchases. Finally, Bai and Waldfogel (2012) use a survey of college students in China and find that 75% of Chinese movie consumption is through pirate channels, and that each instance of unpaid consumption displaces about 0.14 paid sales.

In the context of "natural experiments," Danaher et al. (2010) use data surrounding NBC's decision to remove its television content from the iTunes store in September 2008. They find that piracy on NBC content increased by 11.4% relative to ABC, CBS, and FOX piracy, after NBC's content was removed from iTunes. They also find that piracy levels on ABC, CBS, and FOX content increased during this timeframe as well, potentially suggesting that NBC's decision to remove it's content from iTunes caused an increase in both NBC and non-NBC piracy as users switched from iTunes to BitTorrent to obtain their television content.

In the context of product level data, Danaher and Waldfogel (2012) and Smith and Telang (2013) both examine the impact of longer international release windows on sales levels. Danaher and Waldfogel use box office data from July 2003 through July 2006 for the top 10 movies in each time period across 17 countries and find that, after the widespread diffusion of BitTorrent, longer lags between the U.S. and the international release date cause about a 1.3% reduction in sales per week, or about a 7% reduction per movie. Smith and Telang (2012) use a similar method, but focus on DVD sales. Specifically, Smith and Telang use data from 2009 through 2011 for seven countries covering DVD sales for over

200 titles. They find that a 1-week longer release delay between the first DVD pirate source and the legitimate DVD release date is correlated with a 2% reduction in DVD sales per movie. Likewise, De Vany and Walls (2007) use piracy and revenue data for a major studio release and find that file-sharing caused a movie to lose $40 million in box office revenue. Finally, Ma et al. (2013) use the timing of the availability of theatrical piracy relative to the box office release date and find that pre-release piracy can cause an 8% reduction in box office sales, and that overall box office revenue would be 24% higher if no piracy were to occur in the theatrical window.

In the context of panel data, Zentner (2012) uses theatrical revenue and home video sales (VHS and DVD) for 36 countries from 1996-2008 to analyze whether Internet penetration reduces movie sales. He finds no statistical relationship between increased broadband Internet penetration and theatrical revenue, but a strong negative relationship between increased broadband penetration and DVD sales. Similarly, Liebowitz and Zentner (2012) use panel data for major cities in the U.S., including Internet penetration, demographics, and television viewership to analyze whether increases in broadband Internet penetration reduces television viewership. While not a direct measure of the impact of piracy, they argue that broadband penetration may represent an indirect measure of piracy given the prevalence of pirated television content on the Internet. Their results are consistent with the results reported above in that they find a moderate reduction in television viewership among younger viewers caused by increased broadband Internet penetration.

Martikainen (2011) is one of the two academic papers we are aware of that find no evidence of harm from file sharing. Martikainen uses BitTorrent download data, collected from March to May 2009 and finds no evidence that increased levels of BitTorrent sharing reduce DVD sales. However, it is also important to note that this paper is only able to analyze DVD sales, not digital sales; and other work in the literature (e.g., Danaher et al. 2010) finds that during a similar time frame, digital sales and digital piracy were strongly related, while there was no statistical relationship between digital piracy and physical sales.

The other academic paper we are aware of that finds no evidence of harm from piracy is Smith and Telang (2009) which analyzes how movie broadcasts on advertising supported

television stations impact demand for the movie through legal (DVD) and illegal (piracy) channels. This paper finds that movie broadcasts stimulate demand in both legal and illegal channels, and uses the broadcast of movies on television as an exogenous shock to demand. The paper finds that movies that have pirated copies available at the time of broadcast have a similar increase in post-broadcast DVD sales versus movies that do not have pirated copies available. The paper argues that this suggests that digital piracy at the time a movie is shown on television does not significantly impact DVD sales. However, the paper also notes that "our results do not speak to the impact of piracy in the earlier part of a movie's lifecycle, where the availability of pirated content may have a negative impact on sales" (p. 336).

Returning to our initial statement that one should determine what the literature says by first looking at individual papers and then looking at the weight of the literature's finding, our review of the literature suggests that while it is fair to say that the results in the academic literature are mixed with respect to whether file-sharing harms sales, we also believe it is fair to say that the vast majority of papers find evidence of harm.

Specifically, restricting attention to papers published in peer-reviewed journals, our review of the literature finds three papers that find no evidence of harm from piracy (see Table 1) and 15 papers finding that piracy results in a statistically significant decrease in sales (see Table 2).

Table 1: Peer-Reviewed Journal Articles Finding No Statistical Impact of Piracy

Citation	Primary Data	Result
Oberholzer and Strumpf (2007, *Journal of Political Economy*)	2002 OpenNap music downloads, 2002 U.S. sales of popular albums	"[F]ile sharing has had no statistically significant effect on purchases of the average album in our sample."
Smith and Telang (2009, *MIS Quarterly*)	2005-2006 Amazon DVD sales ranks and BitTorrent movie file downloads	"[T]he availability of pirated content at [television broadcast] has no effect on post-broadcast DVD sales gains.
Andersen and Frenz (2010, *J. of Evolutionary Economics*)	2006 survey of Canadian customers' file sharing and CD purchasing behavior	There is "no (statistical) association between the number of P2P files downloaded and CD album sales."

