Piracy and Music Sales: The Effects of an Anti-Piracy Law[*]

Adrian Adermon[§] and Che-Yuan Liang[#]

April 23, 2014

Abstract: The implementation of a copyright protection reform in Sweden in April 2009 suddenly increased the risk of being caught and punished for illegal file sharing. This paper investigates the impact of the reform on illegal file sharing and music sales using a difference-in-differences approach with Norway and Finland as control groups. We find that the reform decreased Internet traffic by 16% and increased music sales by 36% during the first six months. Pirated music therefore seems to be a strong substitute to legal music. However, the reform effects disappeared almost completely after six months, likely because of the weak enforcement of the law.

Keywords: copyright protection, anti-piracy law, piracy, file sharing, music sales

JEL classification: D04, D12, K11, K42, M48

[*] We thank Orley Ashenfelter, Niklas Bengtsson, Mikael Elinder, Jon Gemus, Christine Greenhalgh, John Ham, Henrik Jordahl, Tatsiana Ksionda, David Lee, Mikael Lindahl, Erik Lindqvist, Johan Lyhagen, Alexandre Mas, Håkan Selin, two anonymous reviewers, and seminar participants at Uppsala University, the Research Institute of Industrial Economics, the Industrial Relations Section at Princeton University, the Swedish Institute for Social Research at Stockholm University, and the Stockholm-Uppsala Doctoral Students Workshop in Economics, the Institute for Labor Market Policy Evaluation Workshop in Public and Labor Economics at Öregrund, and the National Conference of Swedish Economists for valuable comments and suggestions. We also thank Martin Warberg, Peter Lindqvist, and Tatsiana Ksionda at Growth from Knowledge (GfK) Sweden, Christian Unsgaard at GfK Norway, NetNod, the European Internet Exchange Association (Euro-IX), the Norwegian Internet Exchange (NIX), the Finnish Communication and Internet Exchange (FICIX), IFPI, IFPI Sweden, IFPI Norway, IFPI Finland, GfK Sweden, and GfK Norway for providing the data used in the paper; we thank Aron Berg, Serge Radovcic at Euro-IX and Jorma Mellin at FICIX for help with other various aspects of the data collection. The Jan Wallander and Tom Hedelius Foundation and the Uppsala Center for Fiscal Studies are acknowledged for financial support.
[§] Uppsala Center for Labor Studies, Department of Economics, Uppsala University, P.O. Box 513, SE-75120, Uppsala, Sweden; Phone: +46-(0)18-4717638; Fax: +46-(0)18-4711478; E-mail: adrian.adermon@nek.uu.se.
[#] Uppsala Center for Fiscal Studies, Department of Economics, Uppsala University, P.O. Box 513, SE-75120, Uppsala, Sweden; Research Institute of Industrial Economics, P.O. Box 55665, SE-10215, Stockholm, Sweden; Phone: +46-(0)18-4711633; Fax: +46-(0)18-4711478; E-mail: che-yuan.liang@nek.uu.se (corresponding author).

1. Introduction

Copyright protection of information goods is a topic that has been debated at least since the invention of the printing press in the 15th century. The non-rival nature of information goods means that there is a tradeoff between proper incentives to creators and public benefits from the wider distribution of works (see, e.g., Plant, 1934, Hurt and Schuchman, 1966, Boldrin and Levine, 2002, and Varian, 2005). The level of intellectual property protection has important welfare implications. Inventions such as the photocopier, CD burners, and the Internet have made the copying of books, music, and movies inexpensive and easy and the enforcement of copyright more difficult. It could be argued that music in practice has become a public good. In view of these developments, guidelines for and evaluation of intellectual property rights policies are more important than ever.

During the last decade, the debate has focused on the recorded music industry. Global sales of recorded music amounted to over USD 25 billion in 1999 at the trade level after a steady and continuous rise in the 1990s.[1] In June 1999, however, the online music file sharing service Napster popularized file sharing to a wide audience. Since then global sales fell by 30% between 1999 and 2009. Interest groups for the industry were quick to blame illegal file sharing of copyright-protected material, commonly referred to as "piracy". By now, there are a considerable number of papers supporting the view that piracy displaces legal music sales.[2] Only a few papers have, however, investigated the effects of antipiracy interventions (Bhattacharjee et al., 2006, Danaher and Smith, 2014, and Danaher et al., 2013b).

On April 1, 2009, the intellectual property rights law IPRED was implemented in Sweden. This law made it easier for property right holders to identify file-sharers by requesting their identity from Internet Service providers and therefore increased the perceived risk of getting caught and punished for piracy. The piracy issue attracted wide media coverage around this time and overnight Internet traffic decreased by 40%. Since the law did not concern any Internet activities unrelated to piracy, this drop must be attributed to decreased piracy. At a glance, the reform therefore seems to be a promising measure in combating piracy. We study the effects of this intervention on Internet traffic and music sales up until nine months after the reform. Besides providing an evaluation of the potential effectiveness of anti-piracy interventions, the reduction in Internet traffic also provides an unprecedented large experimental-like variation in piracy which is an opportunity for studying the causal effect of piracy on music sales.

We use a difference-in-differences strategy with Norway and Finland as control groups to estimate the reform effects on piracy and music sales. We find that the reform decreased Internet traffic by 16% during the first six months. Under the assumption that half of Internet traffic consists of piracy (Schultze and Mochalski, 2009), this decrease corresponds to a decrease in piracy of 32%. However, we also find that Internet traffic almost completely recovered after six months. The fact that there were only a few legal cases that stayed in court for a long time (the first court case to finish reached a final verdict 4 years later) could have

[1] International Federation of the Phonographic Industry (IFPI).

[2] See Liebowitz (2006), Dejean (2009), and Danaher et al. (2013a) for overviews of the literature. Oberholzer-Gee and Strumpf (2007) and Andersen and Frenz (2010) are the two exceptions finding no effects.

hampered the long-run effectiveness of the reform. During the first six months of the reform, it increased physical music sales by 33%, digital music sales by 46%, and total music sales by 36%.

Our estimates imply a music sales elasticity of piracy of approximately one on the margin and a large causal effect of piracy on music sales. Under the strong functional form assumption that the marginal effect is constant, music sales would have been twice as large in 2009 in the absence of piracy. Piracy could then account for 80% of the drop in music sales between 2000 and 2008, which would support the music industry's claim that piracy was the main cause of the decline. Our results are corroborated by survey-based evidence on individual file sharing and music consumption behavior in Sweden prior to and after the reform.

Our paper is closely related to a paper by Danaher et al. (2013b) who study the impact on digital iTunes sales of the introduction of the HADOPI anti-piracy law in France in September 2010 and the legislative process back until 1.5 years before the introduction. HADOPI is similar to IPRED by allowing rights holders to identify infringers, but HAPODI is more ambitious as it is backed up by an administrative authority handling the cases. On the other hand, it is softer as two warnings are handed out before any sanctions are imposed. They find that the reform increased iTunes sales by around 25%. Unlike our paper, they neither study the first stage effect on piracy nor the reform effect on physical sales.

The papers by Bhattacharjee et al. (2006, 2007) and Blackburn (2006) are other closely related papers. Rather than studying the effects of government interventions, they study the effects of the industry-driven interventions by the Recording Industry Association of America (RIAA) in 2004, when they announced a lawsuit strategy against illegal file sharers. Bhattacharjee et al. find that the reform decreased the number of available and downloaded files although the overall availability of files continued to be large. Additionally, they find reduced survival times on the charts for less popular albums. Blackburn also finds that the reform reduced the availability of files. However, he finds a negative effect on sales of popular albums and a positive effect on sales of less popular albums. Unlike our paper, these two papers do not use any control groups.

Our paper is also related to the literature studying the effect of piracy on music sales. To the extent that duplicated material available through file sharing is a substitute to legal purchases, economic theory predicts a negative relationship. Theoretical research, however, shows that there are potential mechanisms that may induce an opposite effect, i.e., the sampling effect, where piracy may provide a sampling opportunity for consumers and raise the willingness to pay for goods that they know match their preferences better (Takeyama, 1994, and Duchene and Waelbroeck, 2006). The direction of the relationship is therefore an open empirical issue. The welfare implications of piracy and potential welfare gains of effective anti-piracy interventions depend on the size of the displacement effect because incentives faced by creators depend on music sales.

Estimating the causal effect of piracy on sales is, however, challenging. Observed correlation between piracy and sales may stem from variation in sources such as Internet activity unrelated to piracy, cultural trends, demographics, or macroeconomics. Some papers try to isolate random variation in piracy by exploiting an event (a natural or quasi-experiment) that affects the demand or supply of files. Events such as anti-piracy interventions (Blackburn, 2006, and Bhattacharjee, 2006), the introduction of the first file-sharing service Napster (Peitz and Waelbroeck, 2004, Zentner, 2005, Michel, 2006, Bhattacharjee et al. 2007, Liebowitz, 2008, and Hong, 2013), and German school holidays affecting the availability of files in the U.S. as German kids spend more time on their computers (Oberholzer-Gee and Strumpf, 2007) have been exploited. Another related approach (that sometimes is combined with the former approach) is to exploit variation in file sharing capacity or changes in file-sharing capacity between individuals or groups of individuals. Internet sophistication and connection speed (Zentner, 2005, 2006, and Rob and Waldfogel, 2006) are typical measures of file-sharing capacity that have been exploited.

The literature can also be categorized by the level of data aggregation used. Several papers use individual-level (typically survey) data.[3] Piracy by some individuals may, however, affect music consumption by other individuals, e.g., because friends play music to each other. Other difficulties include measurement errors on infrequent purchases and reliability issues for survey responses about illegal behavior. Another group of papers use city- or country-level Internet penetration measures as proxies for piracy in combination with aggregate sales.[4] Variation in Internet penetration does, however, not fully capture variation in Internet traffic intensity. A third group of papers use estimated downloads at the album level from one file sharing network to measure piracy in combination with album-level sales.[5] A single file sharing network and a selection of albums might, however, not be representative. It is also impossible to identify complete downloads for technical reasons and because illegal file sharers want to hide their activities. Furthermore, piracy of specific albums could affect sales of other albums, e.g., because listening to an album in a certain genre may affect a user's tastes regarding other albums in that genre.

