Economics
The Open-Access, Open-Assessment E-Journal

Vol. 11, 2017-13 | May 18, 2017 | http://dx.doi.org/10.5018/economics-ejournal.ja.2017-13

A Replication of Four Quasi-Experiments and Three Facts from 'The Effect of File Sharing on Record Sales: An Empirical Analysis' (Journal of Political Economy, 2007)

Stan J. Liebowitz

Abstract
The influential piracy paper by Professors Oberholzer-Gee and Strumpf, although mainly based on proprietary data, contained an "important complement" to the main results, consisting of four "quasi-experiments" using publicly available data. This replication examines all of these quasi-experiments by using identical data and statistical methods where possible, as well as sometimes extending or augmenting the data or methods. This study concludes that the quasi-experiments performed by OS each contain important errors, oversights or inconsistencies that most often, but not completely, overturn the results claimed in the original OS article.

(Published as Replication Study)

JEL Z1 O3 L8
Keywords Piracy; file sharing; sound recordings; music; iTunes; Napster

Authors
Stan J. Liebowitz, ✉ School of Management, University of Texas at Dallas, USA, liebowit@utdallas.edu

Citation Stan J. Liebowitz (2017). A Replication of Four Quasi-Experiments and Three Facts from 'The Effect of File Sharing on Record Sales: An Empirical Analysis' (Journal of Political Economy, 2007). *Economics: The Open-Access, Open-Assessment E-Journal*, 11 (2017-13): 1–21. http://dx.doi.org/10.5018/economics-ejournal.ja.2017-13

Received September 30, 2016 **Published as Economics Discussion Paper** December 5, 2016
Revised May 9, 2017 **Accepted** May 10, 2017 **Published** May 18, 2017

Introduction

One of, if not the most influential article on the effects of digital piracy (file-sharing) on the sound recording industry, appeared as a lead article in the *Journal of Political Economy* in 2007.[1] Its authors, Felix Oberholzer-Gee and Koleman Strumpf (OS), using data obtained from a small pirate server, concluded that piracy had essentially no impact on record sales. Although I have elsewhere (Liebowitz, 2016b) discussed issues with their main data set and analyses, direct replication of their main regressions is not possible because OS never made their main data public.[2]

Nevertheless, OS used publicly available data when conducting four "quasi-experiments" described as "an important complement" to their main analysis. In this replication I examine these quasi-experiments and also several pieces of data that OS use to support their analysis. Note that additional details, results, and instructions on replicating my results can be found in an Appendix, to which I will sometimes refer.

1 Does American piracy decrease every summer?

Here are OS describing their first quasi-experiment:

> The first experiment involves variation over time. The number of filesharing users in the United States drops 12 percent over the summer (estimated from BigChampagne [2006]) because college students are away from their high-speed campus Internet connections. If downloads crowd out sales, we should observe that the share of albums sold in the summer increases following the advent of file sharing. [2007, page 36]

OS purport to demonstrate that American piracy fell each summer, supposedly because American college students lost access to their high-speed Internet connections when they went home for the summer.[3] Thus any impact of piracy, such as harming record sales, should have been weaker in the summer than during the rest of the year. OS then compared the yearly summer shares of record sales for four pre-Napster years and seven post-Napster years, expecting the summer share of record sales to be higher in the post-Napster years if piracy harmed sales, everything else equal.

---

[1] According to Google Scholar, Oberholzer-Gee and Strumpf (2007) has received more citations (1004 as of September 29th, 2016) than any other paper on piracy in the economics literature, of which I am aware. According to Web of Science (in May 2016), it was the fifth most cited paper that the Journal of Political Economy published in 2007.

[2] At a public forum in 2004, Koleman Strumpf stated they would make their data set available "soon" [http://web.archive.org/web/20040515202021/http:/cdbaby.net/fom/000004.html, search for "soon"]. Four years after this statement, OS stated that they had signed a non-disclosure agreement, although they did not provide this agreement to reporters who asked to see it. See Glenn (2008) or Häring (2008).

[3] Contrary to OS's claim, most college students would not lose high speed access during the summer in the U.S. since most college students do not live in dormitories. Approximately half of college students go to community colleges, which do not have dorms. Casual estimates found on the web indicate that of the students in four-year colleges, maybe 20% live in dorms.

The BigChampagne data that OS use to measure monthly piracy runs from August of 2002 to May of 2006. OS assume that what happens in the three complete summers during this 46-month period holds for all the summers during the full seven-year post-Napster period.