Table 2: Peer-Reviewed Journal Articles Finding That Piracy Harms Sales

Citation	Primary Data	Result
Hui and Png (2003, *Contrib. to Economic Analysis & Policy*)	1994-98 IFPI worldwide CD sales data and physical piracy rates	"[D]emand for music CDs decreased with piracy, suggesting that 'theft' outweighed the 'positive' effects of piracy."
Peitz and Waelbroeck (2004, *Rev. of Econ. Res. on Copyright*)	1998-2002 worldwide CD sales, IPSOS survey data for piracy downloads	Internet piracy may have been responsible for a 20% decrease in music sales between 1998-2002.
Bounie et al. (2006, *Rev. of Econ. Res. on Copyright*)	2005 survey of movie piracy and purchases from French universities	"[Piracy] has a strong [negative] impact on video [VHS and DVD] purchases and rentals" but statistically no impact on box office revenue.
Michel (2006, *Topics in Economic Analysis and Policy*)	1995-2003 U.S. BLS micro Consumer Expenditure Survey data	"The relationship between computer ownership and music purchases weakened" due to piracy, potentially reducing CD sales by 13 percent.
Rob and Waldfogel (2006, *Journal of Law and Economics*)	2003 survey of U.S. college students piracy and purchase behavior	"[E]ach album download reduces purchases by 0.2 in our sample, although possibly by much more."
Zentner (2006, *Journal of Law and Economics)*	2001 survey of European music purchase and piracy behavior	"[Piracy] may explain a 30 percent reduction in the probability of buying music."
Bhattacharjee et al. (2007, *Management Science*)	1995-2002 Billboard 100 chart rankings, WinMX file sharing post 2000	P2P file sharing technologies have resulted in "significantly reduced chart survival except for those albums that debut high on the charts."
Hennig-Thurau, Henning, Sattler (2007, *Marketing Science*)	2006 survey of German movie purchase and piracy intentions	Piracy causes "substantial cannibalization of theater visits, DVD rentals [and] purchases responsible for annual revenue losses of $300 million in Germany."
Rob and Waldfogel (2007, *Journal of Industrial Economics*)	2005 survey of U. Penn. students' movie purchase and piracy behavior	"[U]npaid first [piracy] consumption reduces paid consumption by about 1 unit."
Liebowitz (2008, *Management Science*)	1998-2003 Census data on broadband Internet use and music purchases	"[F]ile sharing appears to have caused the entire decline in record sales [observed from 1998-2003]."
Danaher et al. (2010, *Marketing Science*)	2007-2008 BitTorrent downloads of television torrents	The removal of NBC content from iTunes resulted in an 11.4% increase in demand for NBC piracy relative to ABC, CBS, and FOX piracy.
DeVany and Walls (2007, *Review of Industrial Organization)*	Box office revenue and the supply of pirated content for an unnamed movie	"[Piracy] of a major studio movie accelerated its box-office decline and caused the picture to lose about $40 million in revenue."
Waldfogel (2010, *Information Economics and Policy*)	2009-2010 survey of Wharton students' music piracy and purchases	"[A]n additional song stolen reduces paid consumption by between a third and a sixth of a song."
Bai and Waldfogel (2012, *Information Economics and Policy*)	2008-2009 survey of Chinese university students' movie behavior	"[T]hree quarters of [Chinese students'] movie consumption is unpaid and … each instance of [piracy] displaces 0.14 paid consumption instances."
Danaher et al. (2013, *J. of Industrial Economics*)	2008-2011 iTunes music sales in France and other European countries	The HADOPI anti-piracy law "caused iTunes music sales to increase by 22-25% [in France] relative to changes in the control group [countries]."

Moreover, if one were to further restrict attention to what are commonly considered "first or second tier" academic journals,[4] and also eliminate the Smith and Telang (2009) paper which focuses on the impact of piracy of catalog titles (12-18 months after the DVD release), the count would be one published paper finding no evidence of harm (Oberholzer-Gee and Strumpf) and 11 published papers finding evidence of harm.

To summarize, while the academic literature is not uniform in finding harm, taken as a whole we see a very consistent story across the academic literature: The vast majority of papers which have been published in peer-reviewed academic journals — papers spanning a variety of methods, time periods, and contexts — find that piracy causes a statistically significant decrease in sales. Moreover, the conclusion that piracy harms sales is particularly strong among papers published in first of second tier academic journals and among papers analyzing the most recently available data.

2.2 The Impact of Piracy on Welfare Creation and Distribution

In this section we analyze two important social welfare questions surrounding piracy: Does piracy reduce overall social welfare by reducing the incentives to create content and does piracy change the distribution of welfare between artists and firms?

Analyzing the impact of piracy on overall social welfare is important because even if one accepts that filesharing displaces sales, it is worth asking whether policymakers should be concerned with this displacement. In instances where illegal downloaders value a product below its purchase price, pirating that product does not result in a lost sale but provides value to the consumer that would not have been obtained without the illegal channel. In instances where pirates value the product above its purchase price but choose to pirate it for free, a sale is lost. However, the surplus lost by the producer of the good (the purchase price) is simply transferred to the consumer (who obtained the good for free instead of

[4] Which in this case would include *Information Economics and Policy, Journal of Industrial Economics, Journal of Law and Economics, Journal of Marketing, Journal of Political Economy, Management Science, Marketing Science, MIS Quarterly, Review of Economics and Statistics, Review of Industrial Organization*.

having to purchase it), resulting in no first-order loss to social welfare. Thus there are three main reasons why one might oppose filesharing:

1. The simplest case in which one might oppose filesharing is when one's personal welfare is tied to profits of the creative industries. For example, a music artist has a personal incentive to care that creators are compensated for their efforts, and an employee of a movie studio may worry that if piracy diminishes studio profits she may lose her job. Because most evidence seems to indicate that filesharing does displace media sales, it is not surprising that these industries are frequently opposed to filesharing.

2. Another argument against filesharing could be made on philosophical grounds. Intellectual property rights are a part of the legal structure of many countries, and without such protection information goods such as movies or songs could easily be replicated and sold by anyone, limiting the ability of content creators to profit from their work. If one philosophically believes that creators have inherent rights as sole executors over how their creations are distributed or sold, then illegal filesharing damages our ability to enforce this right, and therefore should be opposed. Purely philosophical arguments are not the subject of this book however, and so we mention this argument but note that it is not one for which we can provide evidence for or against.

3. Finally, one might worry that if filesharing diminishes the ability of content creators to profit from their creative efforts (which appears to be the case), the incentives to bring new quality works to the market will also be diminished resulting in either fewer or lower quality works of creative media such as music and motion pictures. This is clearly a potential concern not only for producers but also for consumers and policymakers, because if a product that would have existed in the absence of filesharing is never created due to reduced incentives brought about by filesharing, both producers and consumers lose causing an overall net loss in social welfare. A related supply side argument is that while piracy in one region or country may not affect the decision of whether or not to create products, it may affect the availability

of products in those regions. For example, even if high piracy in some country X did not change a music label's decision of what artists to develop and what songs to offer, it may deter some distribution platforms or stores from offering content in country X due to reduced potential profits. If some consumers who would have bought the product cannot purchase due to lack of availability, social welfare could be lost.

