

The Effect of Graduated Response Anti-Piracy Laws on Music Sales: Evidence from an Event Study in France

Brett Danaher[‡]
bdanaher@wellesley.edu

Michael D. Smith[†]
mds@andrew.cmu.edu

Rahul Telang[†]
rtelang@andrew.cmu.edu

Siwen Chen[‡]
siwychen@gmail.com

This Version: August 2012

Acknowledgements: The authors thank four major record labels for generously providing data to support this research. Smith and Telang acknowledge generous support from Carnegie Mellon's iLab. While this research was undertaken independently, the IFPI has compensated Danaher for his time to present the results of this study to the IFPI and to produce a non-technical version of this paper for their internal use. We also thank Adeline Tan and Ljubica Ristovska for excellent research assistance.

[‡] Department of Economics, Wellesley College, Wellesley, MA, 02481

[†] School of Information Systems and Management, Heinz College, Carnegie Mellon University, Pittsburgh, PA, 15213.

The Effect of Graduated Response Anti-Piracy Laws on Music Sales: Evidence from an Event Study in France

ABSTRACT

Despite the problem that filesharing poses to creative industries, there is little research on the effectiveness of governmental anti-piracy policies. This study analyzes how the HADOPI graduated response law in France affected music sales on the iTunes store.

We obtained a panel of iTunes sales data from the four major labels across a broad set of countries. We applied a difference-in-difference approach, comparing sales trends in France to a control group of European countries. Our results suggest that increased consumer awareness of HADOPI caused iTunes music sales to increase by 22-25% relative to changes in the control group.

We find that these sales changes are similar for each of the four major labels, suggesting that our results are not peculiar to any particular label. We also find that the observed sales increase is much larger in high piracy genres than low piracy ones, strengthening the causal interpretation of our results.

Keywords: *Piracy, regulation, digital distribution, music industry, natural experiment.*

1. Introduction:

Since the rise of Napster, global recorded music sales and licensing have plunged from nearly \$27 billion US dollars in 2000 to \$15 billion in 2010,¹ with some countries witnessing a coinciding decrease in investment in developing local talent (IFPI 2010). Concurrently, a strain of economics literature has emerged demonstrating that at least some of this decrease was due to online filesharing. (see for example, Liebowitz (2006), Rob and Waldfogel (2006), Zentner (2006), Hui and Png (2003), OECD (2009)). The media content industries have lobbied strongly to governments asking for policies that better protect intellectual property in light of Internet filesharing, resulting in what has been a very controversial debate over if and how copyright policy should be adapted to the digital era. One prime example was the recently proposed Stop Online Piracy Act (SOPA) in the US, a proposal that was portrayed by the opposition as a Draconian law that would inhibit the growth of the Internet and modern culture (the act failed to pass a Congressional vote). Bearing in mind the importance of copyrighted works to markets and to culture, it seems important to discuss not only which policies are too restrictive but also what sorts of government policies are actually effective in inducing filesharers to turn to legitimate consumption channels. However, to date we are aware of no economic evidence on the effect of direct government policy intervention on filesharing and media sales. Our goal in this paper is to provide such evidence by evaluating the effectiveness of a highly publicized anti-piracy law passed in France.

In May of 2009 the French Parliament passed an anti-piracy law known as HADOPI, or the Creation and Internet Law. The purpose of this law is to “promote the distribution and protection of creative works on the Internet.” The law empowers the HADOPI administrative authority to send warnings to identified infringers and transfer the case to the court in cases of repeat infringement. HADOPI acts on the basis of information submitted by rightholders and has the power to monitor online infringement and verify information with ISPs. When a rightholder submits a notice of infringement, it is verified by HADOPI

¹ Source: IFPI. This includes both digital and physical sales of recorded music.

and matched against information held by the relevant ISP. Valid infringement triggers a notice of infringement sent by email from HADOPI to the account holder. When the same account is identified again as being used for infringement within a period of six months of the first warning, a second warning is sent by HADOPI to the account holder, this time by registered mail. Where the account is identified for a third time within a period of one year, HADOPI may escalate the case by referring it to the criminal court, where a judge is empowered to order a range of penalties, including account suspension for up to one month.² This law has been controversial on several major fronts, including the cost of the law, suggestions that Internet access may be a human right, potential violations of the principle of net neutrality, and the possibility that the law could hold Internet users responsible for copyright violations even if their computers have been hijacked.

The purpose of this study is not to debate the broader social and policy merits of this law. Instead, our purpose is to analyze whether this law had an impact on consumer behavior. This question is important for two reasons. First, music industry profits are clearly important to individuals working in that industry. Second, a broader social concern is that if the media industries are less able to recoup profits on their investments in creative works they will likely decrease their investment in bringing new music, television, and films to society, thus reducing overall social welfare. While more restrictive proposals such as SOPA do not seem to have enough political support to pass into law, it may be that relatively less heavy-handed laws like HADOPI can influence consumer behavior.

In this regard, we evaluate the effectiveness of HADOPI using a panel of iTunes sales data for the four major music labels (Universal Music, Warner Music, EMI Music and Sony Music) across a broad set of countries.³ We employ a difference-in-difference approach, using sales trends in selected European countries to simulate the counterfactual of what music sales in France would have been if HADOPI had

² See the HADOPI law itself in articles L 331-25 and L 336-2. Or an explanatory memorandum from the French government at <http://www.culture.gouv.fr/culture/actualites/conferen/albanel/creainterenglish.pdf>.