The logic of this quasi-experiment, therefore, holds *only if piracy regularly fell in each of the three summers for which they have data,* since the results for these three summers are to be generalized to the other four summers for which OS do not have data. Because a regular decline in summer piracy is the key requirement for this quasi-experiment, OS's examination of the declines in summer piracy is a worthy target for replication.

To begin, we can simply chart the monthly number of pirates using the BigChampagne data, as in Figure 1. The three complete summer periods in these data are denoted by circles around them.[4] Visual inspection of Figure 1 would seem to indicate that only the first summer clearly reflects lower piracy than the non-summer months. Indeed, if one calculates piracy levels in the summer months relative to the other months in each of those three years, as was done in my prior replication (Liebowitz, 2007), it is only the first summer that has a notable decline. The average summer piracy decline is 12%, matching the value reported by OS.[5]

Nevertheless, OS have explained, in response to the prior replication, that they were comparing the summer decline *relative to a trend.* They stated "[a]t 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."[6]

Figure 1 does not seem to comport any better with this latter explanation. It does not appear to reveal a "clear break" of an otherwise upward trend during the summer of 2005, since piracy in that summer is higher than in the previous winter/spring, and is basically the same as during the following winter. But let's bring this claim more formally to the data.

In their response to my earlier replication, OS claimed that they regressed the number of pirates:

> 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.

The singular term "indicator variable" suggests that they used a single dummy variable to represent all three complete summers. This seems strange, since using a single dummy would *not* allow an answer to the question of whether a summer piracy decline is a regular, routine occurrence *each* summer. Answering the latter question would require separate dummies for each summer.

---

4 American university summers tend to take place during June through August, although OS also include May and September. I use their 5-month definition of summer vacation in all the calculations below.

5 The first summer has a piracy drop of 38.8%, the second summer a drop of 0.2% and the third summer an increase of 3.8%.

6 This quote comes from a response by OS to an editor from an earlier submission of this replication.

Figure 1: Number of simultaneous pirates (BigChampagne)

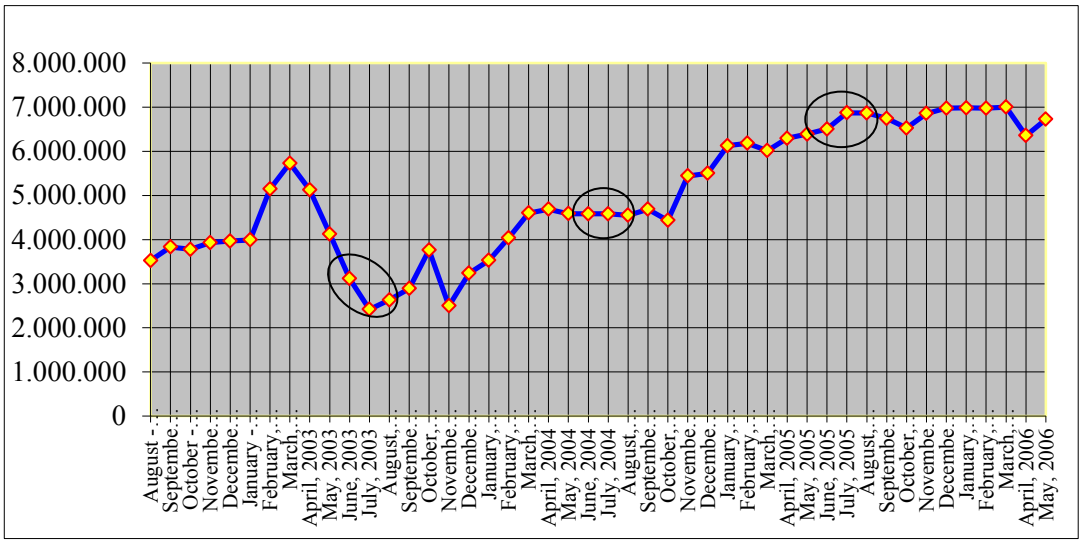

Nevertheless, it is easy to replicate the regression with the single summer dummy which is found in the first two data columns of Table 1. With the time trend included, the summer dummy is negative although of borderline statistical significance (8%), indicating that on average piracy falls in the summer, but the average decline appears to be 8.8%.[7]

If we are going to test the claim that piracy declines in *each* summer due, say, to college vacations, it is necessary that we examine each summer individually. The two rightmost columns of Table 1 provide the coefficients from including separate yearly dummies (all statistically significant) for each of the three summers, with and without a time trend.

