Information Economics and Policy 37 (2016) 61–66

Contents lists available at ScienceDirect

Information Economics and Policy

journal homepage: www.elsevier.com/locate/iep

# The effect of file sharing on record sales, revisited[☆]

Felix Oberholzer-Gee[a,*], Koleman Strumpf[b]

[a] Harvard Business School, MA, United States
[b] University of Kansas School of Business, KS, United States

ARTICLE INFO

*Article history:*
Received 13 October 2016
Available online 10 November 2016

ABSTRACT

Even as we approach the twentieth anniversary of widespread file sharing, its impact on the sale of copyrighted material remains in dispute. We contributed to this debate with an early study, "The Effect of File Sharing on Record Sales: An Empirical Analysis," that was published in the Journal of Political Economy in 2007. Perhaps surprisingly, we found that piracy contributed to the decline in music sales but was not the main cause. In this article, we review and respond to recent criticism of our work by Stan Liebowitz in Econ Journal Watch. We show how the use of proxies for file sharing can result in misleading conclusions. We close by reviewing what we know about the impact of file sharing on record sales today. In our view, new music formats are an important if understudied channel through which changes in technology influence the demand for entertainment.



## 1. Introduction

Today's technology infrastructure greatly facilitates the collection and dissemination of information. One consequence of the ease with which we can access and share information is a substantial weakening of copyright protection (Waldfogel, 2012a). While there is a perception that the unauthorized use of copyrighted material is diminishing, the volume of material accessed via file sharing has increased by over 40% between 2008 and 2014 (calculated from Cisco, various years). This raises the question whether illegal use diminishes sales and hence the incentives to create new works (Waldfogel, 2012b and 2012c). We contributed to this literature ten years ago with an article titled "The Effect of File Sharing on Record Sales: An Empirical Analysis" (Oberholzer-Gee and Strumpf, 2007) in which we studied the effect of piracy on record sales in the United States at a time when piracy was in its infancy. To the surprise of many, we found that piracy had a limited impact, reducing record sales by no more than three percent, less than one-third of the sales decline that we observed in 2002.[1]

Many authors participated in the ensuing debate,[2] some finding limited effects while others provided evidence of a greater role of piracy in explaining the observed decline in record sales (for surveys, see Oberholzer-Gee and Strumpf 2010 and Waldfogel 2012a). Some researchers questioned our approach more directly, arguing that we failed to identify a more significant impact of file sharing due to methodological shortcomings (Liebowitz 2016a and 2016b; Rob and Waldfogel 2006; Smith and Telang 2012).

In this paper, we seek to address the concerns raised by other scholars, in particular the recent critique by Stan Liebowitz in Econ Journal Watch (Liebowitz, 2016a). In our discussion, we emphasize the importance of observing file sharing directly, as has become more common in recent contributions to the literature (Aguiar and Martens, 2016; Hammond, 2014; Lee, 2016), and the role of unobserved time-varying heterogeneity. We conclude by discussing what we know about the impact of file sharing as a result of a decade of academic research.

## 2. The debate

In discussions of our study, scholars raised a number of issues, including the reasonableness of the file-sharing data that we employ, the validity of the main instrument, and the plausibility of

☆ We thank three anonymous reviewers and our editor, Lisa M. George, for helpful comments on a previous version.

* Corresponding author.
*E-mail address:* foberholzer@hbs.edu (F. Oberholzer-Gee).

[1] The theory literature points out that piracy need not negatively impact the owners of intellectual property. In particular file sharing facilitates sampling which may stimulate new purchases. See for example Shapiro and Varian (1999) and Peitz and Waelbroeck (2006).

[2] Contributions to the literature include Adermon and Liang, 2014; Aguiar and Martens 2016, Andersen and Frenz 2010, Bhattacharjee et al. 2007, Danaher et al., 2014; Hong 2013, Lee 2016, Liebowitz 2006 and 2008, Michel 2006, Peitz and Waelbroeck 2004 and 2006, Rob and Waldfogel 2006, Waldfogel 2010, Waldfogel 2012b and 2012c, Zentner 2006. We assess the various approaches taken in these papers in the final section.

http://dx.doi.org/10.1016/j.infoecopol.2016.11.001
0167-6245/

the published estimates (see Liebowitz, 2016a; Rob and Waldfogel, 2006; Smith and Telang, 2012). We will discuss these questions in turn.

In the discussion below we consider four issues related to the data and the instrument used in our main analysis, and two more considering quasi-experiments which use more aggregate data. The original paper also considers a wide variety of alternative instruments, identification strategies, and robustness checks which are often ignored in critical assessments but which are consistent with our main empirical results.