Our paper is the first to use a measure of aggregate Internet traffic in a country to capture aggregate piracy. Combining this with aggregate sales allows better estimation of the full effect, which accounts for any cross effects between different individuals and different albums. Compared with the papers studying the effect of piracy on music sales, our natural experiment approach is also the first to incorporate genuinely untreated control groups.

The paper proceeds as follows. The next section provides a background on file sharing and the reform. Section three describes the data and graphically investigates the development of Internet traffic and music sales. That section also provides an informal and nontechnical analysis that discloses the main results in an intuitive way. Section four outlines in detail our empirical strategy. Section five reports the regression estimates. The last section concludes.

[3] For instance, Bounie et al. (2005), Michel (2006), Rob and Waldfogel (2006), Zentner (2006), Bhattacharjee et al. (2006), Andersen and Frenz (2010), and Hong (2013).
[4] For instance, Peitz and Waelbroeck (2004), Zentner (2005), and Liebowitz (2008).
[5] For instance, Oberholzer-Gee and Strumpf (2007), Blackburn (2006), Bhattacharjee et al. (2007), and Hammond (2013).

2. Background

2.1 File sharing

Although people have been copying music illegally since the 1960s using analog tapes, digital piracy (which allows easy reproduction without a loss of quality) became feasible in the mid to late 1990s through the availability of affordable CD burners and the Internet in conjunction with the mp3 file format for compressing digital music. Piracy exploded when Napster was introduced in June 1999 and allowed individuals to easily search for and download music files indexed on a central server. Although Napster was shut down in 2001 after a legal battle with the recording industry, it was followed by many similar services, e.g., Gnutella, Kazaa, and Grokster. These modern services use peer-to-peer (P2P) technology in which users download files directly from each other rather than from a central server.

In 2001, the BitTorrent protocol was released. With previous P2P protocols, popular files were sometimes hard to get because users queued up to download from the suppliers. With BitTorrent, all users downloading a file both upload and download pieces of the file simultaneously in a decentralized fashion so that the more popular a file is, the faster the download becomes. In 2002, the first major BitTorrent tracker, Suprnova, was started, and the Sweden-based site The Pirate Bay followed in 2003, becoming the largest BitTorrent site in terms of traffic in 2008 (Alexa) when it reached 25 million unique peers (The Pirate Bay). The speed and availability provided by the BitTorrent protocol made it the most popular protocol used by P2P networks in the period we study (Schultze and Mochalski, 2007, 2009).

During the last decade, the recorded music and movie industries ramped up their legal battle against file sharing networks by shutting down several large tracker sites from 2004 onwards. In Sweden, the piracy issue has received great attention since the police raid against The Pirate Bay on March 31, 2006, wherein servers were confiscated. Law suits and trials followed with hearings starting on February 16, 2009. Four site operators were found guilty on April 17, 2009, less than a month after the law reform that we use, but the site continues to operate.

Sweden is one of the most developed countries in the world in terms of Internet penetration. It is among the top ten with respect to density of broadband subscribers, percentage fiber among broadband connections, and broadband access density among households. Sweden is therefore of interest for other countries because it is on the IT frontlines and provides a forecast for countries less developed in this respect.

2.2 The IPRED law

The Intellectual Property Rights Enforcement Directive 2004/48/EC (IPRED) is a European Union directive that was passed on April 29, 2004 under article 95 of the Treaty of Rome. The directive regulates the enforcement of intellectual property rights, especially concerning evidence, and was set to be implemented before April 29, 2006. However, by October 2006 only 12 of the 25 member states had implemented the directive (Ipeg, 2006). The Swedish government declared its intention to implement the IPRED directive as Swedish law[6] in

[6] Adelsohn Liljeroth, Lena, “Kulturarbetare får stöd av domstol”, Svenska Dagbladet, March 14, 2008.

March 2008. The law was passed on February 25, 2009, and it was implemented on April 1 of the same year.

IPRED[7] allows an intellectual property rights holder to request the identity of the person behind a specific Internet Protocol (IP) address from an Internet Service Provider (ISP) if there is a reasonable suspicion that this person infringed on the intellectual property in question. The copyright holder must present evidence tying the infringement to a specific IP address in court, which may then order the ISP to provide the information necessary to identify the infringing individual. The court will only take this step if the individual has uploaded copyright protected material or has downloaded large volumes of copyright protected material. With BitTorrent technology, however, all users simultaneously upload and download data. Thus, most file sharers were theoretically affected by IPRED. Before the reform, the police handled piracy cases and prosecutions ended up in the criminal court. In practice, the reform made it easier for property rights holders or their representatives to gather evidence and sue offenders in civil court or to assist the police.

In mid to late April, several large Swedish ISPs[8] declared that they would start destroying their IP logs so that their customers would not have to be concerned about IPRED. The largest ISPs[9] have, however, not followed suit. So far, there have only been a few legal cases pleading the new law. On the day the new law was implemented, five audiobook publishers requested the identity of a person behind an IP address associated with piracy. In late July, four movie companies filed a motion for details about the operators of a BitTorrent tracker, and in early December the Swedish section of the International Federation of the Phonographic Industry (IFPI) filed to get information on an IP address that has been used for music piracy. The audiobook case was appealed several times, until the Supreme Court of Sweden, guided by a ruling from the European Court of Justice, finally ordered the ISP to hand over the relevant information on December 21, 2012. The BitTorrent tracker case also led to a ruling in favor of the copyright holders on January 3, 2013. Although the law seemed fairly harmless to individual file-sharers ex post, the perceived threat ex ante was large as the 40% reduction in Internet traffic overnight shows.

Although IPRED is not the first regulation of relevance for file sharing, it is much more forceful than any previous regulations in Sweden due to the relatively strong formulation. In connection to the legislative process, there was a wide and extensive public debate about whether the law was too invasive on privacy. Around the weeks of implementation, the piracy issue also received wide media coverage in Sweden and internationally because of the highly publicized trial of The Pirate Bay. Public opinion was in general against the law, and this culminated in the Swedish elections to the European parliament on June 7, 2009, where the single-issue Pirate Party received 7% of the total vote and a seat in the European parliament.

Besides the widespread public awareness, incidental evidence and public perception also suggest that IPRED had a strong deterring effect on piracy around the date of implementation. The institutional development, however, suggests some need for carefulness in the empirical analysis and in the interpretation of the results. Because the law was passed more than a month before the implementation, it is plausible that there might have been

[7] Henceforth IPRED will refer to the Swedish implementation of the directive.
[8] Bahnhof, Alltele, and Tele2.
[9] Telia Sonera, Bredbandsbolaget, and Com Hem.

anticipatory effects. File sharers may have tried to stock up on downloadable music and movies when anticipating the implementation. We have examined this hypothesis and find no support for anticipatory behavior in the data. Because the legal processes have been slow, it is plausible that people came to reevaluate the risk of getting caught, reconsidered their initial fears, and started to share files again a few months later. The reform effect may hence have been dynamic over time, and we allow for this in the analysis.

Laws similar to IPRED have been implemented in several other countries. Versions of a three-strikes graduated response system, where file sharers are given three warnings before any legal sanctions are used, have been proposed or implemented in France, the UK and New Zealand. The French HADOPI law, introduced in 2009, created a new government agency tasked with administrating such a system. After three strikes, the individual's Internet connection could be temporarily suspended.[10] In July 2013, the suspension of Internet connection was repealed and replaced with a monetary fine. At the same time, the government announced that focus would shift to fighting commercial piracy.[11] Under the UK Digital Economy Act 2010, individuals who receive three warnings during a 12-month period are put on a list that copyright holders can access in order to take them to court.[12] The law has been delayed, and is expected to be implemented in 2015.[13] New Zealand has already implemented a very similar law, the Copyright (Infringing File Sharing) Amendment Act 2011.[14]

Other proposed laws around the world include the international Anti-Counterfeiting Trade Agreement (ACTA), which would create a common legal framework for enforcing copyright protection but was rejected by the EU parliament in 2012, and the controversial US Stop Online Piracy Act (SOPA), which focuses on suspending websites that infringe copyright but has not yet been enacted.

While the three-strikes laws are similar to IPRED in making monitoring and legal action easier, their gradual response design contrasts with IPRED, which makes prosecution of one-time offenders possible without prior warning. Thus, IPRED is likely to have had a more immediate deterring effect on file-sharers, while three-strikes laws probably generate more gradual reactions. The SOPA and ACTA initiatives are quite different from IPRED in that they largely focus on going after large-scale copyright infringers. Because of its unique design, knowledge about the impact of the IPRED law could be useful for policy makers everywhere.

---

[10] "Government Bill promoting the dissemination and protection of creative works on the internet - Explanatory Memorandum", French Republic, Ministry of Culture and Communication. Available online: http://www.culture.gouv.fr/culture/actualites/conferen/albanel/creainterenglish.pdf.

[11] "Publication du décret supprimant la peine complémentaire de la suspension d'accès à Internet", French Republic, Ministry of Culture and Communication, July 9 2013. Available online: http://www.culturecommunication.gouv.fr/Espace-Presse/Communiques-de-presse/Publication-du-decret-supprimant-la-peine-complementaire-de-la-suspension-d-acces-a-Internet.

[12] "Online Infringement of Copyright and the Digital Economy Act 2010", Ofcom, 26 June 2012. Available online: http://stakeholders.ofcom.org.uk/binaries/consultations/online-notice/summary/notice.pdf.

[13] "UK piracy warning letters delayed until 2015", BBC News, 6 June 2013. Available online: http://www.bbc.co.uk/news/technology-22796723.