³ We study digital sales for two reasons: first, the data are more readily quantifiable than physical sales are (unlike with digital sales, there are no comprehensive sources of physical sales data). Second, previous studies have shown that illegal Internet downloaders prefer digital sales channels over physical ones when purchasing legally. For example, see Danaher et al. (2010).

not been passed. Using Google Trends, we find that public awareness of HADOPI became widespread in Spring 2009, and our difference-in-difference model suggests that HADOPI awareness caused a 22.5% increase in iTunes song unit sales in France (over and above any change in the control group), as well as a 25% increase in iTunes album unit sales (over and above the change in the control group). Closer examination reveals similar trends separately for each of the four major music labels, suggesting that our industry-wide results are not driven by one label's advertising campaign or marketing activity.

To test the validity of our results, we add another level of difference to the model. Previous research and new survey data reveal that music genres differ in their tendency to be pirated. One would expect that if the observed relative increase in French sales is caused by HADOPI, that high-piracy genres would experience a larger increase in sales than low-piracy genres do. Our results are consistent with this hypothesis: low piracy genres experienced only a 7% difference-in-difference sales increase in France after HADOPI, while high piracy genres experienced a 30% difference-in-difference increase in sales, a result that is consistent with the hypothesis that the observed increase in French sales after HADOPI is due to a reduction in Internet piracy.

2. Background on Music Industry and HADOPI

Looking at aggregate sales reports it's easy to see why the music industry might be concerned about the impact of piracy on sales. Forrester research and the Recording Industry Association of American (RIAA) have reported that music industry revenue in the United States dropped by 46% from \$14.6 billion in 1999 to \$7.7 billion in 2009.⁴ Worldwide sales have seen a similar drop of 44% from \$27 billion in 2000 to \$15 billion in 2010. Studies by the IFPI have found corresponding decreases in investment in local talent in some countries. However, the economic literature is only just beginning to address the question of whether diminished music industry returns due to piracy cause a decrease in the amount of creative

⁴ <http://www.riaa.com/faq.php>

works brought to the market.⁵ Nonetheless, academic studies on the effect of piracy suggest that online file sharing can explain anywhere from one fifth to all of the decrease in music industry revenues since 2000,⁶ and because of this it is important to understand what sorts of policies or strategies can act to mitigate this negative impact.

In June 2008 the HADOPI Law was first presented to the French Senate by several politicians, and in October 2008 the Senate backed the law, meaning that it would next go before the French National Assembly. (However, as we will show in more detail below, there was relatively low publicity around the law or national awareness of it during this time). In March 2009 the HADOPI law was presented to the National Assembly, where it was at first supported and then rejected in 2009. This debate gave rise to a number of media articles generating awareness and controversy.⁷ Importantly, this could have led to some confusion among the general populace over whether the law was yet effective or not. In May 2009 both the National Assembly and the Senate backed an amended version of the law, leading opposition Parliamentarians to send it to the French Constitutional Council for review. In June 2009 the Constitutional Council rejected the main part of the law (again potentially adding to confusion over whether the law was yet in effect), largely over the issue of judicial review for penalties imposed by the third strike (in this early version of the law, penalties could be applied to individuals on the third infringement without judicial review). This section was then amended to require judicial review, and the Constitutional Council accepted the amended law in October 2009, putting the law into effect.⁸

From that point onward, filesharers in France could theoretically begin to receive notices of infringement. However, it was over a year before the HADOPI agency began sending out first notices, with the first wave of infringement notices going out in September 2010. Later, in Spring 2011, the HADOPI agency

⁵ See, for example, Oberholzer-Gee and Strumpf 2009 or Waldfoegel 2011.

⁶ See the music piracy studies cited in our literature review.

⁷ For example, <http://news.bbc.co.uk/2/hi/europe/7992262.stm>

⁸ For references regarding the political timeline of the HADOPI law, see law <http://www.senat.fr/dossier-legislatif/pjl07-405.html> and <http://www.conseil-constitutionnel.fr/decision.42666.html>.

began the initial wave of second notices, and as of August 2011 no third notices had been sent out and no penalties had yet been applied.⁹

It is also important to note that while the most publicized responsibility of the HADOPI agency under the law is to send out infringement notices, as part of the HADOPI legislation the agency started an education campaign to inform citizens about the illegality of sharing copyrighted materials online, the dangers it may pose to content generation in the future, and the various legal channels in which media can be obtained. These campaigns are an ongoing part of HADOPI's responsibilities. These campaigns are ongoing and do not represent discrete events. Thus their effects cannot be separated from the effects of the graduated response/penalty portion of the law, and our study must be about the combined effect of both these education campaigns and the warning and penalty system.

3. Literature Review

Our research is novel in that despite the current debate over copyright policy reform, we are aware of no studies that examine the impact of government anti-piracy legislation on consumer behavior or media sales. Thus our research progresses in a new direction from the large existing literature analyzing how piracy has impacted media sales. This literature has used a variety of empirical methods from cross-country or cross-city variation in piracy levels (for example, Zentner 2005, Hui and Png 2003, Peitz and Waelbroeck 2004, Liebowitz 2008, and Danaher and Waldfogel 2011) to survey data (for example, Rob and Waldfogel 2006, Rob and Waldfogel 2007) to exogenous shocks in the availability of pirated or legitimate content (for example, Oberholzer-Gee and Strumpf 2007, Smith and Telang 2009, Danaher et al. 2010). What this literature has in common is that nearly all of the academic studies find that media piracy has a significant negative impact on sales. The one exception to this rule, Oberholzer-Gee and Strumpf 2007, is also one of the earliest studies of piracy's impact, and has seen some concerns expressed