### *2.1. Do our data exhibit too much variability?*

Our file-sharing data came from the logs of two OpenNap servers, which operated continuously for 17 weeks from September 8 to December 31, 2002. Some commentators wondered why our sharing data look much more variable than one might expect. In fact, a naïve comparison would show that our measure of file sharing appears to increase sharply over time while overall weekly downloads, as measured by other sources, remain fairly stable (Liebowitz 2016a, p. 380). The difference is not difficult to understand. Our analysis omits album-weeks during which an album has not yet been released. Including such weeks would bias our estimates because downloads and sales are zero by definition. The omission is empirically important as more than one out of eight album-weeks is not included for this reason. There are two implications. First, the effective download rate in our data differs from the total weekly number of downloads, and this difference is greater in the early weeks of the sample when many albums had not yet been released. Second, much of the steep rise in the number of downloads in our data is due to changes in the number of albums, but this variation plays no role in our identification strategy as we employ album fixed effects.

Finally, when comparing our data to other sources, it is worth keeping in mind just how poor the quality of much of the available data is. BigChampagne, a popular source of data, observes the supply of files, not downloads. Similarly, NPD data come from a panel of 40,000 individuals who willingly agree to have their PCs monitored. This group of individuals is unlikely to be representative of the overall P2P population.[3]

A better way to see if the temporal variation in our data tracks the variation in file sharing more generally is to compare our data to download activity on the Internet2 backbone (Internet2 Netflow Statistics, 2004). Netflow counts all packets due to file sharing. Since our sample predates the period of video downloads, illegal transfers should be roughly linear in the number of music downloads. After normalizing by the number of album-weeks, we find the correlation between our weekly downloads and the number of file sharing packets on Internet2 to be 0.49, which suggests the temporal variation in our data matches that of overall US file sharing.

### *2.2. Is our main instrument valid?*

To address questions of simultaneity – albums that are popular among file sharers are likely to sell well – we studied how school vacations in Germany influenced downloading activity in the United States. In 2002, German students were an important source of files for the United States. A natural question to ask is whether this instrument is valid. One way to check this assertion is to see whether an increase in the popularity of music in the U.S. magnifies the effect of the German supply shock. To run this test, we interact the vacation variable with an album's rank on the U.S. MTV charts. We find a positive coefficient on the interaction in the first stage – increases in popularity heighten the effect of German vacations on U.S. downloads – but the coefficient is statistically insignificant and economically small (Oberholzer-Gee and Strumpf 2007, Table 8). We can run this type of test with other measures of music popularity (e.g., Billboard Airplay) and find similar results. Even better, we can allow the effect of U.S. popularity on downloads and sales to vary by album. We find that these three-way interactions (German kids×MTV×Album FE) are individually and collectively insignificant in the sales equation. As these results show, there is no consistent and predictable relation between the vacation variable and U.S. sales.

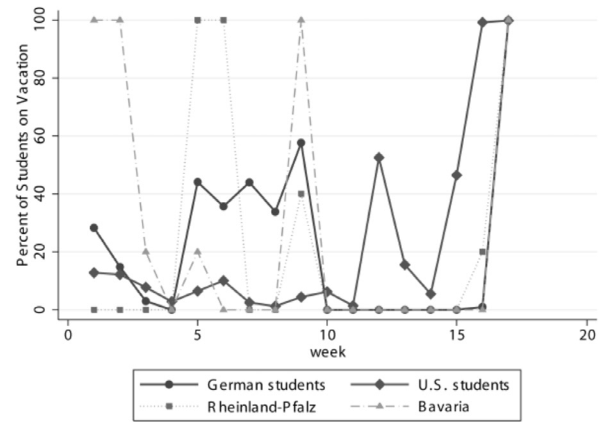


**Fig 1.** Timing of German and US vacations reproduced from Oberholzer-Gee and Strumpf (2007).

The misperception may have been the result of comparing changes in weekly U.S. album sales and the share of German students on vacation (Liebowitz 2010; Fig 1). Sales rise throughout the fall and peak before the holidays. Similarly, all German students are off around Christmas. But this general impression is misleading: The direction of the correlation between sales and students on vacation changes from period to period. It is negative over the entire study period (corr = −0.023) and positive for the second half (corr = +0.093). More importantly, this is not the variation that we study in our models, which include a polynomial time trend (or week fixed effects) and album fixed effects. Even if the raw data suggested a systematic correlation, our results are identified from supply shocks, some of them album-specific, relative to a common time trend. Similarly, it is straightforward to make sure that our results do not reflect the Christmas period. If we omit the December data from our sample, the effect of downloads on sales remains statistically insignificant.