[14] "Copyright (Infringing File Sharing) Amendment Act 2011", New Zealand Legistlation, 18 April 2011. Available online: http://www.legislation.govt.nz/act/public/2011/0011/latest/whole.html#DLM3331813.

3. Data

3.1 Internet traffic

We use data on aggregate Internet traffic through Internet Exchange Points (IXPs). We capture changes in piracy by assuming that changes in Internet traffic after the reform (after accounting for time effects) can be attributed to changes in piracy alone. If piracy and legal online activities complement each other, e.g., because file-sharers spend additional time on online activities while downloading illegal content, we would overestimate the change in piracy. This would, however, not affect our estimates of the reform impact on music sales. Because the reform provides variation at the country level, we use data at this level to avoid complications in the estimation due to correlation between local units.

The Internet is, loosely speaking, a network of networks. An IXP connects the smaller networks to each other. For a more detailed description of how the Internet and IXPs work, see Appendix A. Traffic between two users on the same network does not, however, pass through IXPs. For our results to be representative for all traffic, it is necessary that the relative reform effect on piracy is the same for between- and within-network piracy. This is most likely to hold because BitTorrent, the most popular P2P protocol by far during this period, does not allow users to select where to download from. All downloads usually involve within- and between-network traffic, and users cannot select the share of each type of traffic to use. Furthermore, the law does not discriminate between different types of traffic. It is possible that different shares of within- and between-network traffic are piracy. This is, however, not a problem as long as the relative response in piracy due to the reform is the same.

It is possible that some file sharers simply started hiding their activities using technical solutions such as proxy servers or VPN networks[15] in response to the reform. A strength of our indirect measure of file sharing is that this traffic would still be picked up as Internet traffic. If anything, a switch to these technologies would make file sharing traffic more likely to pass between networks causing a downward bias of our estimate of the reform effect on piracy.

We have a weekly Internet traffic panel for Sweden, Norway, and Finland for 2009. The data are in weekly averages and are measured in gigabits per second (Gbps). The data come from the largest IXPs in Sweden, Norway, and Finland. The Swedish data are provided by Netnod, the Norwegian data by Norwegian Internet Exchange (NIX), and the Finnish data by Finnish Communication and Internet Exchange (FICIX). These IXPs provide the vast majority of between-network connections in their respective countries.

The development of Internet traffic is plotted in Figure 1. A solid line is used for Sweden, a dashed line for Norway, and a short-dashed line for Finland. A vertical line marks the last time point before the reform. Data series plots follow this structure throughout this section.

[15] For more details on these technologies, see Appendix A.

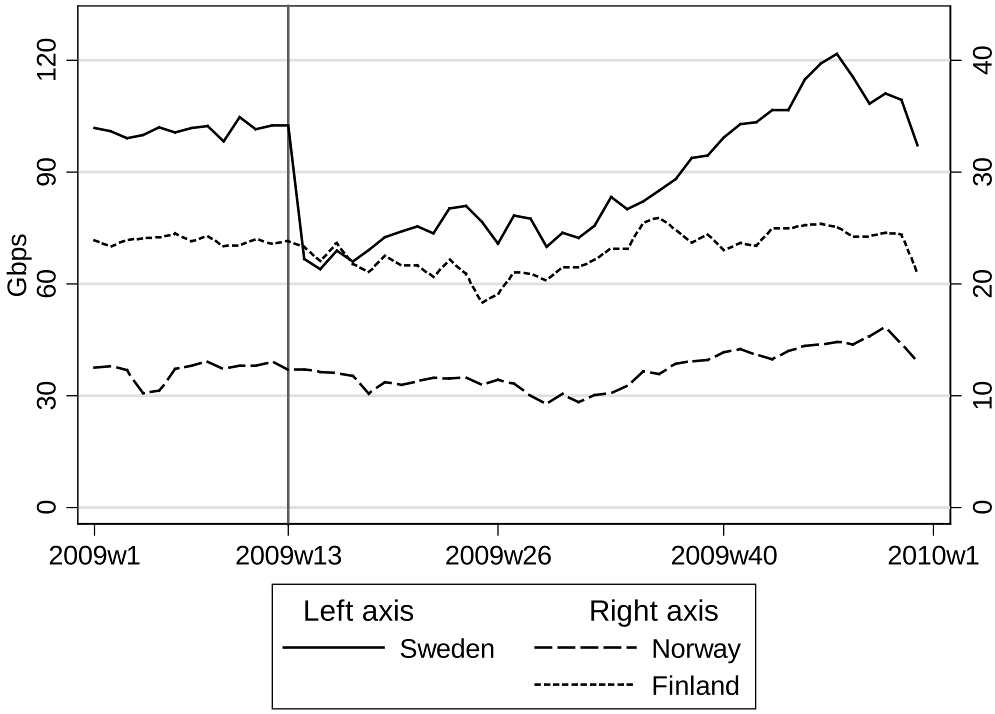

Figure 1. Internet traffic

Average Internet traffic in 2009 was 90.83 Gbps in Sweden, 12.29 Gbps in Norway, and 22.81 Gbps in Finland. We see that Swedish Internet traffic decreased sharply in the first week of April at the time of the reform.[16] Neither Norway nor Finland shows such a drop at this date. In fact, there are no other changes of this magnitude between adjacent points in time in our data. Swedish Internet traffic recovered toward the end of the year. Some of this increase seems to be common for all countries. There are some time effects such as trends and seasonal effects when disregarding the post-reform period in Sweden. The trends look reasonably parallel for all three countries in the pre-reform period and between Norway and Finland in the post-treatment period. There are no signs of an anticipation effect near the end of February when the law was passed.

3.2 Music sales

For aggregate music sales, we use monthly data from 2004 to 2009. From 2007, data split by physical and digital recorded music are available. Before 2007, digital sales were negligible. All sales numbers are at the trade level, and they are adjusted for inflation and converted to 2009 SEK.[17] The music data are provided by local branches of IFPI. IFPI members have a 95% market share in Sweden, 90 to 95% in Norway, and 86% in Finland. The development of sales at the quarterly level is plotted in Figure 2 and Figure 3.

[16] The drop was 40% on the reform day. Since the whole drop can be attributed to piracy, this share can be used as a lower bound on piracy's share of pre-reform Internet traffic.
[17] The average SEK/USD exchange rate was 7.6 in 2009.