The results of these regressions confirm the intuition from Figure 1. The key finding is that the summer of 2005 has a significant *increase* in piracy. This runs counter to OS's claim that piracy routinely *falls* in each summer, thus destroying the logic of this quasi-experiment.

But note, also, that the 2003 decline might be due, at least in part, to special circumstances. In June of 2003 the organization representing American record companies, the RIAA, officially announced that it was gathering evidence for future lawsuits against individual pirates, and in September it began to bring suit against individuals. Bhattacharjee et al. (2006) demonstrate that this threat of lawsuits reduced piracy. Therefore, we cannot even be sure that the 2003 summer piracy decline was not partially or completely an artifact of *sui generis* piracy politics.

In a recent defense of this quasi-experiment (Oberholzer-Gee and Strumpf, 2016), OS ignore my main discussion about the summer regressions, and instead focus on my secondary conjecture that the 2003 decline in summer piracy might be due to the lawsuits. They argue,

[7] There are various defensible methods that can be used to calculate the percentage summer piracy decline, so it is possible that some other method might match OS's stated 12% value, although none of three or four other methods I used provided a match. My calculation for the single summer coefficient in the table divides it by the average number of downloads during non-summer months for the three years with summers, although the full 46-month period was used for the regression. For the specification with three summer dummies, I divide each coefficient by the average piracy values for the 7 non-summer months in the same year, and average those three values.

Table 1: Regressions explaining piracy levels; robust t-stat in brackets

	Merged Summers, Time Trend	Merged Summers	Separate Summers, Time Trend	Separate Summers
Summer03			–1.29e+06	–2.12e+06
			[–3.42]	[–5.67]
Summer04			–611,668	–556,702
			[–4.14]	[–2.22]
Summer05			578,467	1.52e+06
			[4.37]	[5.72]
Summer3yr	–449,687	–384,482		
	[–1.82]	[–0.81]		
Time Trend	87,885		74,084	
	[10.8]		[9.91]	
Constant	3.11e+06	5.16e+06	3.43e+06	5.16e+06
	[10.4]	[21.1]	[12.1]	[20.6]
Observations	46	46	46	46
R-squared	0.70	0.02	0.78	0.38
Average # pirates in 21 non-summer months over 36 months	5,082,597	5,082,597		
Average Summer Impact	–8.8%	–7.6%		
Summer impact relative to non-summer, 2003			–30.6%	–50.3%
Summer impact relative to non-summer, 2004			–13.3%	–12.1%
Summer impact relative to non-summer, 2005			9.0%	23.6%
Average of 3 summers			–11.6%	–12.9%

erroneously, that the cause of the 2003 decline was irrelevant to their quasi-experiment. See Liebowitz (2017) for more details.

Finally, OS have argued, in their response to an earlier version of this replication, that another data set (Internet2) of supposedly high quality also found that American piracy fell in the summer, thus providing support for the basis of the quasi-experiment.[8] The problem with this claim is that Internet2 traffic is unsuitable to be a measure of how piracy fluctuates during the year, because Internet2 is a high-speed network which primarily connects American universities and research centers. Because American university activities slow during the summer, with most students and many faculty members on vacation, any decline in summer piracy on that network would be unrelated to overall changes in piracy.

[8] OS also claim that Internet2 monthly download data have a positive correlation of 0.49 with their download data. In Liebowitz (2017), I find that this correlation is –0.68 (I also provide the raw data so that the calculations can be verified by others). If Internet2 data were negatively related to OS's download data, and if Internet2 really was a sound data set for these purposes, that the data underlying OS's main regressions would have to be considered unreliable.

2 East Coast sales versus West Coast sales

OS argue that due to the time zone differences, East Coast Americans have greater access to European pirate-files than do West Coast Americans. The supposed reason for this is that prime-time American East Coast pirating occurs between 1am and 5am in Europe, whereas prime-time West Coast pirating takes place between 4am and 8am in Europe. OS take this to mean that Europeans are more likely to be sleeping, and thus have their computers off, during prime piracy times on the American West Coast compared to the East Coast. OS argue that if piracy harmed record sales, the East Coast should suffer a larger sales decline than the West Coast because East Coast users have more access to European files.