### *2.3. Are our first-stage estimates plausible?*

The first-stage estimates in our analysis – the influence of German vacations on downloads in the U.S. – seem large, raising the question whether the instrument proxies for an unobservable. In assessing this possibility, it is important to keep in mind that downloads are heavily skewed: the median number of downloads for an album-week is zero, and it is the heavily downloaded album-weeks which drive the first-stage estimates. For a change in the number of kids on vacation from its minimum to its maximum, we find an additional 8.3 downloads. This value is smaller than the

[3] Liebowitz (2016a) also compares our download data to information from a private BitTorrent tracker. This comparison is complicated by differences in technology and timing – our data are from OpenNap servers which operated six years prior to the private tracker – and sample composition. Unlike our data, the private BitTorrent tracker is neither representative of all BitTorrent users nor of all of the content which is available online.

mean number of downloads in album-weeks which have non-zero downloads, and it seems quite plausible.[4]

A more formal approach to resolving the issue of spurious correlation is to add week fixed effects to our models – the main effect of vacations is now subsumed in the time fixed effect – and interact additional album-specific, time-varying instruments with the vacation variable. One such instrument is whether the band which recorded the album is on tour in Germany that week. These specifications, which provide similar results, use a valid instrument even if spurious correlation were an issue in the specifications with a polynomial time trend (Oberholzer-Gee and Strumpf 2007, Table 6).

We gain further confidence in our main instrument by exploring the mechanisms that link German vacations and U.S. downloads. For example, we can show that increases in the German vacation instrument as well as the vacation-concert interaction reduced the time it took to download a file. Unlike today, in 2002 it was time-consuming to transfer files, and many transfer requests remained incomplete. The many shortcomings of peer-to-peer file-sharing systems at the time of our study are one reason why file sharers responded flexibly to supply shocks. In our empirical analyses, we find that users download more files when supply conditions improve temporarily (Oberholzer-Gee and Strumpf 2007, Table 6). We also provide supplemental evidence for this mechanism, showing that the time effect is largest for albums and genres which are popular in Germany, and that similar time savings occur in other countries whose times zones are complementary to Germany.

### *2.4. Are german users overrepresented in our data?*

BigChampagne data (cited by the OECD and taken up in discussions of our paper) suggest that we have too many German users in our data. This impression is false. As discussed above, BigChampagne reports the number of users who share files, not the number of users who download content. This is an important distinction: many users downloaded music but did not share files. BigChampagne is also not a reliable source for global activity. For instance, the company did not include users from networks popular in Japan (OECD, 2004, p.189), an important source of files. The Japanese share of users is 8% in our paper and an implausible 0.7% in the BigChampagne data.

Better sources of data confirm the importance of German users. For example, Expand Network, a file sharing monitoring company, passively observed the KaZaA network during the same weeks that we examine (Leibowitz et al., 2002). The U.S. share of file sharers is 34.1% in the Expand Network and 30.9% in our study; the German share is 12.6% according to Expand and 13.5% in our analysis.

### *2.5. How significant is the vacation-related supply shock?*

In a recent contribution, Liebowitz (2016a: 391) argued that the number of students influenced by holidays was too small to plausibly have an effect on the availability of files in the United States.[5] His calculation assumes that only students between the ages of 12 and 17 are affected by school holidays, and that children in Germany have “very limited” interest in English language songs because most do not speak English. Both assumptions are incorrect. German vacation schedules remain relevant for students older than 17, in particular the millions of students who attend professional schools (Berufsschulen) and those who attend Gymnasium.[6] In 2002, more than 70% of students in professional schools were older than 17 and 36.8% were older than 20. More than 350,000 students in gymnasium were older than 17. A second misconception is that students younger than 12 years of age do not share files.[7] In the 6–13 age group, 36% of German students used a computer at least once a week in 2002. 17% report they use a PC every day. 77% say they are “interested” or “very interested” in music. In a survey of this age group, 11% said they downloaded music at least once a week. For these reasons, our data include students in three educational tiers: Primarbereich (ages 6–10), Sekundarbereich I (ages 10–16), and Sekundarbereich II (ages 15–20).[8] We do not include university students because German universities follow a different vacations schedule.