These three categories are the basic sources of most opposition to filesharing.[5] However, the first category is limited to a particular set of agents and the second category relies on a philosophical argument that is difficult to address empirically. Because this chapter is focused on policy we will base our discussion on the third category of opposition to filesharing.

A necessary condition for filesharing to affect the supply of creative works is that content creators must be incentivized by the profits they can obtain from their works. While this may seem tautologically true to an economist, reasonable debate exists as to whether such an assumption is valid in the world of artistic goods. First, if profits surpass a level beyond which additional profits do not impact incentives to create, then some reduction in profits may not lead to a change in supply. Second, it may be that not all content creation is incentivized by profits. Hypothetically, there may exist a musician creating highly-valued music who would continue to work in exactly the same manner even in the face of diminished profitability. However, one might argue that the products brought to market involve not just the work of this artist but also others involved in the production chain, from sound engineers who master the album to label executives who see that the music is marketed and promoted to distributors who actually sell it. How necessary are these parts of the production process, and how motivated are they by profits?

[5] The notion has also been advanced that because piracy allows consumers to acquire goods such as songs or movies for free, it has the psychological effect of lowering their perceived value of such products, diminishing their willingness to pay. Rather than evaluate the validity of this claim, we simply note that if it is true it again serves to diminish the ability of creators to profit from their works and as such can be evaluated in light of the three categories of arguments already described.

Perhaps blockbuster movies represent a clearer example where creative works are likely incentivized by profits. Film studios are profit-maximizing firms who frequently invest hundreds of millions of dollars into the creation and promotion of movies in the hopes of making returns on these investments. If the profits on films are significantly reduced by piracy, then these studios will have less incentive to invest in the creation of such films. Reduced incentives might have a particularly large impact on more innovative films that might be considered riskier in terms of their chance of success. And thus there are reasons to believe that filesharing may impact the quality or quantity of goods brought to the market, but there are also reasons to consider that it may have less impact than one would expect in situations where profits are less of an incentive or where reduced profits are still enough to incentivize similar innovation. Thus the question of whether reduced profits from filesharing affects the supply of creative works is an empirical one and the answer may differ between industries or between different types of products.

In spite of this, there exists very little evidence as to whether filesharing impacts the supply of creative works. The question is difficult to answer for several reasons. First, like measuring the impact of piracy on sales, it is difficult to measure the impact of piracy on supply because we cannot observe the counterfactual — what creative works would have been created in the absence of piracy? Second, even if an experiment were to allow us to measure the counterfactual, the hypothesis is that piracy may impact the quantity or *quality* of creative works, and quality is difficult to measure. For example, if the number of songs created were to remain the same when piracy rises but the average quality of these songs were lower, how would we observe this lower quality in order to measure it? Third, even in situations where natural experiments cause shocks to piracy (such as the adoption of BitTorrent leading to a surge of movie piracy, or the shutdown of Limewire leading to a reduction in music piracy), while the impact on sales might be expected to be immediate, the impact on supply of creative content (if there is one) will likely be longer term as firms slowly respond to the changed environment, and adjust to new profit levels. Thus, these effects are harder to observe, as we do not know how much lag to expect between a technology or policy shock and any potential resulting change in supply of creative works.

One piece of evidence that begins to examine this question can be found in Waldfogel (2011). The goal of the Waldfogel study is to examine the supply of new music over time, looking at the development of Napster as a natural experiment that began a large rise in filesharing behavior, and a subsequent lowering of music industry profits. Waldfogel observes that the number of new albums is increasing over time, but that many of these newer albums are of lower quality as technological advancements have lowered barriers to entry. To fix quality at a certain time-invariant level, Waldfogel constructs several indices of the number of albums produced each year that eventually make it onto a critical "best of" list, such as Rolling Stone's 500 Best Albums list. Thus he is able to measure the number of new albums above a constant quality threshold and ask how this number changes over time. He finds that the post-Napster levels of this index are similar to pre-Napster levels, and that post-Napster trends do not depart significantly from pre-Napster trends. Some may argue that this definition of a constant threshold of quality could be flawed. Waldfogel also shows similar results when the term "quality" could be interchanged with "demand" or "appeal" by producing indices of the number of songs from each year that have sold over a certain number of copies.

Figure 3: Supply of New Albums Above Quality Threshold

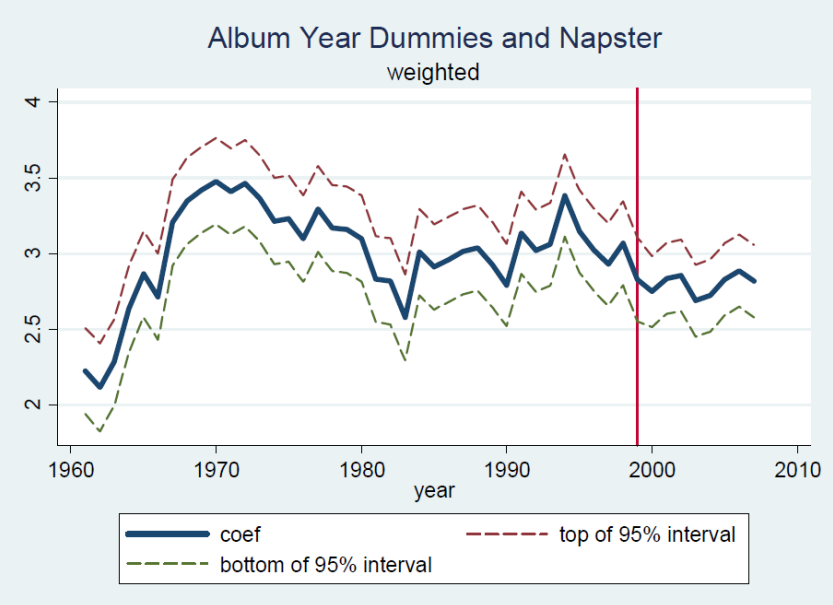

Source: Waldfogel (2011), p. 34

However, Waldfogel also acknowledges that the spread of filesharing coincided with advances in technology and the digitization of music, both of which reduced the production costs for music. Thus his conclusion is not that piracy does not impact the supply of creative works, but rather that the net effect of both technological advances in music production and distribution (including but not limited to filesharing) that occurred after 1999 had little impact on the supply of quality music. The counterfactual — what would have happened to the supply of quality music in the absence of piracy, but in the presence of other unrelated cost-reducing technologies — still remains unclear.