In their words:

> A second experiment considers spatial variation…Because of time zone differences, such [piracy] transfers are easier for East rather than West Coast users. [2007, page 36]

OS present some statistics that they take as supporting their conclusion that piracy has not harmed East Coast sales relative to West Coast sales. Here is their empirical summation:

> In 1998, the last year in the pre-P2P [pre-piracy] period, the share of album sales in the eastern time zone was 43.9 percent. This share has hardly moved since then. In 1999–2002, the mean was 43.5 percent and the range was 42.7–44.0 percent. This is consistent with some common national factors, rather than file sharing, driving sales trends.

OS note that the share of East Coast sales has remained in a narrow range and conclude that this small variation somehow demonstrates that sales in the East did not fall relative to sales in the West. OS do not compare actual market share changes in these two U.S. time zones. Surely, to test their hypothesis, it would seem to require an examination of the market shares of the East Coast relative to the West Coast after piracy became popular in 2000, since that is what the experiment claims to test.

Using the same Nielsen SoundScan 'album-sales-by-city' data for the same years as OS, plus an additional year's data, 2003, allows a comparison of yearly East Coast and West Coast market shares (with the shares of the middle time zones being calculated but not shown in the tables). After examining the data, I find, in common with OS, only small changes in East Coast market shares in the pre-Napster and post-Napster periods.

But for the purposes of testing their hypothesis, I examined the relative shares for the East and West Coasts before and after piracy's birth, which is shown in Table 2.

When the data are arranged to provide an answer to the question this quasi-experiment posed, the results (in Table 2) are contrary to OS's conclusion and are entirely consistent with hypothesis that piracy is harmful to sales.[9] The shares of the East relative to the West decline

[9] SoundScan data contain a large category ("others") of unallocated sales that could not be fit into their list of 100 DMAs (cities). OS put those sales into the middle time zones, which is an arbitrary procedure that obviously would incorrectly overstate the market shares for the middle time zones and understate the market shares for the East and West Coast. I removed those "others" sales from the analysis, which allows the calculation of less biased market shares for all time zones.

over these years, and they decline by a surprisingly large amount that is statistically significant when examined in a regression with a time trend.[10]

This replication, therefore, supports a conclusion that is the opposite of that reached by OS. I should note that I do not consider this to be evidence supporting a claim that piracy harms sales. I find the suggestion that a variation in the small number of Europeans awake late at night could measurably influence the piracy behavior of Americans to be sufficiently farfetched to attribute these results to some other unknown factor.

Table 2: Album Sales by Area and Year

	East	West	Middle	East Share	West Share	East/West
1998	328,917	105,287	198,820	51.96%	16.63%	3.12
1999	339,882	108,768	206,935	51.84%	16.59%	3.12
2000	354,671	115,373	210,732	52.10%	16.95%	3.07
2001	342,753	115,302	203,863	51.78%	17.42%	2.97
2002	299,574	106,158	182,074	50.96%	18.06%	2.82
2003	287,700	102,556	176,820	50.73%	18.09%	2.81

3 The relationship between monthly changes in file-sharing and record sales

The next quasi-experiment is a simple regression relating American record sales to the number of American pirates, with monthly fixed effects, over a 46-month period. In their words:

> When we use monthly data from August 2002 to May 2006 (N=46) and define Sales and Users in millions (with respective sample means of 56.0 million and 5.0 million), the estimated [coefficient] = –0.427 with a robust standard error of 0.33. There is little evidence that growth in the number of users has had a statistically or economically significant effect on sales. [page 37]

Table 3 contains the reported OS results (in the first column) along with what should be an identical regression on identical data (in the second column).[11] and a regression that also includes a simple measure of the health of the economy It is somewhat difficult to categorize this attempted replication. The coefficient in the replication (middle column) is 48% larger than OS's coefficient (left column), although they should have been identical.[12] It is disconcerting that the measured coefficients differ by such a sizable percentage. Nevertheless, the qualitative results are largely the same – a negative relationship between piracy and record sales that is not

[10] This result is statistically significant in the sense that regressing the ratio of East Coast to West Coast market shares on a time trend provides a statistically significant coefficient (t=6.8), providing apparent support for the piracy hypothesis, even though there are only 6 observations.

[11] The full regressions, including monthly dummies and the constant term, can be found in the Appendix

[12] Stata, SPSS, and Excel each generated identical values.