There is also a strong overlap in musical taste in the U.S. and in Germany.[9] Popular music in Germany is generally dominated by U.S. and international (mostly British) artists. There is surprisingly little German music. For instance, on German radio stations, only 6% of songs are by German artists. During our sample period, only 24% of all positions on the German Top 100 album charts were occupied by German artists.[10] Because file-sharing activity is heavily focused on popular albums and songs, it is important that Germans supply this type of song. Of the albums that entered our sample via the U.S. Billboard 200, 62.65% are also on the top 100 German charts. There is also a substantial overlap between the availability and sales ranking of music titles in our sample on Amazon U.S. and Amazon Germany or JPC, a leading German music retailer.[11] As we explain in the published article, we conducted Wilcoxon matched-pairs signed-ranks tests to compare the Amazon rankings in the two countries. It is not possible to reject the null of equal distributions for the entire sample. In genre-by-genre comparisons, equality is rejected only for Latin and Country music.

Perhaps most importantly, the data used in Liebowitz's calculation do not measure file-sharing. The relevant Pew Internet survey question reads: “Ever downloaded music files onto your computer so you can play them at anytime you want?” In the survey, 19% of respondents over 65 say they downloaded music onto their computer. It seems likely that this age group, and perhaps others, was thinking of copying CDs to their PC when it answered the question. (18% of respondents 65 and older did not have access to the Internet at the time of the survey.) Multiplying the two values (19%x18%), as Liebowitz's calculation does, will not give us an accurate measure of file sharing. If we assume, more plausibly, that file sharing in the U.S. is essentially an activity of teens and college students (adults younger than 24), the share of 12–17 year olds rises to 50%. And this share further increases when we take into account that the data for the youngest group comes from 2000, a full two years before the survey of adults. This is a period of rapid

[4] This point follows from basic econometrics. When the dependent variable is skewed as in our data, the change in the regression fitted value as we move from the max to the min can roughly equal, or even exceed, the dependent variable mean. For example, when the data are (x,y) = {(1,3), (1,0), (1,0), (2,24), (2,0), (2,0)} the mean of y is 4.5 while the difference in fitted value as x goes from one to two is 7.

[5] Based on a series of calculations, he argues that pirate-file availability to Americans caused by a typical German school holiday was a mere 0.14 percent.

[6] Data in this section come from Statistisches Bundesamt (2005). “Bildung im Zahlenspiegel.” Wiesbaden.

[7] Sabine Feierabend and Walter Klingler (2002). “KIM-Studie 2002: Kinder und Medien Computer und Internet. Basisuntersuchung zum Medienumgang 6- bis 13-Jaehriger” Baden-Baden.

[8] Sekretariat der Ständigen Konferenz der Kultusminister der Länder in der Bundesrepublik Deutschland (2002). “Vorausberechnung der Schüler- und Absolventenzahlen 2000 bis 2020.” Statistische Veröffentlichungen der Kultusministerkonferenz Nr. 162, Bonn.

[9] For a more general discussion of US music exports to Europe, see Ferreira and Waldfogel (2013).

[10] The album chart information used in our calculations is taken from www.musikmarkt.de.

[11] All data were collected in May 2006. Amazon data come from www.amazon.com and www.amazon.de, respectively. Information for the German retailer JPC is taken from www.jpc.de

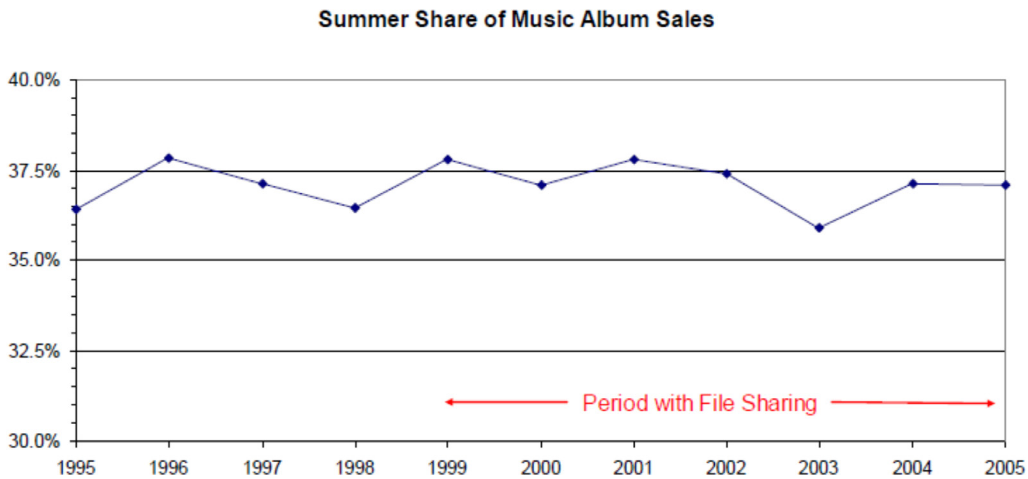


**Fig 2.** Summer share of music album sales.

growth in the number of file sharers and the value used in the calculation of the size of the supply shock surely understates the true fraction as a result.[12]