If anything, the Waldfogel study highlights the challenge of why it is so hard to estimate the link between filesharing and supply. Even with a good measure of supply and a strong experiment like the introduction of Napster, many things are changing over time that affect the production of creative content. Given this, it is not clear what lag we might expect to see between an experiment and a supply response, so we are left without a clear hypothesis test. Finally, Waldfogel highlights the need to consider not just the quantity but also the quality of what is supplied. We believe that his indices provide the best standard currently available for measuring quality with respect to music, but it is not clear how to transport such a concept to movies or other media. And in the end, his findings can speak only to the link between digitization and supply, where filesharing is only one aspect of digitization that may be counterbalanced by other positive aspects of digitization.

An interesting potential approach to studying how content industries change the supply of new quality products in the presence of reduced profits from filesharing is to consider product availability rather than product supply. While generally one believes that piracy accelerates the release of legal products, high levels of piracy may discourage firms from making their products available in markets with high piracy. Anecdotal evidence in countries like Spain, which is generally considered a high piracy country, suggests that firms are reluctant to release legal versions of their products or to invest in marketing and promotion of their products (New York Times, August 20, 2011). If the presence of piracy causes these responses, the outcome will be reduced social welfare to any consumers who would have benefitted from the market availability of these products. Marketing,

promotions and other services are also an important component of the product experience and a decision of reduce investments in such activities would also constitute social loss by reducing the information available to consumers about products they might benefit from consuming.

However, while current empirical evidence of a connection between piracy and product availability is sparse, given standard economic theory it seems reasonable to assume that some connection will exist. Hypothetically, imagine a world in which the moment a major blockbuster film is produced, piracy prevents the studio that developed it from extracting any revenue from theater ticket sales, DVD sales, or other legitimate channels. In other words, imagine a world where everyone can easily obtain motion picture content for free through piracy channels and as a result no one pays for this content. In this hypothetical world, is seems highly unlikely that studios would invest the large fixed costs necessary to develop blockbuster movies. Indeed, this is the very logic behind the copyright protections afforded to content creators by the U.S. Constitution. And thus we propose that, even in the absence of conclusive proof that piracy reduces the incentives to supply quality content, that there is theoretical rationale for governments to consider the possibility that filesharing may negatively impact social welfare.

A related welfare question is whether the availability of piracy impacts the distribution of welfare between artists and firms. In the context of the music industry, Mortimer, Nosko, and Sorensen (2012) analyze the impact of piracy on both sales of music and on the demand for live concert performances by artists. They find that while piracy reduces the demand for recorded music, it simultaneously increases the demand for live concerts, particularly among less well-known artists. Given that most artist contracts have allowed the artist to retain their revenue from concerts (even while they share revenue with the record label for recorded music sales), this shift away from recorded music sales toward concert sales effectively represents a shift of welfare away from labels (through reduced recorded music revenue) and toward artists (through increased concert revenue). However, it is important to note that labels have responded to this shift in revenue toward concerts by asking artists to sign “360 degree deals,” where artists agree to give the label a

percentage of their income through all channels (including concerts). It is also worth noting that we know of no studies in the literature that suggest a similar shift in revenue away from firms and toward artists in the context of the movie industry.

3. Responding to the Threat from Piracy

If piracy represents a threat to media sales, and potentially to creative incentives, the next logical question is how should governments and those in the creative industry respond? In this section we attempt to review what the academic literature can tell us about this question.

To frame this discussion, we start with the widely held view that media companies "can't compete with free." This view is most commonly advanced by those in the industry arguing for stronger anti-piracy legislation to eliminate the ubiquitous availability of pirated content. However, those arguing against strong copyright enforcement have also advanced this argument. For example, Nick Bilton (2012) in a New York Times editorial argued against copyright enforcement on the basis of the fact that copyright enforcement will never be fully effective and therefore should be abandoned.

In our discussion below we argue that the view that firms can't compete with free (pirated) content is flawed. Rather, competing with free can be seen as a special case of price competition. Firms are well acquainted with price competition: differentiate a higher priced product along a set of attributes that are valuable to the market and consumers will willingly pay more to purchase your product rather than purchasing a lower priced competing product.

We believe that there are direct analogies between this example of market-based competition and competition in the context of pirated products, but with a twist. In a typical market-based environment, firms can only differentiate their own product from their competition, while in the context of pirated goods, firms can both increase the attractiveness of their own products relative to pirated products (through attributes such as availability, convenience, and reliability), and they can also use anti-piracy interventions to reduce the attractiveness of competing pirated content relative to paid content. We

discuss each of these alternatives below, starting with the academic literature on the effectiveness of anti-piracy interventions.

3.1 The Effectiveness of Anti-Piracy Interventions

As we see it, anti-piracy interventions can be categorized along two axes. The first axis concerns whether the intervention is regulatory (government-driven) or voluntary (industry-driven), and the second axis relates to whether the intervention targets the supply- or demand-side of piracy.

Along the first axis, regulatory interventions include any government-sponsored effort to reduce piracy, while industry interventions include approaches through industry cooperation or market-based efforts to address piracy. Along the second axis, supply-side interventions target the supply of pirated content while demand-side interventions target individuals demanding pirated content. Figure 4 summarizes this categorization and includes example interventions in each category.

Figure 4: Categorization of Anti-Piracy Regulations

	Demand-Side	Supply-Side
Regulatory (Government-Driven)	HADOPI in France, IPRED in the EU	Megaupload Shutdown
Voluntary (Industry-Driven)	Copyright Alert System in the U.S., various industry lawsuits against file-sharers,	Voluntary changes in search engine ranking algorithms to reduce prominence of piracy in search results

In the context of regulatory, demand-side interventions, Danaher et al. (2013) analyze how the HADOPI law in France impacted French music sales by comparing music sales in France and a set of control group countries before and after the HADOPI legislation was passed. Their paper finds that HADOPI law in France caused a 20-25% increase in French music sales relative to the control group countries (see Figure 5).