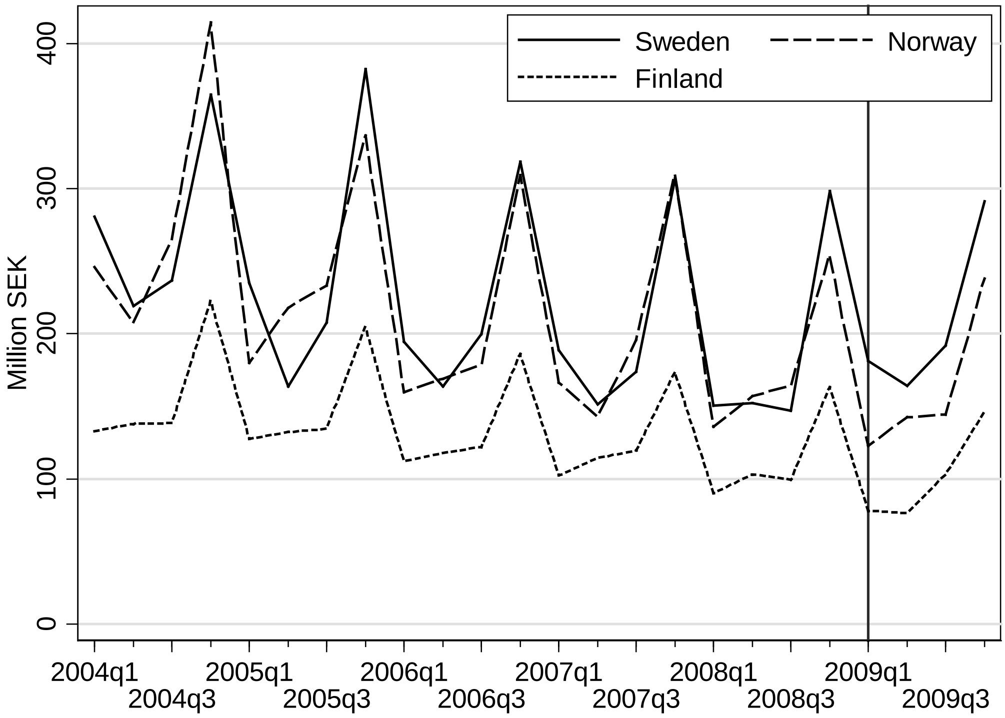

Figure 2. Total music sales

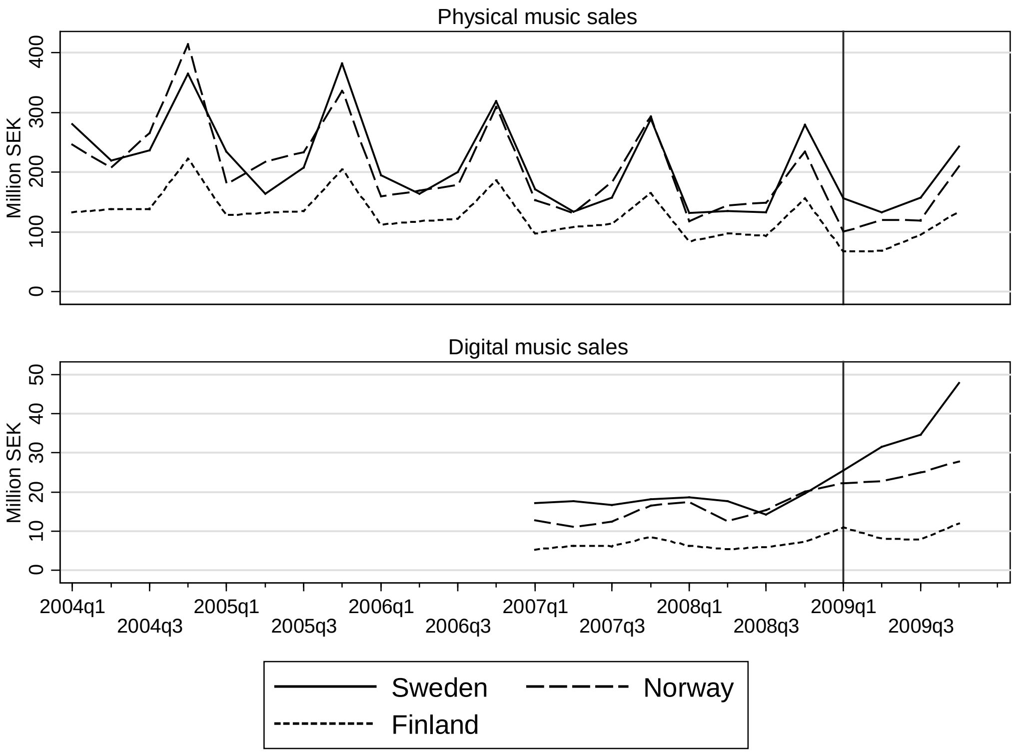

Figure 3. Music sales by format

Total sales were SEK 828 million in Sweden, SEK 648 million in Norway, and SEK 404 million in Finland in 2009. We see that the sales trends were negative. The seasonal patterns of second quarter troughs and last quarter peaks are also clear. Furthermore, the trend and seasonal patterns are, to some degree, country specific. The second quarter trough is, e.g., sharper in Sweden than in the other countries. Seasonal patterns may differ, e.g., if release patterns of new domestic albums are different between countries. In each of the first two quarters after the reform, sales increased compared to the same quarters in the previous year in Sweden, which is a pattern that cannot be found for any two consecutive quarters in the countries without a reform or in Sweden prior to the reform.

Physical sales[18] were SEK 689 million in Sweden, SEK 550 million in Norway, and SEK 365 million in Finland in 2009. Since physical sales make up most of total sales, the pattern is mostly similar to that discussed above. Digital sales[19] were SEK 139 million in Sweden, SEK 98 million in Norway, and SEK 39 million in Finland in 2009. We see that sales trends were positive and not unlikely country-specific. The sales increase in Sweden after the reform is, however, a clear trend break. The increase in the three post-reform quarters does not have a counterpart elsewhere in the data.

For physical sales, we also have sales by subcategories. We have a partition along price levels for Sweden and Finland, i.e., full-price, mid-price, and budget albums, and along origin of artist for all three countries, i.e., domestic and international artists.

4. Empirical Strategy

4.1 Difference-in-differences

We now turn our attention to quantifying the reform effects on the Internet traffic and music sales outcome variables, and indirectly the effects of piracy on music sales. The simplest method would be to compare the post-reform outcomes with the pre-reform outcomes in Sweden. However, if the timing of the implementation was based on factors related to Internet traffic or music sales, such factors may confound the reform estimates. We solve this by using Norway and Finland, two countries similar to Sweden in geographical, cultural, and technological aspects but without any reform at this time, as control groups to control for common trends and seasonal factors. Thus, we identify the reform effects using a difference-in-differences approach.

We estimate the reform effect on piracy and sales using the following specification with ordinary least squares (OLS):

$$\ln Outcome_{it} = \beta Reform_{it} + \alpha T_{it} + \eta_i + \sigma_t + \varepsilon_{it}, \qquad (1)$$

where i is a country index and t is a time period index. $Outcome_{it}$ is the outcome variable which may be Internet traffic or music sales variables, $Reform_{it}$ is a post-reform dummy vector, and T_{it} is a vector of country-specific seasonal dummies and linear trends. The variable η_i is a country fixed effect, σ_t is a time fixed effect, and ε_{it} is an idiosyncratic error

[18] Physical sales include albums and singles on all physical formats, e.g., CD, LP, DVD-Audio, SACD, and cassette.

[19] Digital sales include paid downloads of single tracks and albums, streams, mobile music services (including ringtones), and subscription services, such as Spotify.

term. From Figures 1 to 3, we observe that some reform effects seem to vary over time, and we allow for dynamic effects over time in most specifications. The effects are clear only up until six months after the reform, and in our main specification, we estimate the aggregate effect for the first six months, allowing the last three months to have a different effect. We also observe that some variables have country-specific seasonal patterns and trends, motivating the inclusion of T_{it}.[20] For causal inference, we formally require the reform to be uncorrelated with ε_{it} conditional on the other covariates.

We assume a log-linear model, which accounts for scale and level differences in outcome variables between the countries. The reduced-form coefficients from this model have a proportional interpretation at the margin. For discrete variables, the coefficients can be transformed by an exponential to provide a non-marginal relative reform effect with an implicit constant marginal effect assumption (Halvorsen and Palmquist, 1980).

To the extent that there may be spillover effects of the Swedish reform on its neighbors, the reform effect estimates would be downward biased. It is unlikely that the reform would directly influence file sharers in neighboring countries. However, the availability of files in a country may affect piracy in other countries because file sharing is a cross-national activity. Less piracy in Sweden can thus decrease the supply of pirated material in, e.g., Norway. This is, however, to some degree offset by traffic between Sweden and Norway being replaced by heavier traffic within Norway and between Norway and other countries. It is difficult to speculate on the degree of spillover effects. At least, there are no important immediate spillover effects at the reform date according to the development of Internet traffic in Figure 1. For sales, there is no reason to expect any spillover effects, except those working through piracy.

4.2 Specification and sensitivity tests

Causal inference in the typical difference-in-differences specification requires pre-reform trends to be parallel between treated and untreated groups. With many pre-reform periods available for music sales, one can check the assumption by placing and estimating placebo interventions in pre-reform periods, dropping the post-reform periods. These placebo estimates should then be lower than the main reform estimate and rarely be statistically significant. Because we include country-specific linear trends, we allow differential trends between countries. Nevertheless, there may remain higher-order country-specific trend effects or other seasonal effects. The placebo procedure, however, still works as an identification specification test. We iterate the placebo interventions 54 times, with the intervention starting in every pre-reform month from January 2004 to June 2008, and as in the main specification we estimate the effect for the first six months during which the law clearly was effective

[20] Because we do not want to let variation in outcomes in the post-reform period due to potential dynamic effects contribute to the identification of the country-specific variables, we proceed in two steps. First, we estimate the nuisance parameters for the control variables using all non-reform observations and adjust the outcome variables of all observations for those components. Second, we regress the adjusted outcome variables against the reform variable(s). We adjust the standard errors for using a two-step procedure as proposed by Lovell (1963). Another solution is to allow heterogeneous effects for every post-reform period. To obtain estimates for several periods, such as a single reform estimate, then requires weighing the estimates by the counterfactual outcomes in each period. This procedure was used in the working paper version of this paper (Adermon and Liang, 2010). The results are similar for these two procedures.

allowing for a different effect for the last three months. Because these placebo tests are estimated on a reduced data set with nine post-reform months fewer, we also re-estimate our main reform regressions on a sample where the first nine months are dropped to make the comparison fair. The specification test is then done by comparing the coefficient estimate from the actual reform with the distribution of the placebo estimates.

When long time-series are used, serial correlation in the error term may bias the standard errors downwards.[21] The standard method to account for serial correlation by clustering the standard errors does not work well when there are only three groups. An option is to use a parametric correction such as the Newey-West standard error correction (Newey and West, 1987), which requires specifying the length of the correlation over time. We have performed this correction allowing for correlations up to six months, and it turns out that those standard errors are very close to and usually smaller than the conventional OLS standard errors. Consequently, we report the more conservative uncorrected standard errors in the main results.[22]

In the presence of potentially biased standard errors, the placebo estimates can also be used to test the inference as discussed by Bertrand et al. (2004). In this test, we focus on the statistical significance of the placebo point estimates that should be statistically significant at the 5% level in 5% of the cases as a result of random fluctuations. A larger fraction of significant placebo estimates could indicate problems with type I errors.[23] Assuming that the placebo point estimates are correctly estimated and reflect the distribution of estimates when there is no reform effect, they can also be used for an alternative method to make inference following the approach by Abadie et al. (2010). The empirical statistical significance level of the reform effect being not zero can then be obtained as the probability of obtaining a more extreme reform point estimate from the distribution of placebo estimates.

An additional inference sensitivity test that we implement is Donald and Lang's (2007) two-step estimator for inference with few groups. This method allows for any within-group error correlation including serial correlation. As shown in the simulations in Bertrand et al. (2004) this procedure suffers from low power, but it is useful by providing very conservative inference, and thus provides an upper bound on p-values. In the first step, we regress the dependent variable on all control variables. The residuals are collapsed into country means separately for three time periods – before the reform, the first six months after the reform, and the last three months after the reform.[24] In the second step, we use the collapsed data to regress the dependent variable on a reform indicator and country and period fixed effects, including a separate last-quarter reform dummy. To test if our reform point estimates are statistically significantly different from zero, we use a t-distribution with $n - k = 1$ degree of

[21] Another closely related standard error clustering problem is that there may exist within-group correlation. This is known as the Moulton problem (Moulton, 1990). In our setting, this is not an issue, since we use group-level aggregate data.

[22] Angrist and Pischke (2008) recommend reporting the largest of the conventional and robust standard errors.

[23] Serial correlation makes it more likely to estimate an extreme and therefore statistically significant placebo and reform estimate for the intervention period because random noise tend to be clustered within the intervention period.