⁹ This timeline is outlined in a short report titled "Hadopi, cultural property and Internet usage: French Internet users' habits and points of view" available at <http://www.hadopi.fr>.

about its findings in recent years (see Liebowitz 2008, Liebowitz 2011). Notably, there exist some studies on the effectiveness of private sector legal responses to filesharing. Bhattacharjee et al. (2008) find that the RIAA's legal threats (in the form of highly publicized lawsuits) against file sharers had a statistically significant negative impact on the availability of pirated content, but that a substantial amount of illegal content was still available even after the lawsuits and that piracy supply eventually returned to original levels

4. Theory

Our main theory of the impact of HADOPI is closely tied to the original intent of the law: HADOPI may cause consumers who otherwise would have pirated music to avoid piracy, and some of these consumers may instead purchase music through legitimate channels. There are, however, two important theoretical questions about the impact of the law that warrant discussion. First, when should we expect the impact of the HADOPI law? Will it occur primarily around national awareness of the law, or when the law actually goes into effect? Second, will HADOPI cause increased consumption of legal music and will this consumption occur in digital or physical channels?

With respect to the second question, as noted above, the literature seems to suggest that consumers are strongly tied to either the digital or physical channel such that if a consumer is forced to stop consuming digital piracy, the literature suggests that they are more likely to switch to other digital channels than they are to return to CD purchases (see, for example, Danaher et al. 2010, Hu and Smith 2011). Because of this, we focus on the impact of HADOPI on digital music sales,¹⁰ reflecting our belief that if HADOPI impacts individuals' ability to pirate online, they are more likely to turn to digital music channels than they are to go back to physical purchases of CDs.

¹⁰ Data from the Syndicat National de L'édition Phonographique (SNEP) indicate that 20% of French music sales in the first 3 quarters of 2011 were digital, with physical making up the remainder (<http://www.disqueenfrance.com/fr/catalogpage.xml?id=420906&pg=1&cat=251362>).

However, since nearly all prior papers in the literature indicate that filesharing displaces sales of physical music, it is a limitation of this paper that we are only able to examine iTunes sales data as it is possible that HADOPI could affect physical sales in addition to digital sales. Unfortunately, physical retail sales data were not available to us at the time of this study. Our data also do not reflect revenues from newer legal music streaming platforms such as Spotify or Youtube channels, and we discuss this limitation further in the conclusion of this paper. Thus, our study asks the question of whether a law like HADOPI can stimulate sales of one form of media (music) in one channel (iTunes), and as such likely does not capture the total impact of the law.

With respect to the first question, there is no strong theoretical basis in the literature for whether HADOPI's impact will begin primarily when the public becomes aware of the law or with the actual dates of passage, legal notifications, and legal penalties. Many economic studies of policy changes focus on the date of passage of the new policy as the treatment date. However, because HADOPI went through a significant public and political debate before being passed, and because citizens may not have even been aware of the actual effective date of the law, we believe it possible that public awareness and salience of the law could drive a change in behavior before it actually became effective. Lacking a strong theoretical basis, in our analysis below we test for both effects, finding a stronger impact around the peak awareness of HADOPI than from the specific passage of the law or the dates associated with first notifications of violations. Notably, some research in the criminology literature suggests that a policy change (such as an announced police crackdown on crime) may have an impact on behavior prior to actual enforcement of the policy, particularly when there is uncertainty among offenders of the timing of enforcement or the probability of being caught (Sherman 1990).

5. Data

For this study, we obtained a panel of total weekly iTunes sales units for a number of European countries including France. Our data extend from July 2008 to May 2011, and we observe separately both track unit

sales and album unit sales. The data were obtained directly from the four major music labels — EMI, Sony, Universal, and Warner — and aggregated to reflect total iTunes sales for the majors. According to the IFPI, the four majors reflect roughly 70% of music industry sales, with independent labels reflecting the other 30% not observed in our data. We chose the five European countries (other than France) with the highest iTunes sales levels as our control group for France, under the theory that overall market trends would have the most similar impact on countries with closer sales levels.¹¹ Thus, in this study we observe weekly iTunes sales units for France, the UK, Italy, Spain, Germany, and Belgium. This yields 918 country-by-week observations of total iTunes sales units, broken down into albums and tracks.

We also have a separate dataset provided to us by just two of the four major labels. This dataset is similar to the dataset above except that for each country-week, we observe total iTunes song sales units for each genre of music. We use this dataset as a further test of our main effect, focusing on the following genres: Rap, Hip Hop, Rock, Pop, Classical, Christian, Folk, and Jazz. Our purpose in choosing these genres was to keep only genres which had a significant share of the market and for which we had reasonable priors regarding relative piracy levels for the genre.¹²

Table 1 shows the average sales levels per week (both track and album) across each of the countries as well as providing some indicators of the level of variance within each country.

The UK is clearly the largest country in terms of iTunes sales, and Spain is the smallest. We note from Table 1 that while there is significant variance across countries in terms of sales levels, the variance within countries (across weeks) is relatively smaller and less skewed. In spite of the variance in sales across countries, we will show that the average time trend of our control group — on a logarithmic scale — closely maps the time trend of France.

¹¹ However, any results reported in this paper are not sensitive to this selection of control group. Results remain roughly the same (and completely the same in sign and significance) when choosing a variety of other control groups).

¹² For example, we dropped the Country genre not because its market share was too low, but rather because we could find no research or data giving us information as to its a priori tendency to be pirated.