Figure 5: iTunes Music Sales Before and After HADOPI

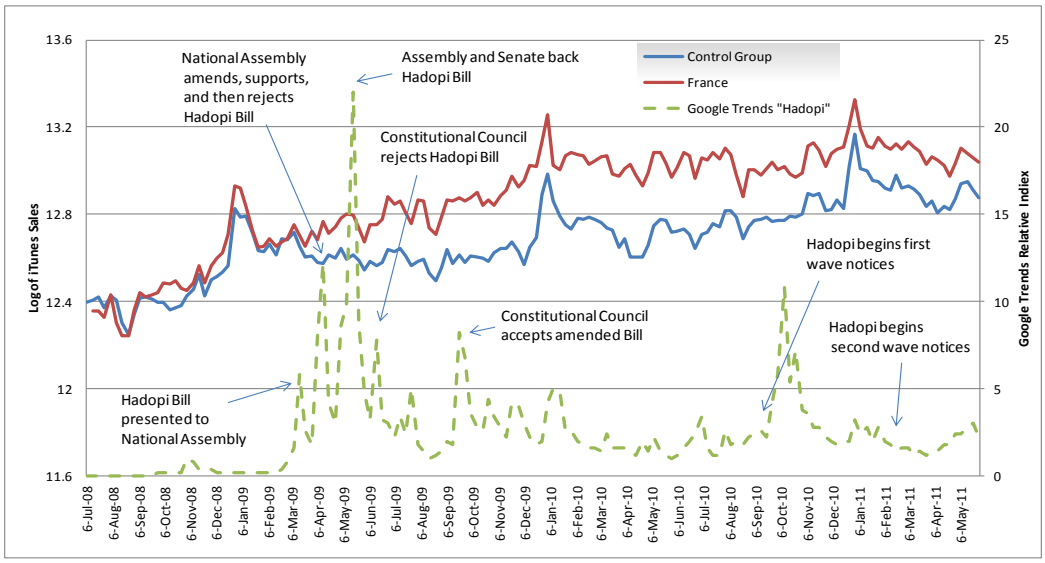

Source: Danaher et. al. 2012

The authors also run an additional level of differencing in their models by comparing changes in high piracy (i.e., rap, hip-hop), medium piracy (i.e., rock, pop) and low piracy genres (i.e., Christian, classical, folk, and jazz). The logic here is that higher piracy genres should have a higher number of "treated" consumers, and therefore should have a larger increase in sales after the HADOPI legislation was passed than other genres do. And, indeed, the authors find the increase in sales for the low, medium, and high piracy genres was 7%, 16% and 30% respectively.

Finally, the authors show that their results are robust to including iOS device sales, to eliminating any individual control-group country from their analysis, and to eliminating any individual record label from their analysis, suggesting that these results are not driven by changes in iOS device sales around the introduction of HADOPI, or by factors peculiar to any particular control group country or label. Finally, it is important to note that the reported results suggest that sales in France have increased *relative* to what they would have been (as proxied by the control group countries). Thus, while total recorded music sales (which include both physical and digital channels) have continued to decline in France after the law's introduction, the results suggest that absent the HADOPI law, the decline in sales would have been even more severe than what has been observed.

These results, showing an increase in sales following a regulatory demand-side intervention are similar to those reported by Adermon and Liang (2010), who find that enforcement of the European Union IPRED directive in 2009 caused a 27% increase in CD sales and a 48% increase in digital music sales in Sweden.

In the context of regulatory, supply-side interventions Danaher and Smith (2013) analyze the impact of the recent shutdown of Megaupload and its sister sites on digital motion picture sales. Megaupload and Megavideo were shutdown on January 18, 2012 after the U.S. Department of Justice obtained a grand jury indictment against the site's founders for copyright infringement. This analysis differs from the HADOPI analysis in that the HADOPI law affected French consumers, and as such, one could construct a set of control group countries to proxy for what sales in France would have been in the absence of HADOPI. In contrast, the Megaupload shutdown affected users in all countries worldwide. Thus, in the absence of a suitable control group of countries, Danaher and Smith use the fact that there was heterogeneity in Megaupload usage across countries to identify the impact of the Megaupload shutdown.

Specifically, the authors obtain data from two major motion picture studios documenting sales through their major digital channels for 12 countries where data were available. They then use Google AdPlanner data (showing the number of unique Internet users who visited Megaupload.com in the month before its shutdown) and International Telecommunication Union data (showing the number of Internet users in a particular country) to calculate the proportion of Megaupload users in a set of 12 countries. This "Megaupload penetration ratio" varies between a high of 17% in Spain to a low of 1.8% in the United States.

The author's identification strategy compares sales changes before and after the shutdown of the Megaupload site across countries with different penetrations of Megaupload usage. The logic is that countries with a higher pre-shutdown penetration of Megaupload usage, have more "treated" users than other countries do, and therefore should have a higher (relative) increase in sales compared than those other countries if the shutdown had an impact.

Indeed, the authors find that this is the case. Their results suggest that a 1% reduction in Megaupload usage within a particular country causes a 2.5-3.8% increase in digital sales. They calculate that this increase translates into a 6-10% increase in revenues from digital movie sales and rentals for two major studios in the 18 weeks after the Megaupload shutdown.

Lauinger et al. (2013) have also examined the impact of the Megaupload shutdown on the overall availability of pirated content online. They find that while shutting down Megaupload may have temporarily reduced the availability of content, this content quickly appears on other sites. We note, however, that while this may be true, it doesn't necessarily mean that shutting down prominent sites will be ineffective when it comes to reducing the utility of using piracy relative to legitimate channels. It may be that shutting down prominent sites increases consumer search costs to find new sites enough that some consumers choose to use legitimate channels (with lower search costs) instead of investing the effort to find the content on one of a number of pirate sites.