Liebowitz further reduces the size of the German shock in his calculation by taking into account that in a typical week 20%–40% of German kids are out of school, arguing the youth share influenced by the vacation instrument should be deflated correspondingly. This reasoning is incorrect since it is the variation in the percent of kids on vacation, not the typical or average value, which makes vacations a suitable instrument. As the figure below shows, the vacation share is quite volatile and bounces around the entire range of 0% to 100%. There are even two consecutive weeks where no kids and then all kids are on vacation.[13]

As this discussion shows, the claim that only a miniscule share of files was supplied by German students and that there is little variation in this type of supply is incorrect. One reason for the influence of our original paper is that we actually measure file sharing instead of relying on difficult-to-interpret proxies and back-of-the-envelope calculations.

### *2.6. File sharing and sales during summer*

While our main results come from our panel data, we look for additional evidence from a number of quasi experiments. The first of these considers the effect of changes in the number of file sharers during the summer months. We show that file sharing declines during summer months when many students leave their colleges. In a review of our quasi experiments, Liebowitz (2016b) argued that the summer of 2003 should be disregarded because it followed the announcement of the first RIAA lawsuits. This is a misunderstanding. What the data show (see Fig 2 below) is that record sales are relatively insensitive to drops in the number of file sharers regardless of the source of variation, whether it is due to college students leaving campus or lawsuits.

While 2003 provides the strongest experiment, the summers of 2004 and 2005 are interesting as well. At a time of rapid growth in the number of file sharers – the number of US users doubled between January 2004 and January 2006 – these summer months represent clear breaks from the growth trend in this period. We can show this more formally using regression analysis, analyzing the 2002–2006 period in a specification which includes a time trend term and an indicator variable for summer. The regression results imply file sharing activity dropped 12% in the summer when we include a time trend and fell by 8% when we do not. Despite these changes, however, the share of summer sales remains the same before and after the advent of file sharing.

### *2.7. Change in the number of file sharers*

A final question is how to infer the displacement from piracy based on regressing sales on the number of file sharers. A common approach is to calculate this as (parameter estimate)×(number of users). For example, Liebowitz (2016b) uses this approach to argue that the correct sales impact is at least twice as large as in our estimates. Unfortunately, these calculations use the 2004 number of file sharers (5.5 m users) which is incorrect.[14] To determine the cumulative sales reduction, one should use the actual number of users each month and then sum the total. Equivalently, one could use the mean number of users (5.04 m users). That means the calculated numbers are too large since 5.04 m/5.5 m = 0.92. But even the 5.04 m significantly overstate the effect, since this is the average over the observation period (August 2002-May 2006) and the calculation is supposed to cumulate back to 2000 when file sharing began. The very early number of users is unavailable, but a simple approximation is to interpolate starting with zero in December 1999 and ending with the earliest observed value (3.5 m, the total in August 2002). Using this interpolation the average number of users over January 2000-May 2006 is 3.72 m, so the Liebowitz values are overstated by almost a third as 3.72 m/5.5m = 0.68.[15]

[12] Even if the critique that we overstate the number of German file sharers were correct, the entire line of argument is irrelevant in our application because the size of the group used in the construction of the vacations variable will not bias our estimates. Selecting too large a group can lead to a weak-instruments problem. But concerns associated with the use of weak instruments, as the published results clearly show, do not apply to our case.

[13] Liebowitz applies a final correction to the relevant share of German files by noting that vacations will have no impact on weekends. However, because the evening hours are important for our purposes (this is during the day in the US), a Sunday followed by a day in school is likely to be different from a Sunday during the holidays. In addition, many students do homework on weekends, making it likely that vacations free up additional time on Saturday and Sunday.

[14] The source of the estimated number of file sharers is BigChampagne.

[15] There are additional problems with Liebowitz's calculation. For example, he bases his numbers on a specification which includes unemployment as an explanatory variable. But this is highly negatively correlated with the number of file sharing users (corr = −0.87). This specification suffers from multicollinearity and we know that interpreting parameters on nearly perfectly correlated variables is inappropriate. When unemployment is omitted, the parameter on file sharing users drops significantly in size.