Finally, we postulate that because citizens may be confused over when the law actually became effective, we may be more likely to see an effect begin when people first became aware of the law rather than when it finally became effective in October 2009. To measure awareness, we collected Google Trends data on Google searches (from France) for the search term “HADOPI.”¹³ Google Trends reports the “relative search index” for a search term in a given country, meaning that for each week we observe the number of searches for that term relative to the average number of all searches in that country across each week in the date range. So, for example, if there were an average of 20,000 searches per week in our date range for the term HADOPI, then in a week where there were 100,000 searches, the Google Trends index would report “5.0” for that week. Thus, while we do not know the actual volume of searches, we know when awareness of HADOPI peaked as measured by Google searches, and the relative height of that peak. This can serve as a measure of national awareness of the law.

6. Results

Our basic strategy for determining the impact of HADOPI is to use a difference-in-difference approach, comparing the change in French sales before and after HADOPI to the average change in sales across the control group. However, we have two initial challenges to overcome. First, we need to give evidence that our control group truly can simulate the counterfactual of what France’s sales would look like in the absence of HADOPI. Second, we need to determine the appropriate “treatment date” on which HADOPI began.

We begin addressing both of these issues with the following model:

$$\ln Sales_{it} = \beta_0 + \beta_1 \Omega_t + \beta_2 \Omega_t * France_i + \mu_i + e_{it} \quad (i)$$

Where $\ln(sales_{it})$ is equal to the natural log of song sales units in country i during week t , Ω_t is a vector of dummy variables for each week of the data (time fixed effects), $France_i$ is an indicator variable equal to

¹³ Google Trends data for other related search terms such as “French Three Strikes Law” show the same pattern.

1 for French observations, μ_i is a vector of country fixed effects, and e_{it} is the idiosyncratic shock term. With this specification, β_1 tracks the average time trend for log sales units over time for the control group, while β_2 estimates how the French time trend differs from this average.

We estimate this model and then plot the results visually in Figure 1, with $\beta_0 + \beta_1$ representing the average sales time trend for the control countries and $\beta_0 + \beta_1 + \beta_2$ representing the French sales time trend. Both of these time trends are plotted and measured on the left axis in Figure 1. We also added to this graph the Google Trends relative index for the search term HADOPI (in France), measured on the right axis.

Examining this table, we first note that the initial peak for the Google Trends Index occurs in March 2009, corresponding to the presentation of the HADOPI law to the National Assembly. Prior to March 2009, the French sales trend appears to follow closely the sales trend of the control group. A Wald test of joint significance for all β_2 between July 6, 2008 and March 30, 2009 could not reject the null hypothesis that the coefficients are jointly zero at the 95% confidence level. Thus, prior to the first Google Trends spike, the time trend for French sales is statistically indistinguishable from the sales trend of the control group. This lends credibility to the identifying assumption of our difference-in-difference model, namely that the control group simulates the counterfactual of how France's sales would have trended in the absence of HADOPI.

Figure 1 also helps to shed light on the question of when to consider the effective treatment date of HADOPI. We see the French sales trend diverge from the control group starting in March 2009, coinciding with the first peak of HADOPI awareness according to Google Trends. French sales then rise further relative to the control throughout the following several months. During this time, we also observe two more peaks of HADOPI awareness; the first occurs in April when the National Assembly supports and then rejects the law, and the second is in May when the Assembly and the Senate both backed the

law. Because each of these peaks is higher than the last, it seems safe to assume that this reflects growing national awareness of the HADOPI law. Thus, this graph leads us to believe that the effect of HADOPI began with rising awareness of the law and not upon its actual implementation. This seems plausible for the reasons we outlined above. We also note a very similar trend for iTunes album sales. We applied the same model as above to iTunes album unit sales and graph the results in Figure 2.

If we accept the identifying assumption that France would have followed the sales trend of the control group if not for HADOPI, then the average effect of HADOPI can be measured as the average gap between the two sales trend lines after March 30, 2009¹⁴. In order to measure this gap and test for statistical significance, we estimate the following slightly different version of the first model:

$$\ln Sales_{it} = \beta_0 + \beta_1 \phi_t + \beta_2 \phi_t * France_i + \mu_i + e_{it} \quad (ii)$$

The only difference between (ii) and (i) is that in (ii), ϕ_t is an indicator variable equal to one if the observation occurs after March 30, 2009. Thus, this model measures the average post-HADOPI change in sales across the control countries (β_1) and then estimates any change in French sales over and above the change in the control group (β_2). This means that β_2 is the coefficient of interest and under our identifying assumptions, it represents the average causal effect of HADOPI on weekly iTunes sales units in France from March 2009 until May 2011.

In columns (i) and (ii) of Table 2 we display the results of estimating this model for both single tracks and for albums. Accounting for the fact that our independent variable is in log terms, these estimates indicate that iTunes track sales units rose about 25.5% in the control group after March 1, 2009 but by 48% in France, indicating that French iTunes track sales were 22.5% higher on average than they would have been in the absence of HADOPI. Similarly, album sales units rose by 42% in the control group but 67%

¹⁴ We note that the use of March 30, 2009 as the treatment date will likely yield a slightly conservative estimate of the impact of HADOPI. In Figure 1, it appears as if French sales may have begun to diverge from the control group during March and that the awareness of the law began to spike in early March, and so our conservative use of March 30 as the start date of HADOPI will bias our difference-in-difference estimate slightly toward zero.