This conjecture is loosely consistent with evidence in the economics literature regarding price competition and consumer search costs online. This literature has shown that consumers are willing to pay several dollars more to purchase homogenous goods (typically books, CDs, and movies) from heavily branded sites such as Amazon.com, even when those products are available for lower prices from other sites (see for example Brynjolfsson and Smith (2000) and Smith and Brynjolfsson (2001)). Smith, Bailey, and Brynjolfsson (2000) argue that the price advantages enjoyed by these sites, can be explained as rational responses of consumers to the high search costs associated with finding other lower priced sellers online, high cognitive costs associated with learning how to use these new sites, or other preferences for convenience, quality, and reliability offered by heavily-branded, well-known sites. These results are similar to other results in the literature showing that consumers face high search costs when search for or processing information in online markets (see for example Brynjolfsson, Dick, and Smith (2010), Johnson et al. (2004), and Hann and Terwiesch (2003)).

In addition to these papers showing that regulatory interventions (whether supply- or demand-side) can increase the sales of media products, Bhattacharjee et al. (2006) analyze how voluntary demand-side interventions — specifically, the RIAA's lawsuits against individuals in 2003 and 2004 — impacted sales. The authors find that these lawsuits reduced the number of files individuals were sharing on P2P networks, and had a disproportionately large impact on large file sharers.

Thus, we believe that when taken as a whole, the academic literature shows that demand- and supply-side interventions can be effective in increasing media sales by reducing the utility of consuming pirated content relative to the utility of legitimate content.

However, we also note that while these results should be encouraging to decision-makers in the media industries looking for ways to increase revenue in the presence of piracy, the extant research only analyzes the benefit of these policies to legal sales, they do not analyze the net social benefit of these policies after taking into account the costs — both direct and indirect (e.g., costs arising from privacy, human rights, or goodwill costs) — associated with implementing these interventions. These costs are difficult to measure but will be important to evaluating the overall benefits from any anti-piracy intervention.

3.2 The Effectiveness of Market-Based Efforts to Reduce Piracy

The results reported above suggest that anti-piracy regulation can reduce the utility of consuming pirated content relative to content obtained through legitimate digital channels, and in so doing, can cause people who otherwise would have pirated to consume through legitimate channels instead. In this section, we review research results that suggest that there are a variety of ways that media firms can impact the marketing of their products to improve the attractiveness of their products relative to pirated consumption.

Perhaps the most clear manner in which firms can compete with piracy is to sell their content through a variety of digital distribution channels, allowing consumers to obtain the products they are looking for with as much or greater ease as they can pirate them. And there is empirical evidence that the use of legitimate digital distribution channels has a meaningful impact on piracy levels. For example, Danaher et. al. (2010) examined the

impact of NBC's decision to remove their video content from the iTunes store on piracy levels and on DVD sales levels through Amazon.com.

By comparing piracy rates of NBC content to piracy of similar ABC, CBS, and Fox content (before and after the removal of NBC from iTunes), the authors found a statistically and economically significant increase in piracy of NBC content caused by the removal of the legal digital distribution channel. In fact, the increase in the number of pirated NBC downloads was larger than the size of the iTunes market had been, suggesting that when consumers pay the fixed cost of switching to piracy they consume more content than they would purchase. This interpretation was validated by a second experiment — when NBC returned their content to iTunes nine months later, piracy of that content decreased by a much smaller amount than it had increased after the removal. Sales of physical DVD's of this content at Amazon exhibited no meaningful changes during either of these experiments.[6] Thus the authors make two conclusions: First, the opening of a digital distribution channel can encourage filesharers to migrate to legal channels. Second, consumers may be somewhat tied to physical vs. digital channels, as there seems to be a short run substitution between piracy and paid legal downloads but not between paid legal downloads and physical purchases.[7]

Another issue that impacts the consumer's decision to pirate is the timing of content availability in legitimate versus pirated channels. Many firms adopt a sequential release strategy across channels. For example, movie studios commonly release their content first in theaters, and later on DVD, digital, or television platforms. Likewise, music labels commonly delay the availability of the content on streaming services relative to CD or iTunes release dates. These business practices developed as way to maximize revenue from consumers by exploiting the (negative) correlation between consumers' willingness to pay

[6] We note that this conclusion is generally consistent with Waldfogel (2009) who finds that increased viewing of YouTube content did not significantly reduce college student's television viewing, and Hu and Smith (2013) who find that delayed available of Kindle eBooks did not significantly increase overall demand for print books.
[7] The authors note, however, that in the long run it is likely that digital sales cannibalize physical ones. The managerial point here is not that digital channels will not affect sales in physical channels, but rather that refusing to use digital channels for particular pieces of content is unlikely to forestall this cannibalization in the presence of digital piracy channels.

for content and their willingness to wait for that content to become available in a particular channel. However, the recent academic literature suggests that delayed release windows can also encourage filesharing. As such, it may be time for firms to revisit the optimality of these delayed release windows given the presence of pirated content.

Similarly, firms may wish to revisit strategies that rely on sequential release strategies across countries. For example, Danaher and Waldfogel (2012) document that movie studios have historically delayed box office releases of films across countries for several reasons, including the high cost of film prints and the desire to have the movies' stars present to promote the premiere. But the authors also note that, after controlling for most variables that impact a film's returns, longer lag's between the world premiere of a film and a country's premiere are correlated with lower returns for that film in that country. The authors find that this correlation grew more negative after the widespread adoption of movie filesharing protocols such as BitTorrent, and that correlation grew more negative for highly-pirates genres like science fiction than for less pirated genres. They show that an increasingly large amount of international piracy of films occurs in the window between the U.S. release and foreign release, and conclude that shortening these release windows could result in at least 7% higher returns due to reductions in this pre-release piracy.

In summary, the academic literature suggests that, in addition to various anti-piracy strategies, there are many market-based strategies that content creators can use in order to make their products more appealing relative to filesharing. Specifically, the two studies highlighted above demonstrate that the ease of access to content through legitimate channels, and the timing of legitimate channels relative to pirated availability are important components of the consumer's decision whether to pirate or purchase. But this alone does not mean that the burden of competing with piracy should fall entirely on the shoulders of producers. For example, producers could make their products vastly more appealing by setting the price to zero, but few would suggest that this solution is ideal (in such a scenario, as discussed above it seems difficult to imagine that the quality and quantity supplied would not be affected). Rather, it seems as if a combined approach could be implemented, with government policies acting to decrease the appeal of illegal

filesharing (effectively raising the "price" of piracy) while producers simultaneously attempt to deliver products to consumers with appealing timing, convenience, and quality.