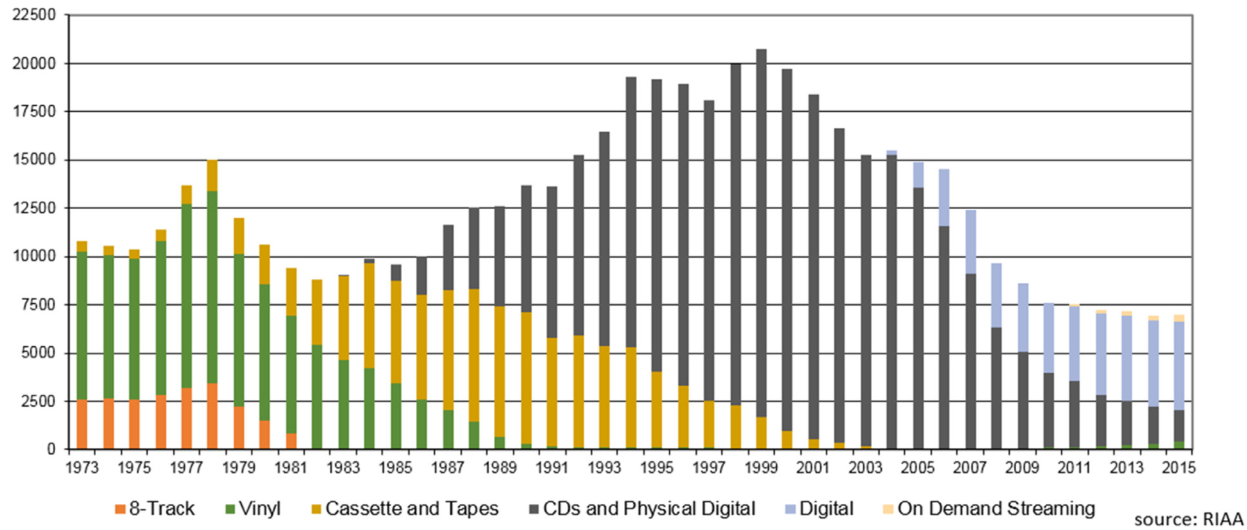


**Fig 3.** US recorded music revenue, by format (million 2015 $).

## 3. How significant was the impact of file sharing, really?

With the benefit of hindsight and more than a decade of research into the question, what can we say about the effect of file sharing on record sales? We have five observations.

First, most studies find that file sharing displaced some sales but the estimated effects vary dramatically, ranging from a negligible (or even positive) influence to a sales displacement rate as high as 30% (Zentner 2006). A typical estimate is a rate of about 20% (Oberholzer-Gee and Strumpf, 2010). In view of the dramatic decline in music sales in the 2000 s, one implication of these findings is that changes in the industry other than file sharing must have had a significant influence on sales.

Second, file sharing research typically follows two designs, a product-level approach that links observations of illegal downloads to sales and a person-focused design that compares paid and unpaid consumption. As Waldfogel (2012a) points out, both approaches can lead to valid estimates of sales displacement if the researchers successfully tackle the issue of unobserved heterogeneity across products (under the first approach) or people (under the second). In practice, however, the two approaches lead to systematically different results. The studies that observe actual file sharing under a product-level approach tend to find no or little displacement (Aguiar and Martens, 2016; Bhattacharjee et al., 2007; Hammond, 2014; Lee, 2016; Oberholzer-Gee and Strumpf, 2007; Tanaka 2004). By contrast, studies that rely on surveys, the person-specific approach, document far greater negative effects. Papers that look at changes in legal constraints (e.g., Danaher et al., 2014; Adermon and Liang, 2014) also tend to find short-term displacement. It is not immediately obvious why the various approaches lead to such different results. One possibility is that person-specific datasets exhibit more time-varying unobserved heterogeneity. Comparing difference-in-difference estimates – a popular means to strip out unobserved differences that do not change over time – with models that explicitly account for unobserved changes in cohort characteristics, Hong (2013) finds that the former attribute the entire decline in record sales to file sharing. Once time-varying changes in unobserved heterogeneity are taken into account, however, the sales displacement rate declines from 100% to a mere 20%.[16]

Third, it seems natural to study the effect of file sharing by comparing sales in the 1990 s and in the 2000 s. The prima facie case for the negative impact of file sharing on the music industry has always been the timing coincidence: as Napster and its later descendants KaZaA and BitTorrent took hold, aggregate sales declined. But this view overlooks just how unusual the 1990 s were. As Fig 3 shows, it is not the 2000 s which are the anomaly but the 1990 s. Real revenues increased by about 75% during this decade after almost no sustained increase at all over the prior twenty years.

Fourth, it is an often forgotten fact that long and painful declines in recorded music sales are nothing new. Fig 3 shows that changes in music industry revenue in the post-Napster era (Napster was launched in summer 1999) bear substantial similarity to changes in revenues in the late 1970 s to early 1980 s. In both periods, inflation-adjusted sales fell steadily by about thirty to forty percent (the 1970 s also had an anomalous run up in sales prior to the decline just as occurred in the 1990 s).