in France, indicating that HADOPI increased iTunes album sales an average 25% per week in France. To deal with serial correlation within each country, we compute robust clustered standard errors as suggested by Bertrand et. al. (2004) and we report these in parentheses. However, because there is only one treated cluster, even these standard errors may be inappropriately small for the coefficient of interest. Thus for β_2 we conservatively calculate p-values from a permutation test: we run the model five additional times specifying each country as the “treated” country and obtain estimates for the coefficient of interest in each case. We then impute p-values for the French coefficient of interest as the probability that the French estimate was drawn from a t-distribution (with four degrees of freedom) containing the estimates for the other five countries. These p-values are presented in square brackets below the robust clustered standard errors. Our, difference-in-difference results are statistically significant at the 95% confidence level.¹⁵

In addition to these aggregate results, we also estimated the same models for each of the four major labels individually, and the results for each individual label exhibit roughly the same patterns as the aggregated results (particularly in sign). Thus, it is unlikely that our results are driven by the marketing efforts or release schedule of an individual firm. As well, marketing, advertising, and release schedule decisions must — by law — be made non-cooperatively, making it less likely that they will be correlated across labels. In short, the effects we have observed thus far appear to be industry-wide and not specific to any particular label’s music.

We also conducted a series of placebo tests. We ran the same difference-in-difference model five more times, each time supposing one of the countries in our control group to be the “treatment” country. We were unable to reject the null hypothesis for any country in these tests; in other words, for each country (other than France) we cannot reject the hypothesis that that country’s sales increased by no more than the average of the other countries. Thus, at least with regard to the countries we studied, our findings are unique to France.

¹⁵ We also collapsed the data into 12 observations, summing up track sales for each country before and after March 30. When we estimate the model on these 12 observations, the coefficient of interest is still estimated as 0.2 and is still significant at the 95% confidence level.

6.1 Genre Analysis

Policy changes such as these are often difficult to study due to a lack of experimental power. Indeed, although we observe five control countries over time, we only observe one experiment: the passage of HADOPI in France. One could argue that some other French-specific factor may have coincidentally affected France's music sales levels at exactly the same time as public awareness of HADOPI. While we believe that it is unlikely that the timing of this would coincide with the passing and awareness of HADOPI, we cannot rule out this possibility based on the preceding tests.

In an attempt to partially address this concern we add an additional level of difference to the model based on priors about the popularity of piracy across various genres of music. Specifically, EMI surveys of French citizens show that that Rap and Hip Hop are the most heavily pirated genres, even relative to popularity in legal sales channels. While Rock and Pop experience average levels of piracy, the data also indicate that genres such as Classical, Christian, Folk, and Jazz experience significantly lower levels of piracy.¹⁶ We also note that other published studies suggest a similar distribution of pirated downloads across genres.¹⁷ This allows us to increase the number of experiments in that we can treat each genre as a unique experiment based on its *a priori* tendency to be pirated. If the observed increase in French sales is due to a reduction in piracy triggered by HADOPI, we would expect the increase in Rap sales to be larger than that for Rock and Pop and the increase for Classical, Christian, Folk, and Jazz to be quite low.

In columns (iii), (iv), and (v) of Figure 2 we see that that the rise in French sales after HADOPI (relative to the control) is indeed highest for the most heavily pirated genres (column v) and is small and statistically insignificant for the least pirated genres (column iii). We conducted an F-test across each estimation to test whether the estimate of the effect of HADOPI on French sales is different across genres. Due to relatively large standard errors we cannot reject the null hypothesis the post HADOPI French

¹⁶ These genre findings are based on surveys performed by EMI Music in France during October and November 2010. 8,173 interviews were completed and results were weighted to nationally representative figures. The data are not publicly available but were available to the authors.

¹⁷ For example, Liebowitz (2008).

increase is the same across genres.. Nonetheless, the point estimates indicate that that while the effect of HADOPI increased iTunes sales of Classical, Christian, Folk, and Jazz genres in France by 7%, the point estimate of the effect on Rock and Pop was 17% and the effect on Rap and Hip Hop was 30%.

Unobserved changes that might increase French music sales (around the same time as HADOPI) would be unlikely to have this same pattern across genres. For example, if Apple began to heavily promote the iTunes store in France more so than in the control countries around the same time as HADOPI, one would expect such promotions to affect each genre equally, at least proportionally to prior sales. Instead, we see sales in France for heavily pirated genres rise faster than for less pirated genres, which suggests that this sales increase is due to a reduction in French piracy levels. This is consistent with the idea that the difference-in-difference increase we observe in France is actually attributable to HADOPI.

6.2 Challenges to Identification

In this study, causal inference rests on the assumption that sales of the control group would have trended similarly over time to French sales in the absence of HADOPI. Indeed, the time trend for French sales is statistically indistinguishable from the control group's time trend prior to HADOPI. Any phenomenon "X" that would have led to the trends we observed in the data would have to meet two conditions: First, it must be true that X began to change differentially in France relative to the control around March/April of 2009, and that X would have an impact on iTunes sales.. Second, X would have to have a stronger effect on Rap/Hip hop sales, and a weaker effect on Jazz/Classical/Christian/Folk sales, than on sales of other genres. For example, if Apple coincidentally began a general iTunes marketing campaign in France in March 2009 (and did not do so in the control countries), this still would be unlikely to explain our observations unless the marketing campaign were primarily directed toward audiences of heavily pirated genres.¹⁸

¹⁸ We note that we are aware of no Apple campaign during this timeframe that was specific to France.

One phenomenon that might meet these criteria would be French adoption of devices connecting to the iTunes music store. For example, if sales of mobile iOS devices in France diverge substantially from sales in the control group countries starting with HADOPI, this in turn might affect iTunes sales, undermining our results.¹⁹ In Appendix I we address this possibility and show that the French adoption rate of iOS mobile devices (devices that connect to the iTunes store) was no higher during this period than it was for the control countries.