Table 3: Replicating Coefficients on Sales; Includes Monthly Fixed Effects

	OS Reported	Direct Replication	Add Unemployment
Number of pirates (Coefficient)	–0.43	–0.638	–2.54
t–stat (robust)	–1.29	–0.96	–2.16
Unemployment rate (0–100)			–6,326,399
t-stat (robust)			–1.87
Implied impact on sales, 2004	–23,586,017	–35,229,887	–140,300,898
Implied percentage drop from 2000 to 2004	–3.0%	–4.5%	–17.9%
Share of drop due to piracy, 2004	–26.4%	–39.4%	–156.9%

statistically significant. This replication fails to match the OS results, but it does not change the qualitative conclusion. This is one instance where a simple replication does not overturn their result.

A very straightforward broadening of the replication would be to include some measure of the health of the economy. To this end, a monthly unemployment rate was included in the regression, seen in the last column of Table 3.[13] The coefficient on piracy quadruples in size while achieving statistical significance (4%). The unemployment rate is of borderline significance (7%) with a negative sign, implying that record sales fall when the economy worsens, consistent with expectations.[14] The addition of this variable does overturn their result.

Even though the OS result is not statistically significant, OS make a separate claim that the coefficient on the number of pirates is economically unimportant. To assess this claim, the last three rows of Table 3 contains the predicted decline in record albums for 2004 (with similar results for other years as seen in the Appendix), the inferred percentage decline compared to the value in 2000, and the inferred share of the overall decline due to piracy since the peak of sales in 2000, as Napster was ramping up.[15] The middle column, for example, implies that piracy reduce 2004 record sales by about forty million units, a reduction equal to 4.5% of sales in the year 2000, and that this amounted to almost forty percent of the sales reduction that occurred between 2000 and 2004. Whether this is economically unimportant appears to be in the eye of the beholder, but I do not believe that declines of this size are economically unimportant. When the unemployment rate is added, the entire decline in album sales could be attributed to

[13] These values come from the U.S. Bureau of Labor Statistics. I used the non-seasonally adjusted values since the purpose is to measure the overall strength of the economy, and if it is always a little stronger during the Christmas season we would want to include that extra strength. Using the seasonally adjusted values would slightly increase the absolute value of the coefficients and t-statics for both the number of pirates and unemployment rate.

[14] OS have argued, in their response to an earlier version of this replication, that multicollinearity between the number of pirates and the unemployment rate makes these coefficients unreliable (the correlation is –.68). Their concern has some merit, but the VIFs are 5.91 and 4.78 for unemployment and piracy and respectively. These VIFs are moderately high, but in a range normally thought not to be indicative of a serious multicollinearity problem (typically thought to require a VIF above 10).

[15] The predicted decline in sales due to piracy is calculated as the product of the average number of pirates for the year 2004 and the piracy coefficient. In the Appendix I provide these measurements separately for each year. The overall sales decline is based on Neilsen SoundScan average record sales over this period, compared to 2000, which was the peak year of sales using SoundScan statistics.

piracy,[16] a result that happens to match what the majority of the economic studies have found about the impact of piracy on record sales (Liebowitz, 2016a).

In conclusion, a narrow replication produced a coefficient 48% larger than the coefficient reported by OS although it should have been identical, and although the results are qualitatively unchanged, I find the difference disconcerting. However, a slightly broader replication, including a measure of the economy's health, indicated a considerably larger economic impact of piracy that is inconsistent with their conclusion.

4 The relationship between genre piracy and genre sales

The final quasi-experiment examines whether genres of music that are most prone to piracy suffer larger sales declines than less piracy prone genres.[17] Here is OS's explanation:

> A third experiment, which also provides a test of the dropout hypothesis, is to see whether download [piracy] intensity influences long-run sales effect of file sharing growth after explicitly controlling for trends in [radio] music format popularity. [page 36]

For the period 1999–2005, OS found that genres with high piracy intensities were less likely to suffer sales declines than genres with low piracy intensities, although this result was very far from statistical significance [note that their measure of piracy intensity is from their 2002 data, which they assume remains representative for later years through 2005]. They provided few details of the underlying analysis other than to say that they ran regressions on a genre's sales changes controlling for genre piracy intensity and changes in the genre's popularity as measured on radio.

These missing details are important because, for the set of genres used throughout their paper, it appears to be impossible for OS to have correctly done what they claim to have done. Throughout their paper they include 11 "genres" into which albums were classified when in fact several of the "genres" are merely categories of sound recordings that are not musical genres.[18] Four of those categories (*catalog, new artists, current hits,* and *soundtrack*)[19] are, respectively, based on sales volume, age of recording, experience of the artist, or whether the music was in a movie.