[24] The results are not sensitive to including the last three months into the treatment period or dropping the last three months. We focus on the three-period scheme in order to replicate the point estimates from our main specification.

freedom, n being number of observations, 9, and k being number of estimated parameters, 8.[25]

4.3 Implied effect of piracy on music sales

The reform effect estimate of Internet traffic $\beta_{Internet}$ provides a lower bound of the reform effect on piracy. To get an estimate of the reform effect on piracy β_{Piracy} we need to scale $\beta_{Internet}$ by piracy's share of all Internet traffic prior to the reform. We know of no estimates of this number for Sweden, but a report from the German Internet traffic analysis firm Ipoque (Schultze and Mochalski, 2009) provides estimates of the proportion of Internet traffic that consists of file sharing in several European countries. Estimates for the Nordic countries are not provided, but the share is 53% in Germany, 70% in Eastern Europe, 55% in Southern Europe, and 54% in Southwestern Europe. Although there are legal uses for P2P technology, the vast majority of all file sharing consists of piracy. Since Sweden is probably similar to Germany in this respect, we assume the share to be 0.5. The results can easily be rescaled by a constant for the reader that finds this assumption implausible.[26]

The reform effects are of intrinsic interest. Under the assumption that the reform had the same relative effect across media, it is also possible to infer the effect of piracy on music sales as $\beta_{Sales}/\beta_{Piracy}$. Estimating this quotient corresponds to an instrumental variables strategy where the reform is the instrument for piracy. β_{Sales} then corresponds to the reduced-form parameter and $\beta_{Internet}$ is the first-stage parameter. The identifying assumption is that the reform did not affect sales other than through its effect on piracy. There are several technical complications with this strategy and we only report the reduced-form estimates here as they are clear enough.[27]

Notice, however, that our natural experiment approach to infer the effect of piracy on music sales avoids endogeneity issues related to omitted variables and reverse causation. For instance, an observed positive correlation between piracy and music sales could reflect that music interest is high (leading to a lot of piracy and high sales), and a zero or negative correlation may be caused by a recession (leading to a lot of piracy but low sales). The IPRED reform provides an exogenous *ceteris paribus* variation in piracy that is as good as random. The identifying assumption is that the reform affects piracy (by making it seemingly riskier), but not music sales, except through piracy.

Behavioral responses to the reform are likely to be heterogeneous across individuals. Casual, small-scale file-sharers were probably more likely to be deterred by the reform than

[25] This approach requires the assumption that the second-step error term is homoscedastic. This holds if the number of observations per group is large. Because our data are aggregated on the country level, we implicitly have a large number of observations for each group. To see this, consider the case where we would have used individual-level data instead. The problem of serial correlation would not have changed, and absent any individual-level covariates, the two-step estimation procedure would have been the same.

[26] A hard lower bound would be 0.4, i.e., all piracy ceased after the reform. Conversely, a hard upper bound would be 1, i.e., all Internet traffic is piracy. We believe that these assumptions are both unreasonable and focus on a number that has at least indirect empirical support.

[27] Because of the need to rescale the first-stage coefficient and the fact that we use different samples and partially different control variables, it is not possible to use the standard two-stage least squares (2SLS) estimator. The two-sample instrumental variables procedure developed in Angrist (1990) and Angrist and Krueger (1992, 1995) is a possible route. This procedure was used in the working paper version of this paper (Adermon and Liang, 2010).

large-scale file sharing enthusiasts who have better knowledge of and access to technologies that make it harder to detect them. The implicit IV discussed in the previous paragraph would then be a local average treatment effect (LATE), with casual file sharers more likely being the compliers, while heavy users more likely are “never-takers” who do not reduce their activities in response to the reform. The casual users make up the vast majority of file sharers, and for them file sharing is probably a closer substitute to buying music than for heavy file sharers, since many of the heavy file sharers would not afford to buy all the music and movies they download illegally. For copyright protection reforms, our LATE is close to the most policy relevant parameter.[28]

5. Results

5.1 Internet traffic

The reform effect estimates for Internet traffic, estimated using Equation (1), are reported in Table 1. The quarterly relative reform effects are first reported followed by an estimate for the first half year after the reform during which the law clearly was effective in the *Apr-Sep* row, which provide a summary estimate of the reform effect. Columns report different specifications. *Baseline* is the basic difference-in-differences specification. *Trend* adds a country specific time trend.

Table 1. Internet traffic

	(1) Baseline	(2) Trend
Apr-Jun	-0.225***	-0.209***
	(0.0150)	(0.0149)
Jul-Sep	-0.143***	-0.106***
	(0.0166)	(0.0168)
Oct-Dec	-0.0152	0.0478**
	(0.0191)	(0.0197)
Apr-Sep	-0.185***	-0.159***
	(0.0123)	(0.0127)
Country fixed effects	X	X
Time fixed effects	X	X
Country*Trend	-	X
No. of Observations	156	156

Note: Dependent variable is $\ln(Internet\ traffic)$. Standard errors are in parentheses. Relative reform effect estimates $e^{\alpha_t} - 1$ are reported. Data is weekly for 2009 only. *Time fixed effects* are monthly dummies common across countries. *Country*Trend* are country-specific linear time trends. * p<.1, ** p<.05, *** p<.01.

[28] It would have been interesting to translate the reform effect on piracy into bounds on the number of songs stolen and also the effect at the level of pirated units, i.e., the effect of a pirated unit on sales of that type of units. However, this is complicated by several factors that together make the bounds too widespread to be informative. First, our measure of Internet traffic does not measure all traffic: it only measures traffic between networks. As discussed in the background section, the relative change in between-network piracy is a good proxy for the relative change in all piracy. To obtain the absolute levels of total piracy requires, however, an estimate on the share of traffic that is between-network traffic. Second, music files come in different formats with a large variation in size of files, and the relative shares of different types of files are not known. Third, there are general equilibrium effects between different types of units. We cannot quantify any of those factors with any precision. We therefore refrain from attempting to derive such bounds. This contrasts with many other papers (e.g., Blackburn, 2006, Rob and Waldfogel, 2006, and Oberholzer-Gee and Strumpf, 2007) that estimate the individual behavioral effect at the level of pirated units but cannot say much about aggregate piracy.

The Internet traffic estimates are negative and statistically significant at the 1% level in each of the first two quarters and in the two quarters jointly in the baseline specification. The results are similar when trends are accounted for. The immediate reform effect in the first quarter is -20.9%, and falls to -10.6% in the second quarter after which the effect is not statistically significantly negative. The entire effect for the first six months prior to recovery is -15.9% and statistically significant at the 1% level. Under the assumption that half of Internet traffic consists of piracy (Schultze and Mochalski, 2009), this decrease corresponds to a decrease in piracy of 32%.

Although we can be very confident in interpreting the initial drop as a drop in piracy, subsequent effects are harder to interpret. The composition of Internet traffic other than piracy could have changed differentially between the countries as a result of the behavioral responses due to the reform. If individuals buy more physical music and listen to more music as a result, that may lead them to spend less time on legal online activities. On the other hand, individuals could also have migrated to bandwidth-intensive legal Internet music streaming services such as Spotify. Such effects could have occurred to a greater extent in Sweden than in its neighbors, producing a bias when the entire effect of the decrease in Internet traffic is attributed to a decrease in piracy. Because these effects are likely to have increased over time, we refrain from interpreting the estimates for the last three months as true reform effects. Although Internet traffic returned to its previous trajectory in late 2009, it is not necessarily the case that piracy also returned to its previous level and that the reform became ineffective.

5.2 Music sales

The reform effect estimates for music sales are reported in Panel A of Table 2 (Panel B of Table 2 is discussed in the next subsection), which is similarly organized as Table 1. We also add country-specific seasonal (monthly) dummies from specification *Season* and on. In column (4) and (5), we report estimates for physical and digital sales separately.

Table 2. Recorded music sales

	(1) Baseline	(2) Season	(3) Trend	(4) Physical	(5) Digital
			Panel A: Main results		
Apr-Jun	0.231**	0.381***	0.360***	0.333***	0.254***
	(0.122)	(0.103)	(0.102)	(0.108)	(0.0960)
Jul-Sep	0.255**	0.389***	0.368***	0.333***	0.689***
	(0.124)	(0.103)	(0.102)	(0.108)	(0.129)
Oct-Dec	0.209*	0.166**	0.148*	0.0942	0.924***
	(0.120)	(0.0865)	(0.0857)	(0.0886)	(0.147)
Apr-Sep	0.243***	0.385***	0.364***	0.333***	0.455***
	(0.0874)	(0.0729)	(0.0722)	(0.0765)	(0.0861)
			Panel B: Donald and Lang inference (n=9)		
Apr-Sep	0.243*	0.385*	0.364*	0.333*	0.455
	(0.0344)	(0.0504)	(0.0656)	(0.0409)	(0.122)
t-statistic	7.869	8.958	6.462	9.374	4.470
p-value	0.0805	0.0708	0.0977	0.0677	0.140
2008 Placebo	-0.0948*	-0.00601	0.0445	0.0463	0.0336
	(0.0130)	(0.0359)	(0.0564)	(0.0570)	(0.0236)
t-statistic	-6.934	-0.167	0.806	0.831	1.445
p-value	0.0912	0.895	0.568	0.559	0.385
Country fixed effects	X	X	X	X	X
Time fixed effects	X	X	X	X	X
Country*Season	-	X	X	X	X
Country*Trend	-	-	X	X	X
Observations	216	216	216	216	108

Note: Dependent variable is $\ln(Sales)$. Standard errors are in parentheses. Relative reform effect estimates $e^{\alpha_t} - 1$ are reported. Data is monthly, 2004-2009 for columns 1-4, and 2007-2009 for column 5. *Time fixed effects* are monthly dummies common across countries. *Country*Season* are country-specific monthly dummies. *Country*Trend* are country-specific linear time trends. Panel B reports estimates from the Donald and Lang (2007) two-step estimator. *2008 placebo* reports estimates from moving the reform back one year and dropping 2009. * p<.1, ** p<.05, *** p<.01.

The reform estimates for the first six months are all positive and statistically significant at the 1% level across specifications. The effect increases when seasonal effects are accounted for and stays similar when trends are accounted for. The reform effect is 36.4% for the first six months in the preferred specification. The composition of the total increase is made up of a physical sales increase of 33.3% and a digital sales increase of 45.5%.[29] The quarter estimates are all positive and in most cases statistically significant at the 1% levels for the first two quarters.