7. Discussion

Combined, the four major labels sold an average of 491,000 tracks per week in France after March 2009. Our findings suggest that in the absence of HADOPI, if France followed the same trend as the control group,²⁰ sales would have averaged only 401,000 units per week. Thus, our results suggest that the HADOPI law (and the education and media attention surrounding it) increased iTunes single sales by 90,000 units per week on average. If we assume an average song price of €1 per song, this equates to an increase of €4.7 million (\$6.3 million) in annual iTunes track revenues.

Likewise, average iTunes album sales were about 56,000 units per week in France after March 2009. In the absence of HADOPI, we estimate that iTunes album sales would have averaged only 44,800 albums per week. Thus, our results suggest that HADOPI causally increased French digital albums sales by an average of 11,200 units per week. Assuming an average of €8.5 per album, this equates to an increase of €4.9 million (\$6.7 million) in annual iTunes album sales due to HADOPI. Together, our estimates suggest that HADOPI increased annual iTunes revenues (tracks plus albums) by about €9.6 million (\$13 million)

¹⁹ We note that even in this scenario, the increase in iOS device may be related to HADOPI if HADOPI causes consumers to be more likely to purchase from iTunes and purchasing from iTunes causes consumers to be more likely to purchase iOS devices. Thus, even if sale of iOS devices substantially increased in France around HADOPI, it may not necessarily undermine our estimates. However, as we show in the Appendix, there is no evidence that the uptake of these devices was different in France during this time period.

²⁰ Which they did for the period of time from July 2008 until March 2009.

per year for the four majors combined.²¹ Under the assumption that the four majors make up 70% of the industry, if sales for the remaining 30% of artists experienced the same change in sales as we observe for the major labels, then the impact of HADOPI was to increase overall digital iTunes sales by €13.8 million (\$18.6 million) per year for the entire music industry.

We also note that the effect of HADOPI was larger for more heavily pirated genres like Rap and smaller for less pirated genres like Christian music or Jazz, which is what one would expect if the increase in sales were causally related to HADOPI. It is also worth noting we observed this sales pattern for each of the four majors when analyzed separately, providing some support our assumption that HADOPI may have had a similar effect on the independent labels (the other 30% of the market). Finally, we note that our results do not appear to be explained by an increase in Apple iOS devices in France relative to the control group countries.

The most interesting, and potentially surprising, part of this conclusion is that the study occurs before anyone received a third notice (that is, before any cases have been referred to the criminal court), and that the increase in sales is observed starting with the high level of publicity around the law and continuing with the sending of notices. Thus our study demonstrates that relatively a less heavyhanded policy like HADOPI (which provides education, awareness, and notices before any penalties) can influence consumers to choose paid channels over filesharing.. Disentangling the effects of the positive reinforcement actions like education from the negative reinforcement actions like sanctions is a fruitful area for future research.

A finding that education and salience have a strong impact on user behavior is consistent with those of Bhattacharjee et al. (2008), who find that public awareness of a very small number of major lawsuits conducted by the RIAA against music file-sharers reduced national supply and demand for pirated content on P2P networks. However, unlike Bhattacharjee we do not find that the effect is necessarily short-lived.

²¹ According to sources inside the major labels, about 30% of this goes to iTunes while the other 70% is split in some manner between the label and the artist.

We suggest that with regard to mitigation of sales displacement by piracy, a national anti-piracy policy combined with educational efforts may be much more effective in the longer term than are a small number of high-profile lawsuits.

As noted above, a limitation of our study is that we only observe data for one industry (music), in one channel (iTunes). However, this also means that our study likely understates the true sales impact of HADOPI. A number of other channels provide revenues to the music industry, including cd sales and legal music streaming services – these channels (such as Deezer, MusicMe, or YouTube) may also have benefitted from HADOPI. However, some of these services — such as Deezer — were adopted after the impact of HADOPI, leaving us with no pre-HADOPI period to study. As well, growth trends in these services seem to vary greatly across countries, making selection of a control group difficult at best. Thus \$18.6 million per year represents a lower bound on the total effect that HADOPI had on music industry revenues. It is also quite possible that other industries such as motion pictures or television have been impacted by HADOPI.

For policy-makers, our results may have important implications in other countries that are considering passing similar graduated response laws, as well as in France where a number of parties oppose the continued existence of the law. Likewise, our results may inform industry practice in some countries, like the United States, that have seen the voluntary agreement between the music industry and Internet Service Providers on the application of a graduated response system. Though of course, generalizing our finding to other settings and countries would require additional considerations unique to those settings.

In the context of ongoing policy debates, we also wish to point out two important limitations of our study. First, our results only address the impact of this legislation on industry revenue. Our study does not address the costs of this or other anti-piracy interventions. These costs could accrue either through the direct costs of implementing and enforcing the legislation, or through indirect social costs and potential side-effects associated with implementing such anti-piracy legislation, and such costs should be balanced

against potential benefits in any discussion of policy change. Second, we wish to point out that our results should be viewed only in the context of efforts to influence the demand-side of piracy (that is, through educating and influencing consumer behavior). It is unclear whether efforts to influence the supply-side of piracy (such as site blocking that was envisioned in SOPA) would be similarly effective in altering consumer behavior.

Appendix I: iOS Devices

As noted above, it is possible that our results are explained by other phenomena, however, these phenomena would need to affect French iTunes sales more than sales in the control group countries, would need to do so during the HADOPI timeframe, and would need to have a larger impact on heavily pirated music genres than on less pirated genres.