Because these four genres are not based on musical characteristics, there are no radio station music genres that match these four record categories. For example, consider the record category

---

[16] A value over 100% means that there would have been an increase in sales from 2000 to the 2003–5 period, except for the negative consequences of piracy.

[17] See Liebowitz (2008) for a detailed examination of the impact of piracy on musical genres over a five-year period across 100 cities, where those genres with the greatest piracy intensity were also the genres that lost the most sales.

[18] OS's record genres include: Catalog(ue), Alternative, Hard, Jazz, Latin, New, R&B, Rap, Current, Country, Soundtrack.

[19] As an example of the heterogeneity with any of these four genres, consider the soundtrack from the movie "Frozen" and the soundtrack from the movie "Straight Outta Compton." Although both are successful soundtracks, one is a Disney movie for children and the other is a biography of the first gangsta rap group. There are no radio station genres which would play both albums.

"new artists." A radio station may play new rap artists, or new country and western artists, or new rock artists, or even new classical or jazz artists, but these stations would also play established artists and there are no radio stations that play, without respect to musical style, all types of new artists. Further, there are no radio genres that match the 'new artists' category (a list of radio genres is found in the appendix). Nor are there radio stations playing only soundtracks of every type of music. Ditto for "current hits" and albums older than 18 months (catalog) of any style of music. With no radio stations specializing in these genres, how could OS control for radio format popularity trends for these genres that make up more than a third of their genres? OS do not specifically mention whether they include these four genres in their genre regressions, although it appears that they did.[20] Nor do they tell us what radio genres they "matched" to these four record categories that do not have similar radio genres. It should be noted that even for actual musical genres, it is surprisingly difficult to match radio genres to sound recording genres.[21]

A separate issue is the OS measure of piracy intensity using their 2002 proprietary piracy data (the quality of which I have questioned elsewhere[22]). OS construct a variable intended to measure piracy intensity which they name "downloads per album" but this is apparently a ratio of the number of pirate downloads per album *title,* in a genre,[23] not downloads per album *sold*. I create an alternative measure of piracy intensity (which I refer to as "OS corrected"), which is downloads per albums sold. The Appendix contains a discussion of why the latter is the superior variable, and also addresses how all the piracy intensity variables used here differ from more ideal measures. The correlation between the OS variable and my "corrected" version of this variable, for all OS genres, is only 6.5%, although it increases to 62.3% when the four non-music record categories are removed.

Putting these difficulties aside, we can try to replicate their regressions using the seven actual musical genres in their data. I broaden the replication by also examining the period 2000–2005, which appears to be a more appropriate time frame because file-sharing (beginning with Napster) did not become prominent until 2000[24] and because SoundScan's measured sales

---

[20] It appears that the four genres are not all removed because OS state on page 37 that the mean value of pirate downloads per album *title*, across genres, is 61.2 (which, strangely, does not match the 57.1 value in their Table 3), which is very different than the average value of 35.9 if the four non-musical genres are removed, using the numbers in their Table 3.

[21] For example, one of the OS-chosen sales genres is "jazz" but although there is a radio "jazz" subcategory of New AC/Smooth Jazz, most listening measurements in that category are zero. On the other hand, "new adult contemporary", which is also a subcategory of "NewAC/Smooth Jazz" appears to be the closest to jazz, although it is not clear how close. Similarly, the OS "hard" genre (related to "metal") is not a radio category, with the closest appearing to be the "new rock" subcategory under "Alternative" or the "active rock" subcategory of "Rock." A complete list of American radio genres is in the Appendix.

[22] See Liebowitz (2016b).

[23] As seen in OS's Tables 1 and 3, their album sales averages are more than a thousand-fold larger than their measure of downloads, implying that the mean value of downloads per sold album would be much less 1. Therefore, the average value that OS provide (61.2 downloads per album) implies that the number of album *titles* per genre is in the denominator.

[24] I provide, in the Appendix, details on Napster's size that make clear that Napster was not economically important until the summer of 2000.

of sound recordings peaked in 2000, making this period more representative of the post Napster regime of declining record sales. One further broadening of the replication is to use an additional source of genre piracy intensity data, and I was able to acquire data on piracy intensity from the firm NPD, which is from the year 2005 as opposed to the 2002 data used by OS.[25,26]

Table 4 presents simple correlations between changes in genre sales and piracy intensity. In all cases, genres with higher piracy intensities experienced larger percentage declines in sales, and although the correlations are fairly large, because they are based on only seven observations the correlations would need have an absolute value greater than .75 to be statistically signify-cant.[27]

The actual replicated regressions linking piracy intensity and genre sales changes (controlling for radio genre audience changes) are found in Table 5. Higher piracy intensity is linked to lower sales in all six instances, although the coefficients are only statistically significant in one case.