Fifth, these swings appear to be related to the decline of the leading format for recorded music at the time, vinyl in the late 1970 s, cassettes in the 1980 s, and CDs in the early 2000 s. If we normalize the format-specific peak in revenue, it is striking to see the similarities in the rise and the decline for each format (see Fig 4).

There are two reasons why such format changes can lead to sales declines. First there is a delay in demand as consumers slowly transition to the newer format. Second, some of the build up to the peak involves consumers repurchasing music they owned in the earlier format. This effect does not come into play until a new format is dominant. In the most recent period, consumers experienced two format transitions, first to digital downloads, and then to streaming. The speed with which these changes occurred reduced the incentive to build a new library, and streaming as a technology does not require it by definition. This might be one reason why sales have not increased more in recent years, as Fig 3 indicates typically accompany previous transitions to new formats.

[16] There are other problems with survey data. Respondents are unlikely to be truthful when discussing illegal activities, and those who volunteer answers they are unlikely to be representative of the actual file sharing population. Moreover, teens, the primary file sharing users, are often omitted from these surveys. Many of the papers cited in note 3 consider yet other approaches which have important deficiencies. Some use very poor proxies for file sharing activity (broadband rates at the country- or city-level) and others are event studies, examining sales before and after the implementation of greater punishments for file sharers. But laws are endogenous and their passage reflects the relative strength of copyright owners.

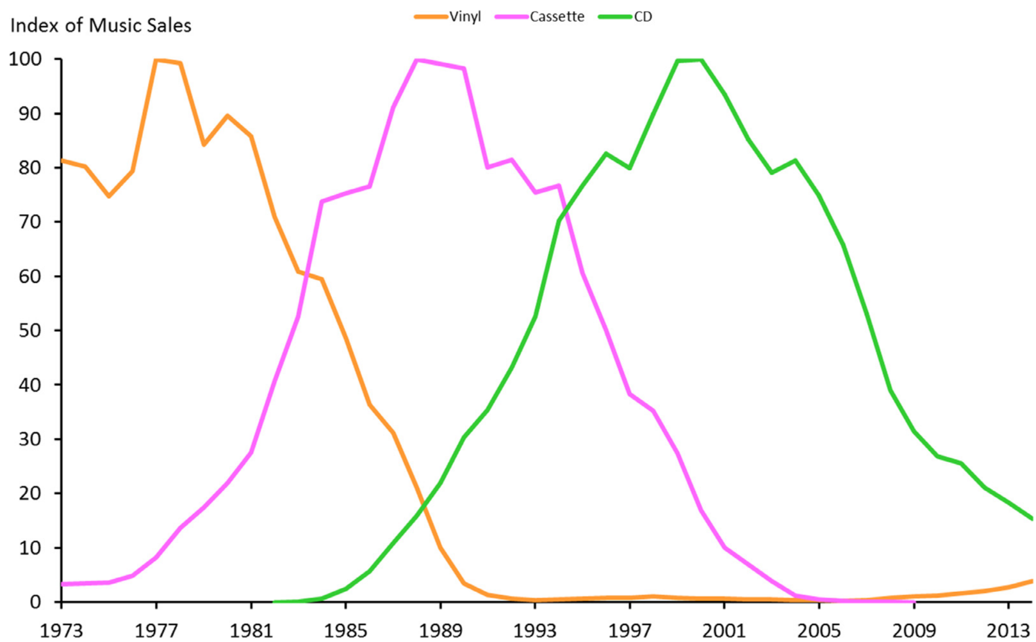


**Fig 4.** Music sales across formats reproduced from Anand (2016) source: RIAA. Rebase peak sales for each format to 100.

In some ways the industry has come full circle. In 2016 sales have stabilized. But, as we pointed out in the introduction, file sharing has hardly disappeared. While there are many explanations for the sharp changes in the recorded music industry over the last twenty years, our sense is that file sharing is but one small facet.