In the last quarter, most of the reform effect on physical sales vanishes. This result and the complete recovery in Internet traffic suggest that the reform was less effective in the long run. A possible explanation is that the perceived threat decreased over time as there were only a few legal cases that took a long time to reach final verdicts.

The digital sales increase is, however, stronger in the third quarter. Capacity constraints in streaming services immediately after the reform could be an explanation for the delayed upturn in digital music.[30] The sudden shift in demand could have triggered or at least enhanced the technological expansion in those services that caused further growth in sales

[29] Although the relative increase is larger for digital sales, the absolute increase is larger for physical sales.
[30] During this period, user accounts for the main music streaming service Spotify were distributed by invite only (except for the premium version of Spotify). The company could thereby maintain tight control on the growth of their user base.

even as Internet traffic (as well as piracy, although not to the same degree) recovered. If the technology expansion involves a fixed cost, it may not reverse back as the demand shifts back.[31] Another additional explanation for the persistency of digital sales in the last quarter could be that once people start to use streaming services regularly, it becomes a habit that they continue.[32] In a wider perspective, the reform may have marked a milestone where the digital market seriously started to compete with the physical market.[33]

The total increase in music sales and piracy of approximately 35% corresponds to a 1% to 1% displacement between piracy and music sales and a music sales elasticity of piracy of approximately 1 on the margin. Under the strong functional form assumption that the marginal effect is constant, music sales would have been twice as large in 2009 in the absence of piracy. Piracy could then account for 80% of the drop in music sales between 2000 and 2008, which would support the music industry's claims that piracy was the main cause of the decline. For countries with similar per capita sales and piracy as Sweden, these figures provide useful guidelines for the impact of piracy.

It is difficult to compare estimates between papers for several reasons. Comparing elasticities between samples is not meaningful since elasticities are not constant at the different levels of piracy between the samples. For example, it is not likely that doubling the amount of piracy has the same proportional effects on sales when piracy is high than when it is close to zero. Comparing marginal effects is more reasonable. However, without knowing the piracy levels in different countries and only having information on relative marginal effects, the comparison is problematic. The same relative effect of piracy on sales translates into different absolute effects of piracy on sales. Most studies find that music sales would have been 0 to 40% higher in the absence of piracy in the U.S. at some point in time during the period 1998 to 2006 (the estimates in Oberholzer and Gee (2006) are on the lower end, and the estimates in Blackburn (2006) are on the higher end). Liebowitz (2011) argues that most of these estimates imply that the emergence of online piracy could account for the entire recent decline in music sales, which is consistent with our results.

We now move on to analyzing physical sales in more detail. The reform effect estimates for different categories of physical music sales using the preferred specification with country-specific time effects are reported in Table 3. The reform estimates for the first six months are all positive and statistically significant at the 1% level across the different dimensions of partition.

---

[31] In September 2009, Spotify replaced their servers in order to accommodate their rapid growth ("Spotify Changes Servers to Reduce Energy Costs", Giorgiana Bursuc, Softpedia, September 19 2009. Available online: http://news.softpedia.com/news/Spotify-Changes-Servers-to-Reduce-Energy-Costs-122122.shtml).

[32] If there is a fixed cost involved in shifting music consumption between technologies, the reform could have provided the necessary incentive for people to make the switch to digital streaming services. When the perceived cost of piracy then falls, it is still optimal for some users to stay with the new technology, rather than shift back to piracy.

[33] Further evidence of the rapid growth of the Spotify service is that the performing royalties paid to Swedish copyright holders increased by a factor of eight between June 2009 and June 2010. ("Spotify-peng åtta gånger större", Mats Rörbecker, Dagens Nyheter, June 24 2010. Available online: http://www.dn.se/kultur-noje/musik/spotify-peng-atta-ganger-storre).

Table 3. Physical recorded music sales

	Albums, Price			Albums, Origin	
	(2) Full	(3) Mid	(4) Budget	(5) Domestic	(6) International
Apr-Jun	0.386***	0.780***	2.210***	0.419*	0.405***
	(0.141)	(0.295)	(0.823)	(0.276)	(0.139)
Jul-Sep	0.0864	1.204***	0.550*	0.862***	0.175
	(0.111)	(0.365)	(0.398)	(0.363)	(0.116)
Oct-Dec	-0.0302	0.194	0.170	0.0851	0.341***
	(0.0989)	(0.198)	(0.300)	(0.211)	(0.132)
Apr-Sep	0.227***	0.981***	1.230***	0.625***	0.285***
	(0.0910)	(0.234)	(0.420)	(0.226)	(0.0906)
Country fixed effects	X	X	X	X	X
Time fixed effects	X	X	X	X	X
Country*Season	X	X	X	X	X
Country*Trend	X	X	X	X	X
Observations	144	144	144	216	216

Note: Dependent variable is $\ln(Music\ sales)$. Standard errors are in parentheses. Relative reform effect estimates $e^{\alpha_t} - 1$ are reported. Data is monthly, 2004-2009. For columns 4–6, only data for Sweden and Finland are used. *Time fixed effects* are monthly dummies common across countries. *Country*Season* are country-specific monthly dummies. *Country*Trend* are country-specific linear time trends. * p<.1, ** p<.05, *** p<.01.

The reform effects are heterogeneous; they are larger for cheaper albums. If file sharers have a lower than average willingness to pay, increasing the cost of piracy from a very low level plausibly makes them start buying cheaper albums. The smaller effects on more expensive albums are also in line with the fact that piracy tends to be concentrated to albums for which file sharers have a relatively low willingness to pay, as documented by Rob and Waldfogel (2006).

The reform effects are also larger for domestic albums. This is a plausible pattern if file sharers perceived the risk of being sued to be higher for domestic albums. We believe that this is the case since previous court cases on piracy in Sweden almost exclusively focused on domestic music.[34]

Our results are complemented with individual survey data evidence, which is presented in Appendix B. A sample of individuals in Sweden has been asked about their file sharing and music consumption behavior prior to and after the reform. Questions on file sharing include how often the respondents used file sharing sites to download music and whether IPRED has changed their use of those sites. Questions on music consumption include whether consumption behavior changed after IPRED and how knowledge and use of digital services changed after IPRED. The answers corroborate our results that piracy decreased and that music sales increased after the reform. They also indicate that the reform really is the cause of the decrease in piracy and that music consumption increased to compensate for the decrease in piracy. Furthermore, there is evidence that streaming services played a key role.

[34] In a series of court cases in 2005 and 2006, four Swedish file sharers were convicted for sharing Swedish music files by mainly Swedish artists (“Tidigare dömda fildelare”, Anna-Karin Gustafsson, ComputerSweden, May 5 2008. Available online: http://www.idg.se/2.1085/1.159681).

5.3 Specification and sensitivity tests

In this subsection, we report results for music sales from the specification and sensitivity tests described in subsection 4.2. The first-six-month effect estimates and 95% confidence intervals for total music sales from placing a placebo intervention in each of the pre-reform months are plotted in Figure 4. The last estimate in the graph is the reform effect. The placebo estimates range between -0.158 and 0.160 compared to the reform estimate of 0.332. A kernel smoothed density plot of the distribution of point estimates is shown in Figure 5. In this figure, the distributions for physical and digital sales (with only 18 placebo estimates for digital sales, due to the shorter time series) are also shown separately, with the vertical lines showing the actual reform effect estimates.