The results of the regressions can be summarized in an intuitive fashion, as found in the bottom two rows of Table 5, which calculates the change in genre sales as piracy levels increase under two hypotheticals. In the first hypothetical (Economic Impact I), this calculation is performed under the assumption that the piracy intensity level increases from zero to the highest piracy intensity genre (which differs for the three measures of genre piracy intensity).[28] The first hypothetical gives an idea of the loss (in units) that might occur as all genres reach the same piracy-intensity as the most pirate-intense genre (measured as of 2002 for OS values or 2005 for NPD values). Five of the six values are unambiguously large and economically important, and the smallest is still a fairly large decline. Note that unlike Table 3, these values are not calculated relative to the actual loss but instead represent a straightforward measure of loss.

Table 4: Correlations between Sales & Piracy Intensity, (7 observations)

Sales change	NPD	OS corrected	OS original
99–05	–0.231	–0.355	–0.284
00–05	–0.488	–0.407	–0.305

[25] NPD data do not necessarily reflect a representative sample of internet users because its web users voluntarily allowed themselves to be monitored. There is no reason to think that this problem would bias the relative amount of piracy *across genres*, however.

[26] NPD provided data for nine genres but we could only match six to the OS genres and seven to the SoundScan sales genres. OS's data did not include "classical" and NPD's data did not include "Latin."

[27] OS have wondered why the NPD correlations here are weaker than in my 2007 working paper. The answer is that in 2007 I equated their "hard" genre with NPD's "rock" but I later decided that NPD's "metal" genre was probably a better match.

[28] The highest piracy-intensity value for the original OS measurement of downloads per album title is associated with "alternative", for the corrected OS measure of downloads per unit sales it is "hard", and for NPD it is "rap".

Table 5: Percentage Sales Change for Genre over 6 or 7-year period with Predicted Impact (constant hidden)

Timeframe	99–05	00–05	99–05	00–05	99–05	00–05
Genre piracy intensity	–0.078	–0.072	–0.12	–0.10	–0.033	–0.100
Robust t-stat	–0.73	–0.86	–0.85	–0.9	–1.09	–2.9
% Change in Genre Radio Audience	–0.089	–0.13	0.15	0.095	–0.15	–0.20
Robust t-stat	–0.45	–0.72	0.33	0.22	–2.7	–3.02
Observations	7	7	7	7	7	7
R-squared	0.093	0.127	0.148	0.176	0.459	0.60
Measure of piracy intensity	OS normalized	OS normalized	OS corrected	OS corrected	NPD	NPD
Economic impact I	–30.7%	–28.4%	–58.2%	–48.5%	–9.4%	–28.4%
Economic impact II	–10.8%	–9.9%	–17.2%	–14.3%	–4.0%	–12.1%

The second hypothetical impact (economic impact II) calculates the expected loss when a genre, starting from zero piracy, achieves the average piracy rate for all genres.[29] The results from the second hypothetical, found in the bottom row of Table 5, are considerably smaller than the first hypothetical since the piracy rate for the average genre is considerably less than for the genre with the maximum piracy rate. Although smaller, note that these average predicted declines are still quite large and are a fairly large portion of the actual decline in unit sales that took place from 2000–2005, which was 16.6%.[30]

The conclusion that the genres with the highest piracy rates suffered the largest sales losses is fairly consistent across specifications (although we cannot have a great deal of confidence in these results), revealing that it makes little difference which time period was used, whether the original OS piracy intensity numbers or the corrected OS piracy intensity numbers were used, or whether the NPD values were used. This is in contrast to OS, who found a positive (statistically insignificant) relationship.

I view these results as economically quite different than the OS genre results, even if the small number of observations makes it difficult to have normal levels of confidence in the results.

5 Data Replication

In this section we compare three of OS's data claims made in support of their thesis, against primary sources.[31]

[29] The average NPD value is only an approximation since we only have data for 67% of the album sales. Also, the difference between average genre and maximum genre differs between the NPD data and the OS data.

[30] The equivalent value for 1999–2005 is 13.3%.

[31] In my earlier 2007 working paper I discussed other claimed facts that I believe to be incorrect, but the demonstration in some of those cases was quite lengthy, so I merely refer the reader to Liebowitz (2007).