4. Discussion

Digital channels create new opportunities for the copyright industries in the form of new channels to reach consumers and lower costs associated with distributing content digitally. However, digital channels can also create a parallel set of challenges for the copyright industries and potentially for society when these same digital channels and cost advantages are used by digital pirates to provide nearly perfect copies of content through piracy channels. The goal of this chapter was to outline the findings in the academic literature regarding the impact of piracy on sales and the effectiveness of both market-based and policy-based approaches to reducing the impact of piracy, and to raise important policy issues that should be addressed in future academic research.

Our analysis concludes that the vast majority of the academic literature that has looked at this question finds that piracy results in a statistically significant reduction in sales, particularly in emerging digital channels. We also note that while there are a few academic studies analyzing the broader social impact of piracy, standard economic theory suggests that if piracy reduces revenue to content creators that, *ceteris paribus*, it will also reduce the incentives to create new high quality content, and that any reduction in the supply of creative content could significantly impact overall social welfare.

Finally, our review of the literature suggests that content creators and governments retain many tools to reduce the impact of piracy on sales. Specifically, we argue that content creators can reduce the impact of piracy on sales by treating piracy as a competing good, and differentiating legal content from pirated content along a set of attributes that are attractive to the consumer such as convenience, reliability, and usability. However, we also argue that government and industry anti-piracy interventions can also serve as effective in increasing sales in legal channels by reducing the convenience, reliability, and usability of pirated content relative to content offered in legal channels.

References:

Adermon, A., C-Y Liang. 2010. Piracy, Music, and Movies: A Natural Experiment. Working Paper, Uppsala University, Uppsala, Sweden.

Andersen, B., M. Frenz. 2010. Don't blame the P2P file-sharers: the impact of free music downloads on the purchase of music CDs in Canada. *Journal of Evolutionary Economics*, 20:715-740.

Bai, J., J. Waldfogel. 2012. Movie Piracy and Sales Displacement in Two Samples of Chinese Consumers. *Information Economics and Policy*, **24**(3) 187-196.

Barker, G., T. Maloney. 2012. The Impact of Free Music Downloads on the Purchase of Music CDs in Canada. Working Paper Number 4-2012, Australian National University College of Law.

Bhattacharjee, S., R. Gopal, K. Lertwachara, J. Marsden. 2006. Impact of Legal Threats on Online Music Sharing Activity: An Analysis of Music Industry Legal Actions. *Journal of Law and Economics,* **49** 91–114.

Bhattacharjee, S., R. Gopal, K. Lertwachara, J. Marsden, R. Telang. 2007. The Effect of Digital Sharing Technologies on Music Markets: A Survival Analysis of Albums on Ranking Charts. *Management Science* **53**(9) 1359-1374.

Bilton, Nick. 2012. Internet Pirates Will Always Win. *New York Times*, Sunday Review, August 4. (Available from http://www.nytimes.com/2012/08/05/sunday-review/internet-pirates-will-always-win.html)

Boorstin, E. 2004. Music Sales in the Age of File Sharing" Senior Thesis, Princeton University, April 2004.

Bounie, D., M. Bourreau, P. Waelbroeck. 2006. Piracy and the Demand for Films: Analysis of Piracy Behavior in French Universities. *Review of Economic Research on Copyright Issues*, **3**(2) 15-27.

Brynjolfsson, E., M. Smith. 2000. Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Management Science*, **46**(4) 563-585.

Brynjolfsson, E., A. A. Dick, M. D. Smith. 2010. A Nearly Perfect Market? Differentiation Versus Price in Consumer Choice. *Quantitative Marketing and Economics*, **8**(1) 1-33.

Chellappa, R. K., Shivendu, S. 2005. Managing Piracy: Pricing and Sampling Strategies for Digital Experience Goods in Vertically Segmented Markets. *Information Systems Research,* 16(4) 400-417.

Conner, K.R., R.P. Rumelt. 1991. Software Piracy— an Analysis of Protection Strategies. *Management Science* **37** 125–139.

Danaher, B., S. Dhanasobhon, M.D. Smith, R. Telang. 2010. Converting Pirates without Cannibalizing Purchasers: The Impact of Digital Distribution on Physical Sales and Internet Piracy. *Marketing Science*, **29**(6) 1138-1151.

Danaher, B., S. Dhanasobhon, M. D. Smith, R. Telang. 2013. Understanding Media Markets in the Digital Age: Economics and Methodology. Working Paper, Carnegie Mellon University, Pittsburgh, PA.

Danaher, B., M. D. Smith. 2013. Gone in 60 Seconds: The Impact of the Megaupload Shutdown on Movie Sales. Working Paper, Carnegie Mellon University, Pittsburgh, PA. (Available from http://ssrn.com/abstract=2229349).

Danaher, B., M.D. Smith, R. Telang, S. Chen. Forthcoming. The Effect of Graduated Response Anti-Piracy Laws on Music Sales: Evidence from an Event Study in France. *Journal of Industrial Economics*, Forthcoming.

Danaher, B., J. Waldfogel. 2012. Reel Piracy: The Effect of Online Movie Piracy on Film Box Office Sales. Working Paper. Wellesley College. Boston, Massachusetts.

De Vany, A.S., W.D. Walls. 2010. Estimating the Effects of Movie Piracy on Box-office Revenue. *Review of Industrial Organization* 30:291-301.

Hammond, R.G. 2012. Profit Leak? Pre-Release File Sharing and the Music Industry. Working Paper, North Carolina State University, Raleigh, North Carolina.

Hann, I., Terwiesch, C. 2003. Measuring the frictional cost of online transactions: The case of a name-your-own-price channel. *Management Science,* **49**, 1563–1579.

Helberger, N., A. Huygen, N. Van Eijk. 2012. Ups and Downs: Economic and Cultural Effects of File Sharing on Music, Film and Games. Working Paper, University of Amsterdam, Amsterdam, Netherlands. (Available from http://ssrn.com/abstract=2009969)

Hennig-Thurau, T., V. Henning, H. Sattler. 2007. Consumer File Sharing of Motion Pictures. *Journal of Marketing*. 71(October) 1-18.