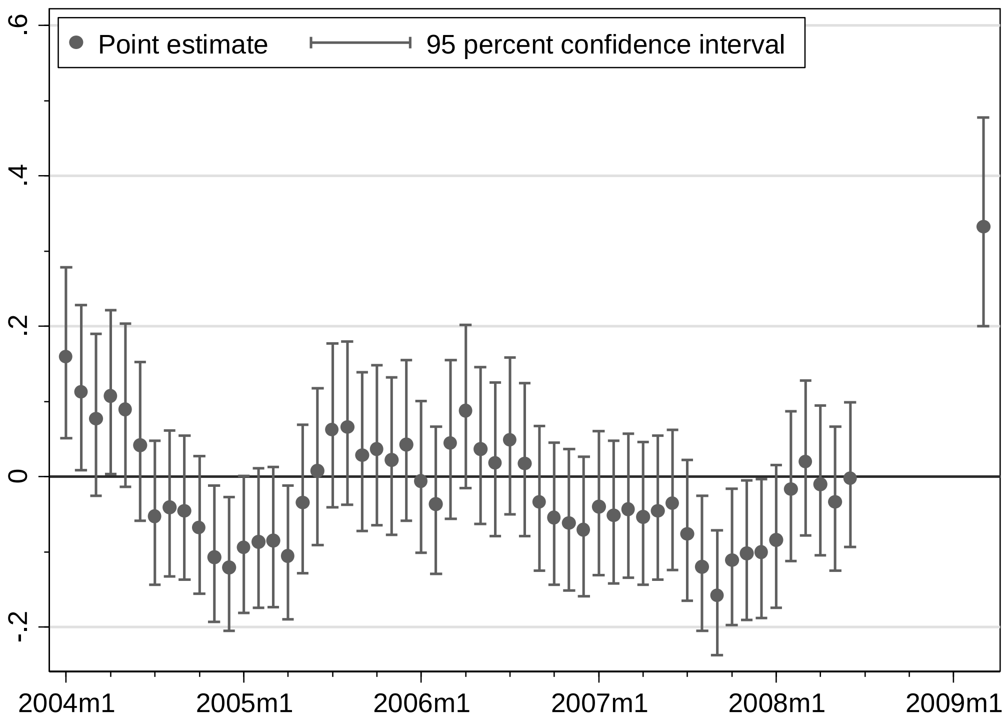

Figure 4. Placebo estimates for total music sales

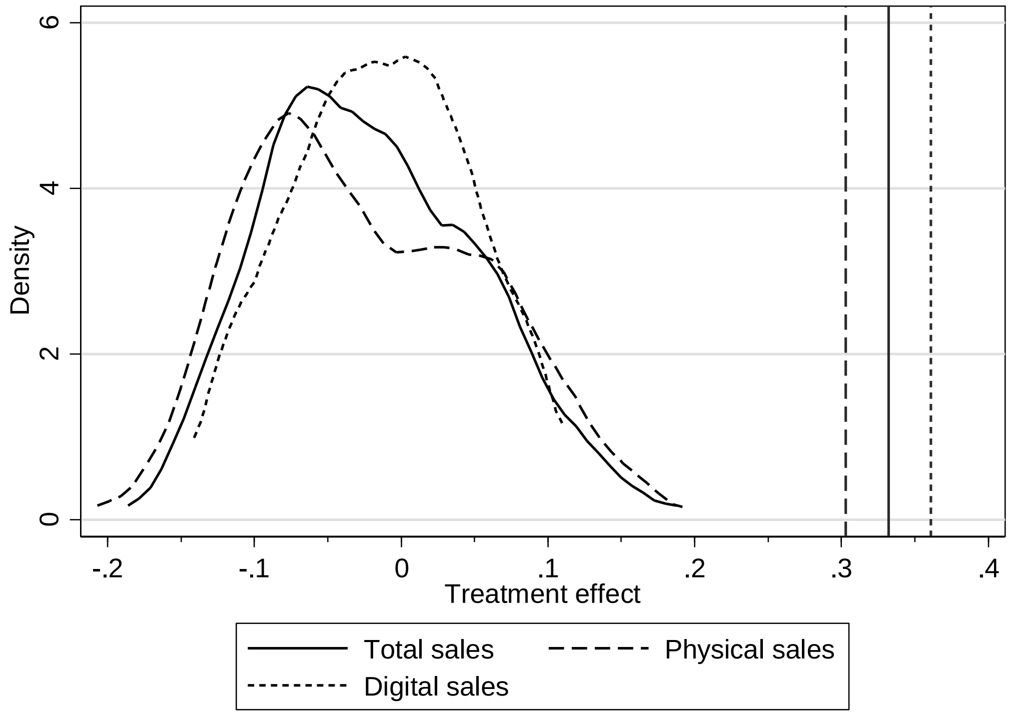

Figure 5. Kernel density of placebo estimates for music sales

We observe that the placebo estimates are much smaller than the reform estimate in all cases. They are also mostly not statistically significant. Furthermore they do not exhibit any clear systematic time related patterns for total sales, which would have been the case if there are remaining higher-order trend effects. This pattern is similar for physical and digital sales separately and is reassuring for our identification strategy.

If the standard errors of the point estimates are unbiased, we expect approximately 5% of the placebo point estimates to be significant at the 5% level. For total sales, 20.4% of the placebo estimates (11 out of 54) are statistically significant at the 5% level. The corresponding shares are 18.5% (10 out of 54) for physical sales and 0% (0 out of 18) for digital sales. These results indicate that standard errors may be a bit downward biased although the number of placebo estimates is too low for any meaningful formal test.

Because of the high rejection rates for total sales and physical sales, we consider alternative ways to make inference. Following Abadie et al. (2010), assuming that the distribution of placebo point estimates reflects the distribution of estimates when there is no actual effect, we ask what the probability of obtaining the estimated reform effect is if there was no effect by comparing our estimated reform effect to the placebo distribution. This gives an empirical p-value of at most 0.018 for total sales (1/55) and physical sales, and at most 0.0526 (1/19) for digital sales. These statistical significance levels are higher than before but still reasonably low. The fact that the reform estimate is much greater than each of the placebo estimates is quite compelling evidence that our results are not artifacts of statistical noise. It reflects that such a large and long deviation from the trend is rather unique.

As a further sensitivity test of potentially biased standard errors due to serial correlation, we implement Donald and Lang's (2007) two-step estimator which provides an upper bound on p-values. The results are reported in Panel B of Table 2. Notice that point estimates are identical to those in Panel A of Table 2. We report t-statistics and the associated p-values from the t-distribution with one degree of freedom. As a comparative exercise, we apply the same method on the placebo estimates placing the intervention in January 2008 and report the results below the reform effect estimates.

We observe that all specifications for total and physical sales are statistically significant at the 10% level, while digital sales are not far from significance at the 10% level. All placebo estimates have high p-values, with the exception of the baseline specification, indicating that controlling for country-specific time factors is important. Given the low power of this method, we believe the results should be viewed as supportive of our main findings.

6. Conclusion

This paper investigated the Swedish implementation of the IPRED law on Internet traffic and music sales using a difference-in-differences approach with Norway and Finland as control groups. We found that the reform decreased Internet traffic by 16% during the first six months. It also increased music sales by 36%. The reform effects on Internet traffic and physical sales disappeared after six months, but continued for digital sales.

IPRED was therefore effective in preventing piracy and in increasing music sales for the first six months. Some of these effects must be attributed to a combination of the law and the widespread public interest. The deterrent effect decreased quickly, possibly because of the few and slow legal processes. Law enforcement through convictions therefore seems to be a necessary ingredient for the long-run success of a copyright protection law. As the first court cases were only settled recently, it is still possible that further convictions would restore an effect that is more long-lasting.

While IPRED caused an immediate increase in the demand for legal music, the steady increase in digital sales, even after piracy recovered, suggest that the reform triggered or enhanced the technological expansion of digital services. The long-lasting effects on digital sales indicate that creating better legal alternatives to piracy is a complementary way to increase sales for the music industry, although some of this effect could just reflect a permanent migration from physical to digital music.

Our estimates imply a music sales elasticity of piracy of approximately one on the margin. Under the strong assumption that the marginal effect is constant, music sales would have been twice as large in 2009 in the absence of piracy. Piracy could then account for 80% of the drop in music sales between 2000 and 2008, which would support the music industry's claims that piracy was the main cause of the decline. Our results indicate that pirated music is a strong substitute to legal music. It is clear that copyright protection is important from the music industry's point of view. The sizeable substitution effects also suggest that there may be important welfare implications of file sharing, but a complete analysis requires estimates of the consumer surplus from piracy and of the sales effects on incentives for creators, which is outside the scope of this paper.

Appendix A – How Internet traffic works

Each device (server, personal computer, etc.) connected to the Internet has a unique IP address. A collection of adjacent IP addresses is referred to as an IP prefix. In order for data packets to reach their destination, the Internet Routing Table contains a list of possible paths by which a packet can reach a specific IP address.

On a more aggregated level, the Internet can be said to consist of Autonomous Systems (AS). An AS is usually a network operated by a single administration such as an Internet service provider (ISP), a government, or a large corporation. Examples of Autonomous Systems include AT&T, TeliaSonera, Microsoft, Google, and the Swedish University Computer Network (SUNET). Formally, an AS is defined as a group of one or more IP prefixes that have a single and clearly defined routing policy (Hawkinson and Bates, 1996). There are over 30,000 Autonomous Systems in the Internet Routing Table.[35]

In order for an Internet user to be able to reach any location on the Internet, Autonomous Systems have to interconnect. This is done through commercial agreements, which are classified as either peering or transit. A peering agreement allows two networks to exchange traffic between each other. In contrast, in a transit agreement one AS pays another AS to carry its traffic to and from the rest of the Internet. Transit agreements are usually used when a small AS needs to connect to the Internet. Traditionally, the Internet is divided into tiers with Tier 1 networks loosely defined as those large networks that tie together the Internet throughout the world. Tier 2 networks are smaller, often regional networks. The Tier 1 networks peer with one another, and the Tier 2 networks buy transit from the Tier 1 networks. Tier 2 networks can often also peer with one another.

Both peering and transit can physically be set up separately between each pair of connected networks, but most of these connections are made at Internet Exchange Points (IXPs). An IXP is a physical network infrastructure operated by a single entity with the purpose of facilitating the exchange of Internet traffic between Autonomous Systems (Radovcic 2009). The number of Autonomous Systems connected should be at least three, and there must be a clear and open policy for others to join.

Many IXPs operate sites in several geographic locations. IXPs are usually run as not-for-profit organizations. Euro-IX lists 121 known IXPs in Europe (Radovcic, 2009) while Packet Clearing House list 512 IXPs worldwide.[36] Netnod operates facilities in Stockholm, Gothenburg, Malmo, Sundsvall, and Lulea, and it connects 56 networks. NIX operates facilities in Oslo, Bergen, Trondheim, Stavanger, and Tromso, and it connects 69 networks. FICIX operates facilities in Helsinki, Espoo, and Oulu, and it connects 28 networks. In 2009, Sweden had the second highest per capita peak IXP traffic in Europe with only the Netherlands having higher per capita traffic (Radovcic, 2009).[37] Finland was ranked 8th and Norway 9th.

There are several technologies available that can potentially help file sharers hide their activities. Proxy servers work as intermediaries between a user and a web site or other server,

[35] CIDR Report, http://www.cidr-report.org/as2.0/.
[36] PCH Internet Exchange Directory: https://prefix.pch.net/applications/ixpdir/.
[37] Netnod members are listed on http://www.netnod.se/connected.shtml; NIX members are listed on http://www.uio.no/nix/nix-ops.html; FICIX members are listed on http://www.ficix.fi/english/member.php.

so that, e.g., a request for downloading a movie appears to come from the proxy server rather than from the user who is actually downloading the movie. The proxy server then forwards the data to the user. There are networks, such as the Tor Project,[38] that allow users to be completely anonymous online, by bouncing communication between many users of the network. VPN networks are, for our purposes, very similar to proxies, the main difference being that proxies work with specific applications, such as web browsers, while VPN routs all of a user's Internet traffic through an intermediate server. VPN also has encryption built-in, unlike proxies.

Appendix B – Individual survey data evidence

To provide additional corroborative evidence to our results, we present some survey-based evidence in this appendix. The survey questions cover file sharing and music behavior prior to and after the IPRED reform in Sweden as well as the reasons for altered behavior. In June 2009, two months after the reform, GfK conducted an investigation of music consumption habits in Sweden. One year later, a follow-up was made. The objective of the surveys was to provide a picture of the music market with respect to consumption behavior and attitudes. The target group was the Swedish population in the age span from 15 to 74 years old. An independent random sample from their Global Online Panel was used for each of the years. 1,006 responses were received in 2009 and 1,060 responses in 2010. The response rate was 66% in 2009 and 56% in 2010.