One possible explanation that would fit these characteristics is that iTunes sales are strongly correlated with the overall penetration of Apple iOS devices (specifically iPhones, iPods, and iPads), and that the penetration of Apple iOS devices increased significantly more in France than in the control group countries during our study period (and that this change was unrelated to HADOPI and that these new iPhone users prefer heavily pirated music genres). In this section we analyze whether sales of iOS devices can explain our results in this way.

To do this, we obtained data from IHS Screen Digest documenting total active iOS devices (iPhones, iPads, and iPods) for each country in our sample for the years 2008, 2009, and 2010. Table A1 displays the total active iOS devices for each country in all three years.

As the table shows, while French adoption of iOS devices increased by 194% from 2008 to 2009 and by 481% from 2008 to 2010, adoption in the control group countries increased by 200% from 2008 to 2009 and by 586% from 2008 to 2010. Thus, if anything, the change in penetration of iOS devices was slightly smaller in France than in the control group countries, a fact inconsistent with the proposed counter explanation for our results.²²

We also note that the increase in iOS penetration in our control group countries is driven by an unusually large increase in iOS penetration in Spain. On one hand because our main results measure changes in French iTunes sales relative to aggregate sales in the control group countries, the relatively large increase

²² We note that because our analysis measures the change in French iTunes sales on a percentage basis relative to the control group countries, percentage changes are the appropriate comparison point to use for whether changes in iOS device penetration can explain our results.

in iOS device sales in Spain does not by itself present a problem for our analysis so long as overall iOS penetration in the control group countries is not significantly smaller than changes in iOS penetration in France.

However, the large increase in iOS device penetration in Spain relative to the other control group countries also presents an opportunity to test whether changes in iOS device penetration is a significant driver of changes in overall iTunes sales. To do this, in Table A2 we report estimates for (ii), except that here we compare the change in iTunes sales in Spain to the control group (not including France), before and after HADOPI. As above, we count March 30, 2009 as the beginning of HADOPI. We note that both track and album sales for the control group (UK, Italy, Belgium, Germany) were increasing during this time. However, in spite of much higher growth in iOS mobile device penetration in Spain than in France or the control countries, iTunes track sales in Spain grew no faster than sales in the control group did, and album sales grew at a slower rate.²³

Thus, we find no evidence that changes in iOS device penetration in France are driving our results. The increase in iOS device penetration in France was, if anything, lower in France relative to the control group countries during the 2008 to 2010 timeframe. Additionally, a comparison of iTunes sales in Spain to sales in the control group countries shows no obvious evidence that changes in iOS device sales are significantly driving changes in iTunes sales.

²³ This finding should not be seen as a test of whether changes in copyright policy in Spain that were enacted in 2011 had any impact. This tests only whether the change in Spanish iTunes sales after March 30, 2009 was greater than the change in the control group over the same period.

References:

- Bertrand, M., E. Duflo, S. Mullainathan. 2004. How Much Should We Trust Differences in Differences Estimates. *The Quarterly Journal of Economics*, 2 249-275.
- 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.
- 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., J. Waldfogel. 2011. Reel Piracy: The Effect of Online Movie Piracy on Film Box Office Sales. Working Paper. Wellesley College. Boston, Massachusetts.
- Hui, K., I. Png. 2003. Piracy and the Legitimate Demand for Recorded Music. *Contributions to Economic Analysis & Policy*, 2(1) Article 11.
- IFPI. 2010. IFPI Response to Commission Green Paper on Creative and Cultural Industries. July 2010.
- Liebowitz, S. 2008. File-Sharing: Creative Destruction or just Plain Destruction? *Journal of Law and Economics*, 49 1-28.
- Liebowitz, S. 2008. Testing File-Sharing's Impact by Examining Record Sales in Cities. *Management Science*, 54(4) 852-859.
- Liebowitz, S. 2011. The Metric is the Message: How much of the Decline in Sound Recording Sales is due to File-Sharing? Working Paper, University of Texas at Dallas.
- Oberholzer, F., K. Strumpf. 2007. The Effect of File Sharing on Record Sales. An Empirical Analysis. *Journal of Political Economy*, 115(1) 1-42.
- Oberholzer, F., K. Strumpf. 2009. Filesharing and Copyright. *NBER's Innovation Policy and the Economy Series*, volume 10. MIT Press.
- OECD. 2009. Magnitude of Counterfeiting and Piracy of Tangible Products. November 2009.
- Peitz, M, P. Waelbroeck. 2004. The Effect of Internet Piracy on Music 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.

Sherman, L. 1990. Police Crackdowns: Initial and Residual Deterrence. *Crime and Justice*, 12 1-48.

Smith, M.D., R. Telang. 2009. Competing with Free: The Impact of Movie Broadcasts on DVD Sales and Internet Piracy. *MIS Quarterly*, 33(2) 321-338.

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. 2005. File Sharing and International Sales of Copyrighted Music: An Empirical Analysis with a Panel of Countries. *Topics in Economic Analysis & Policy*. 5(1) Article 21.