Hong, S., 2004. The Effect of Digital Technology on the Sales of Copyrighted Goods: Evidence from Napster. Stanford University, Mimeo.

Hu, Y., M. D. Smith. 2012. The Impact of eBook Distribution on Print Sales: Analysis of a Natural Experiment. Working Paper, Carnegie Mellon University, Pittsburgh, PA. (Available from http://ssrn.com/abstract=1966115)

Hui, K., I. Png. 2003. Piracy and the Legitimate Demand for Recorded Music. *Contributions to Economic Analysis & Policy,* **2**(1) Article 11.

Johnson, E. J., Moe, W. W., Fader, P. S., Bellman, S., & Lohse, J. 2004. On the depth and dynamics of world wide web shopping behavior. *Management Science, 50*, 299–308.

Liebowitz, S. 1985. Copying and Indirect Appropriability: Photocopying of Journals. *Journal of Political Economy*, **93**(5) 945-957.

Liebowitz, S. 2005. Pitfalls in Measuring the Impact of File-Sharing. *CESifo Economic Studies*, **51**(2-3) 435-473.

Liebowitz, S. 2007. How Reliable is the Oberholzer-Gee and Strumpf Paper on File-Sharing? Working Paper (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1014399).

Liebowitz, S. 2008a. File-Sharing: Creative Destruction or just Plain Destruction? *Journal of Law and Economics,* **49** 1-28.

Liebowitz, S. 2008b. Testing File-Sharing's Impact by Examining Record Sales in Cities. *Management Science,* **54**(4) 852-859.

Liebowitz, S. 2010. The Oberholzer-Gee/Strumpf File-Sharing Instrument Fails the Laugh Test. Working Paper. University of Texas at Dallas, Dallas, Texas.

Liebowitz, S., A. Zentner. 2012. Clash of the Titans: Does Internet Use Reduce Television Viewing? *Review of Economics and Statistics*, **94**(1) 234-245.

Lauinger, T., M. Szydlowski, K. Onarlioglu, G. Wondracek, E. Kirda, C. Kruegel. 2013. Clickonomics: Determining the Effect of Anti-Piracy Measures for One-Click Hosting. Working Paper, Northeastern University, Boston, Massachusetts.

Ma, L., A. Montgomery, P. Singh, M.D. Smith. 2013. Pre-Release Movie Piracy and Box Office Sales: Estimates and Policy Implications. Working Paper, Carnegie Mellon University, Pittsburgh, Pennsylvania. (Available from http://ssrn.com/abstract=1782924)

Martikainen, E. 2011. Does File-Sharing Reduce DVD Sales? Working Paper, University of Turku. Turku, Finland.

Michel, N. 2006. The impact of digital file sharing on the music industry: an empirical analysis. *Berkeley Electronic Press: Topics in Economic Analysis & Policy*, 6, Article 18.

Mortimer, J. H., A. Sorensen, C. Nosko. 2012. Supply Responses to Digital Distribution: Recorded Music and Live Performances. *Information Economics and Policy*, **24**(1) 3-14.

Oberholzer, F., K. Strumpf. 2007. The Effect of File Sharing on Record Sales. An Empirical Analysis. *Journal of Political Economy,* **115**(1) 1-42.

Peitz, M., P. Waelbroeck. 2004. The Effect of Internet Piracy on CD Sales: Cross-Section Evidence. *Review of Economic Research on Copyright Issues*, **1**(2) 71-79.

Rob, R., J. Waldfogel. 2006. Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students. *Journal of Law and Economics*, **49**(1) 29-62.

Rob, R., J. Waldfogel. 2007. Piracy on the Silver Screen. *Journal of Industrial Economics*, **55**(3) 379-393.

Smith, M.D., J. Bailey, E. Brynjolfsson. 2000. Understanding Digital Markets: Review and Assessment, Brynjolfsson and Kahin, eds. *Understanding the Digital Economy*, MIT Press, Cambridge, MA, 99-136.

Smith, M.D., E. Brynjolfsson. 2001. Customer Decision Making at an Internet Shopbot: Brand Still Matters. *The Journal of Industrial Economics*, **49**(4) 541-558.

Smith, M.D., R. Telang. 2009. Competing with Free: The Impact of Movie Broadcasts on DVD Sales and Internet Piracy. *Management Information Systems Quarterly*, **33**(2) 312-338.

Smith, M.D., R. Telang. 2010. Piracy or Promotion? The Impact of Broadband Internet Penetration on DVD Sales. *Information Economics and Policy*, Special Issue on the Economics of Digital Piracy, **21** 289-298.

Smith, M.D., R. Telang. 2013. The Impact of Movie Release Windows on Sales and Piracy. Working Paper, Carnegie Mellon University, Pittsburgh, PA.

Takeyama, L.N. 1994. The Welfare Implications of Unauthorized Re- production of Intellectual Property in the Presence of Demand Network Externalities. *Journal of Industrial Economics,* **42** 155–66.

Waldfogel, J. 2009. Lost on the web: Does web distribution stimulate or depress television viewing? *Information Economics and Policy,* **21**(2) 158–168.

Waldfogel, J. 2010. Music File Sharing and Sales Displacement in the iTunes Era. *Information Economics and Policy*, **22**(4) 306-314.

Waldfogel, J. 2011. Bye, Bye, Miss American Pie? The Supply of New Recorded Music Since Napster. NBER Working Paper No. 16882.

Zentner, A. 2006. Measuring the Effect of File Sharing on Music Purchases. *The Journal of Law and Economics,* 49 63–90.

Zentner, A. 2009. Ten Years of File Sharing and Its Effect on International Physical and Digital Music Sales. Working Paper, University of Texas at Dallas, Dallas, Texas. (Available from http://ssrn.com/abstract=1724444)

Zentner, A. 2012. Measuring the Impact of File Sharing on the Movie Industry: An Empirical Analysis Using a Panel of Countries. Working Paper, University of Texas at Dallas, Dallas, Texas. (Available from http://ssrn.com/abstract=1792615)