Questions and responses on file sharing behavior and the reasons for this behavior are reported in Table A1. Respondents were asked about their file sharing behavior. Although the 2009 survey was conducted after the reform, the responses on this question mainly reflect pre-reform behavior as the response alternatives referred to behavior up to a year ago. The 2010 responses reflect the post-reform behavior. The share of those that never file share increased from 57 to 61% after the reform, a difference that is statistically significant at the ten 10% level. When asked in 2010 about the extent of file sharing compared to the previous year, a statistically significant 52% responded that they file share less than last year (it is not entirely clear whether the respondents had pre-reform or post-reform behavior in mind when thinking of the last year here).

[38] https://www.torproject.org/.

Table A1. File sharing behavior and reasons

How often do you use file sharing sites for downloading music?		
	2009 (n=1006)	2010 (n=1060)
Never	0.57	0.61
Do you use file sharing sites for downloading music more, less, or as much as last year? (2010, n=1060)		
More than last year		0.08
As much as last year		0.40
Less than last year		0.52
Has the IPRED law changed your use of file sharing sites? (2009, n=429)		
I have stopped using file sharing sites		0.23
I use file sharing sites less		0.37
I use file sharing sites as much as before		0.37
I use file sharing sites more		0.03
What are the reasons that you use file sharing sites less to download music than last year? (2010, n=216)		
Spotify		0.56
IPRED		0.34
Better legal services		0.25

Notes: All response options are not reported for the first and last questions. Multiple response options were possible in the last question.

The group that responded that they did file share in 2009 was also asked whether the reform has changed their file sharing behavior. Of this group, 60% responded that they have either stopped or decreased their file sharing activities, which is a statistically significant share. The group that responded that they file share less in 2010 than last year was also asked about the reasons for this behavior. The digital streaming music service Spotify, the IPRED reform, and better legal services were the three main reasons, all having statistically significant shares. Of course, Spotify also existed before the reform, and the increased use of it as well as better legal services can be seen as outcomes driven by the reform.

Questions and responses on how music downloads through file sharing have been replaced are reported in Table A2. The respondents that answered that they file shared less or stopped file sharing due to the reform in 2009 were asked about how they replaced or compensated the music downloaded through file sharing. Almost half of them started to use free ad-financed digital streaming services. However, a statistically significant share did buy more music than previously. To explore their knowledge of different digital music providers, respondents were asked about their knowledge and use of different digital music services in 2009 and 2010. The three largest changes during the year after the reform occurred for digital music streaming services Spotify Free (ad-financed), Spotify Premium, and Sony Ericsson PlayNow. The changes are all statistically significant.

Table A2. How music downloads from file sharing have been replaced

In which way do you replace/compensate your previous use of file sharing sites? (2009, n=258)		
I use free ad-financed streaming services		0.49
I buy more music than previously		0.17
I did not buy music previously but I have started now		0.09
I pay for a music subscription without commercials		0.02
Which of the following do you know about?	2009 (n=1006)	2010 (n=1060)
Spotify Free	0.53	0.81
Spotify Premium	0.39	0.64
Sony Ericsson PlayNow	0.17	0.35
Which of the following do you use?	2009 (n=1006)	2010 (n=1060)
Spotify Free	0.23	0.42
Spotify Premium	0.02	0.07
Sony Ericsson PlayNow	0.04	0.07

Notes: All response options are not reported. Multiple response options were possible.

This survey evidence alone should, however, not be given too much weight. It provides a before-and-after analysis without accounting for other time effects. A survey data analysis also encounters the problems discussed in the introduction. However, the survey evidence provides corroborative consumption side support to the conclusion in this paper that piracy decreased and music sales increased after the reform.

References

Abadie, A., Alexis, D., Hainmueller, J., 2010. Synthetic control methods for comparative case studies: estimating the effect of California’s tobacco control program. Journal of the American Statistical Association 105, 493-505.

Adermon, A., Liang, C.-Y., 2010. Piracy, music, and movies: a natural experiment. Working paper 2010:18, Department of Economics, Uppsala University.

Andersen, B., Frenz, M., 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.

Angrist, J., 1990. Lifetime earnings and the Vietnam era draft lottery: evidence from social security administrative records. American Economic Review 80, 313-336.

Angrist, J., Krueger, A, 1992. The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples. Journal of the American Statistical Association 87, 328-336.

Angrist, J., Krueger, A., 1995. Split-sample instrumental variables estimates of the return to schooling. Journal of Business and Economic Statistics 13, 225-235.

Angrist, J., Pischke, J., 2008. Mostly Harmless Econometrics: An Empiricist’s Companion. Princeton: Princeton University Press.

Bertrand, M., Duflo, E., Mullainathan, S., 2004. How much should we trust differences-in-differences estimates?. Quarterly Journal of Economics 119, 249-275.

Bhattacharjee, S., Gopal, R., Lertwachara, K., Marsden, J., Telang, R., 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., Gopal, R., Lertwachara, K., Marsden, J., Telang, R., 2007. The effect of digital sharing technologies on music markets: a survival analysis of albums on ranking charts. Management Science 53, 1359-1374.

Blackburn, D., 2006. The heterogenous effects of copying: the case of recorded music. Job market paper, Harvard University.

Boldrin, M., Levine, D., 2002. The case against intellectual property. American Economic Review 92, 209-212.

Bounie, D., Bourreau, M., Waelbroeck, P., 2005. Pirates or explorers? Analysis of music consumption in french graduate schools. Telecom Paris Economics Working Paper No. EC-05-01.

Danaher, B, Smith, M., 2014. Gone in 60 seconds: the impact of the Megaupload shutdown on movie sales. International Journal of Industrial Organization 33, 1-8.

Danaher, B., Smith, M., Telang, R., 2013a. Piracy and copyright enforcement mechanisms. In: Lerner, J., Stern, S. (Eds.). Innovation Policy and the Economy 14, forthcoming, Chicago: University of Chicago Press.

Danaher, B., Smith, M., Telang, R., Chen, S., 2013b. The effect of graduated response anti-piracy laws on music sales: evidence from an event study in France. Journal of Industrial Economics, forthcoming.

Dejean, S., 2009. What can we learn from empirical studies about piracy? CESifo Economic Studies, 55, 326-352.

Donald, S, Lang, K., 2007. Inference with difference-in-differences and other panel data. Review of Economics and Statistics 89, 221-233.

Duchene, A., Waelbroeck, P., 2006. The legal and technological battle in the music industry: information-push versus information-pull technologies. International Review of Law and Economics, 26, 565-580.

Halvorsen, R., Palmquist, R., 1980. The interpretation of dummy variables in semilogarithmic equations. American Economic Review, 70, 474-475.

Hammond, R., 2013. Profit leak? Pre-release file sharing and the music industry. Southern Economic Journal, forthcoming.

Hawkinson, J., Bates, T., 1996. Guidelines for creation, selection and registration of an Autonomous System (AS). Network Working Group.

Hong, S., 2013. Measuring the effect of Napster on recorded music sales: difference-in-differences estimates under compositional changes, Journal of Applied Econometrics 28, 297-324.

Hurt, R., Schuchman, R., 1966. The economic rationale of copyright. American Economic Review 56, 421-432.

Ipeg, 2006. Summary of the implementation of the directive 2004/48 on the Enforcement of Intellectual Property Rights (the ”Directive”) in EU Member States as per October 2006. Intellectual Property Expert Group, Hague.

Liebowitz, S., 2006. File-sharing: creative destruction or just plain destruction?. Journal of Law and Economics 49, 1-28.

Liebowitz, S., 2008. Research note - testing file sharing’s impact on music album sales in cities. Management Science 54, 852-859.

Liebowitz, S., 2011. The metric is the message: how much of the decline in sound recording sales is due to file-sharing? Manuscript, http://ssrn.com/abstract=1932518 or http://dx.doi.org/10.2139/ssrn.1932518.

Lovell, M., 1963. Seasonal adjustment of economic time series and multiple regression analysis. Journal of the American Statistical Association 58, 993-l0l0.

Michel, N., 2006. The impact of digital file sharing on the music industry: an empirical analysis. Topics in Journal of Economic Analysis and Policy 6.

Moulton, B., 1990. An illustration of a pitfall in estimating the effects of aggregate variables on micro units. Review of Economics and Statistics 72, 334-338.

Newey, W., West, K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708.

Oberholzer-Gee, F., Strumpf, K., 2007. The effect of file sharing on record sales: an empirical analysis. Journal of Political Economy 115, 1-42.

Peitz, M., Waelbroeck, P., 2004. The effect of Internet piracy on music sales: cross-section evidence. Review of Economic Research on Copyright Issues 1, 71-79.

Plant, A., 1934. The economic aspects of copyright in books. Economica 1, 167-195.

Radovcic, S., 2009. European Internet Exchange Association 2009 Report on European IXPs. Euro-IX.

Rob, R., Waldfogel, J., 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, 29-62.

Schultze, H., Mochalski, K., 2007. Internet Study 2007. Ipoque, Leipzig.

Schultze, H., Mochalski, K., 2009. Internet Study 2008/2009. Ipoque, Leipzig.

Takeyama, L., 1994. The welfare implications of unauthorized reproduction of intellectual property in the presence of demand network externalities. Journal of Industrial Economics, 42, 155-166.

Varian, H., 2005. Copying and copyright. Journal of Economic Perspectives 19, 121-138.

Zentner, A., 2005. File sharing and international sales of copyrighted music: an empirical analysis with a panel of countries. Topics in Economic Analysis and Policy 5.

Zentner, A., 2006. Measuring the effect of file sharing on music purchases. Journal of Law and Economics 49, 63-90.