## References

Adermon, A., Liang, C.-Y., 2014. Piracy and music sales: the effects of an anti-piracy law. J. Econ. Behav. Organ. 105, 90–106.
Aguiar, L., Martens, B., 2016. Digital music consumption on the internet: evidence from clickstream data. Inf. Econ. Policy 34, 27–43.
Anand, B., 2016. The Content Trap: A Strategist's Guide to Digital Change. Random House.
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. J. Evol. Econ. 20 (5), 715–740.
Bhattacharjee, S., Gopal, R.D., Lertwachara, K., Marsden, J.R., Telang, R., 2007. The effect of digital sharing technologies on music markets: a survival analysis of albums on ranking charts. Manag. Sci. 53 (9), 1359–1374.
Cisco. various years. Cisco Visual Networking Index: Forecast and Methodology. http://www.cisco.com.
Danaher, B., Smith, M.D., Telang, R., Chen, S., 2014. The effect of graduated response anti-piracy laws on music sales: evidence from an event study in france. J. Indus. Econ. 62 (3), 541–553.
Ferreira, F., Waldfogel, J., 2013. Pop internationalism: has half a century of world music trade displaced local culture? Econ. J. 123 (569), 634–664.
Hammond, R., 2014. Profit leak? Pre-release file sharing and the music industry. South. Econ. J. 81 (2), 387–408.
Hong, S.-H., 2013. measuring the effect of napster on recorded music sales: difference-in-differences estimates under compositional changes. J. Appl. Econom. 28 (2), 297–324.
Internet2 Netflow Statistics, 2004. Internet2 NetFlow: Weekly Reports http://netflow.internet2.edu/weekly/.
Lee, J., 2016. Purchase, Pirate, Publicize: The Effect of Private-Network File Sharing on Album Sales. Queen's Economics Department Working Paper No. 1354.
Leibowitz, N., Bergman, A., Ben-Shaul, R., Shavit, A., 2002. Are file swapping networks cacheable? Characterizing P2P traffic. Expand Networks working paper presented at the 7th International Workshop on Web Content Caching and Distribution, Boulder, August 14–16.
Liebowitz, S.J., 2006. File sharing: creative destruction or just plain destruction? J. Law Econ. 49, 1–28.
Liebowitz, S.J., 2008. Testing file sharing's impact on music album sales in cities. Manag. Sci. 54 (4), 852–859.
Liebowitz, S.J., 2010. The Oberholzer-Gee/Strumpf File-Sharing Instrument Fails the Laugh Test SSRN working paper 1598037.
Liebowitz, S.J., 2016a. Why the oberholzer-gee/strumpf article on file sharing is not credible. Econ J. Watch 13 (3), 373–396.
Liebowitz, S.J., 2016b. Replicating Four 'Quasi-Experiments' and Three Facts from Oberholzergee/Strumpf's Piracy Article SSRN working paper 2777216.
Michel, N.J., 2006. The impact of digital file sharing on the music industry: an empirical analysis. Top. Econ. Analy. Policy 6, 1–22.
Oberholzer-Gee, F., Strumpf, K., 2007. The effect of file sharing on record sales: an empirical analysis. J. Polit. Econ. 115 (1), 1–42.
Oberholzer-Gee, F., Strumpf, K., 2010. File Sharing and Copyright. In: Lerner, J., Stern, S. (Eds.), Innovation Policy and the Economy 10. National Bureau of Economic Research, Chicago, pp. 19–55.
Peitz, M., Waelbroeck, P., 2004. The effect of internet piracy on music sales: cross–section evidence. Rev. Econ. Res. Copyright Issues 1 (2), 71–79.
Peitz, M., Waelbroeck, P., 2006. "Why the music industry may gain from free downloading – the role of sampling. Int. J. Ind. Organiz. 24 (5), 907–913.
Rob, R., Waldfogel, J., 2006. Piracy on the high c's: music downloading, sales displacement, and social welfare in a sample of college students. J. Law Econ. 49 (1), 29–62.
Shapiro, C., Varian, H., 1999. Information Rules: A Strategic Guide to the Network Economy. Harvard Business School Press.
Smith, M., Telang, R., 2012. Assessing The Academic Literature Regarding the Impact of Media Piracy on Sales SSRN working paper 2132153.
Tanaka, T., 2004. Does file-sharing reduce CD sales? A case of japan. Paper prepared for Conference on IT Innovation.
Waldfogel, J., 2010. Music file sharing and sales displacement in the iTunes era. Inf. Econ. Policy 22 (4), 306–314.
Waldfogel, J., 2012a. Digital Piracy: Empirics. In: Peitz, M., Waldfogel, J. (Eds.), The Oxford Handbook of the Digital Economy. Oxford University Press, Oxford, pp. 531–546.
Waldfogel, J., 2012b. Copyright protection, technological change, and the quality of new products: evidence from recorded music since napster. J. Law Econ. 55 (4), 715–740.
Waldfogel, J., 2012c. Copyright research in the digital age: moving from piracy to the supply of new products. Am. Econ. Rev. 102 (3), 337–342.
Zentner, A., 2006. Measuring the effect of music downloads on music purchases. J. Law Econ. 49 (1), 63–90.