Table 1: Descriptive Statistics

Country	<i>iTunes track unit sales (thousands)</i>			<i>iTunes album unit sales (thousands)</i>		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
Belgium	133.4	130.1	21.3	9.8	9.7	2.2
Germany	728.1	691.6	148.9	87.4	85.0	22.9
Spain	65.7	64.1	11.6	10.1	9.8	2.3
France	447.7	473.9	96.6	49.7	53.4	14.7
Italy	183.9	187.7	37.1	18.7	18.6	4.6
UK	2899.3	2801.9	594.0	270.7	275.2	82.7
Total	743.0	252.6	1022.3	74.4	25.9	98.6

Table 2: Estimate Effects of HADOPI for Tracks, Albums, and Across Genres

	(i)	(ii)	(iii)	(iv)	(v)
	All Tracks	All Albums	Classical / Folk / Modern Christian / Jazz	Rock / Pop	Rap / Hip Hop
After Hadopi	0.228* (0.037)	0.351* [†] (0.033)	-0.042 (0.072)	0.142 (0.068)	0.846* (0.205)
After Hadopi * France	0.203** (0.037) [0.031]	0.223** (0.033) [0.040]	0.068 (0.072) [0.628]	0.158+ [†] (0.068) [0.092]	0.260 (0.205) [0.531]
Constant	12.520* (0.023)	10.168* (0.020)	7.715* (0.044)	11.411* (0.042)	8.731* (0.127)
Observations	918	918	912	912	912
# of Countries	6	6	6	6	6
R-squared	0.361	0.417	0.082 [†]	0.103	0.42

Robust standard errors clustered at country level appear in parentheses

Two-tailed P-values derived from permutation test appear in square brackets

+ significant at 10%; ** significant at 5%; * significant at 1%

March 30, 2009 is counted as the beginning of Hadopi

Columns (i) and (ii) include data from all four majors, while columns (iii) through (v) reflect data from only two.

Table A1: Total Active iOS Devices (in millions)

Country	2008	2009	2010	% increase from 2008 to 2009	% increase from 2008 to 2010
France	1.02	3.02	5.91	197%	481%
Control group avg.				200%	586%
<i>Belgium</i>	<i>0.09</i>	<i>0.25</i>	<i>0.75</i>	<i>194%</i>	<i>776%</i>
<i>Germany</i>	<i>0.96</i>	<i>2.43</i>	<i>5.17</i>	<i>153%</i>	<i>439%</i>
<i>Italy</i>	<i>0.71</i>	<i>1.6</i>	<i>2.96</i>	<i>125%</i>	<i>317%</i>
<i>Spain</i>	<i>0.33</i>	<i>1.72</i>	<i>3.93</i>	<i>416%</i>	<i>1076%</i>
<i>UK</i>	<i>2.1</i>	<i>4.45</i>	<i>8.87</i>	<i>113%</i>	<i>323%</i>

Includes iPads, iPods, and iPhones capable of accessing Apple's App Store.

Table A2: iTunes Sales Before and After HADOPI, Spain vs. Control Group

	(i)	(ii)
	All Tracks	All Albums
After Hadopi	0.223* (0.047)	0.376* (0.026)
After Hadopi * Spain	0.027 (0.047)	-0.129* (0.026)
Constant	12.490* (0.028)	10.134* (0.016)
Observations	765	765
# of Countries	5	5
R-squared	0.277	0.371

Robust standard errors in parentheses

+ significant at 10%; ** significant at 5%; * significant at 1%

Figure 1: iTunes Single Track Unit Sales Trends (4 majors combined), France vs. Control

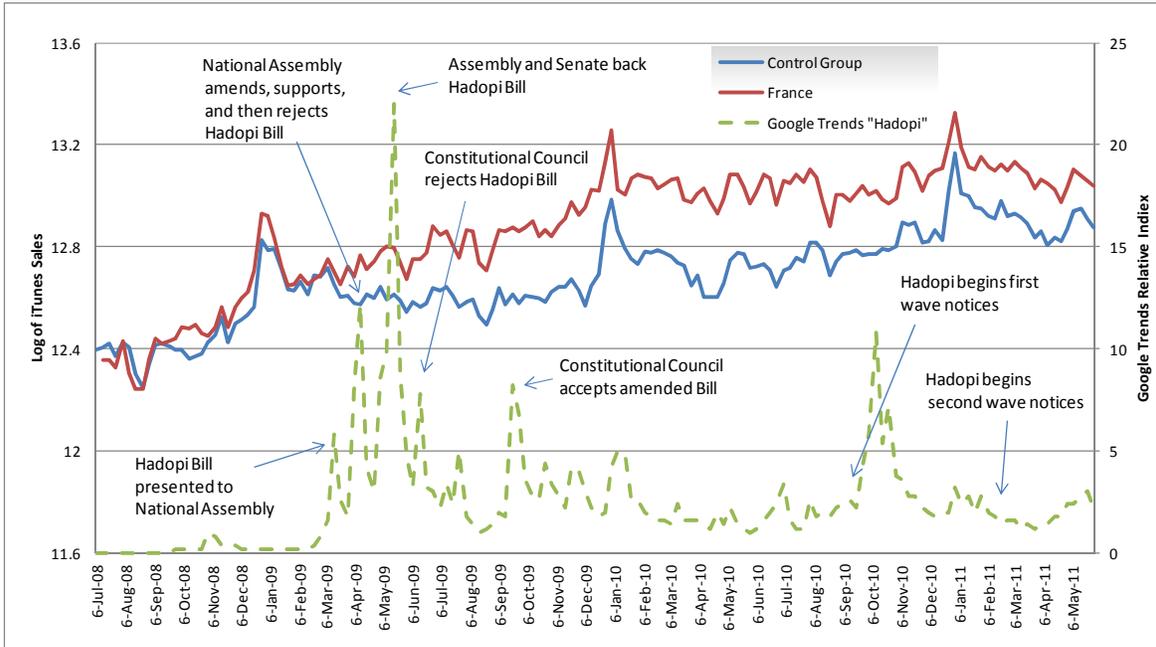


Figure 2: iTunes Album Unit Sales Trends (4 majors combined), France vs. Control