OS made two factual claims disputing whether record sales had experienced a serious decline (with the two claims noted by bracketed numbers):

> there are several trends that are inconsistent with the view that P2P [piracy] now displaces sales on a large scale…[1] music sales have been flat or even rising in major markets with a quickly growing file-sharing population. For example, [2] in 2005 retail music sales rose in four of the five largest national markets. [page 39]

Because OS provide no reference for these factual claims (and the JPE apparently did not ask), we cannot directly source-check their factual claims but instead we must go to our own direct sources. The primary, perhaps only, source for international comparisons of record sales by year is the IFPI (International Federation of the Phonographic Industry).

With regard to the first claim about there being major markets with flat or rising sales, we first need to define "major." I would think that major national markets are presumably larger than Switzerland, which has a population similar to that of New York City. I think most economists would define rising sales as an increase in real revenues.

Using IFPI data and looking at the top 10 markets (Switzerland is number 10), reveals that real retail sales rose in 0 of the top 10 markets during 2000–2005.[32] The IFPI data also reveal that real retail sales rose in only 1 of the top 5 (or 10) markets between 2004 and 2005, not the 4 markets claimed by OS.[33] The Appendix provides more details.

Later, OS claim:

> in the United States the entire drop in 2005 album sales is due to losses at a single firm [Sony-BMG]…If file sharing were responsible for the observed sales decline in the United States, we would not expect this activity to affect the products of only a single firm. (page 40)

Again, OS provide no reference to support their statement. Data from an online news source,[34] combined with SoundScan data on U.S. album sales in 2004 and 2005, allow a reconstruction of unit sales for the four major record companies and the large group of independent record labels which we can treat as a virtual fifth record company. These calculations (found in the Appendix) reveal that 2005 sales fell at four of the five record groups (all except UMG), and that Sony-BMG's decline was only 44% of the aggregated declines in the industry, not the 100% implied by OS.

These data replications clearly fail.

---

[32] Nominal retail sales rose in only 1 out of the top 10 markets, The same results hold whether the starting year is 1999 or 2000.

[33] Nominal retail sales rose in 2 of the top 5 (or 3 of the top 10) markets.

[34] The source can be found here: http://web.archive.org/web/20070117062246/http:/www.undercover.com.au/news/2006/jan06/20060105_universal.html

6 Conclusion

I do not view any of OS's quasi-experiments as holding up well, under replication. In the one case where the qualitative results are the same in both the original paper and the replication, the coefficients are nonetheless distressingly different from one another. And a simple broadening of that quasi-experiment led to a very different qualitative conclusion. In another quasi-experiment the replication results are of the opposite sign (of OS's) and very important economically, although statistically significant in only a small minority of instances (although the same result from multiple data sources enhances our confidence somewhat). Yet the very small number of observations (6) makes the achievement of statistical significance difficult, so that I believe this replication fails to support OS's results. The other two replications of quasi-experiments clearly provided very different results than those reported by OS. I find that the overall evidence from my replications of this "important complement" to their main regressions does not support the conclusions put forward by OS. These findings, together with those of Liebowitz (2016b, 2017), would suggest that other researchers not rely on the conclusions of OS (2007).

Acknowledgement: I would like to thank Steve Margolis, Peter Lewin and Alejandro Zentner for their advice. I acknowledge having received financial support from the Center for the Economic Analysis of Property Rights, at the University of Texas at Dallas. No other entity or person has paid for this research. I note that I often consult on media topics and have been paid to testify as an expert witness both for and against the sound recording industry.

References

Bhattacharjee, S., Gopal R. D., Lertwachara K., and. Marsden, J. R. (2006) Impact of legal threats on music sharing activity: An analysis of music industry legal actions. *Journal of Law and Economics*, 49(1): 91–114. http://www.journals.uchicago.edu/doi/abs/10.1086/501085

Glenn, D. (2008). Dispute over the economics of file sharing intensifies. *Chronicle of Higher Education*, July 17. http://www.chronicle.com/article/dispute-over-the-economics-of/989

Häring, N. (2008). The economists' war about downloading [in German]. *Handelsblatt*, March 3. http://www.handelsblatt.com/politik/konjunktur/oekonomie/nachrichten/journal-of-political-economy-in-der-kritik-der-download-krieg-der-oekonomen/2929688.html

Liebowitz, S. J. (2007). How reliable is the Oberholzer-Gee and Strumpf paper on file-sharing? (September 1). http://ssrn.com/abstract=1014399

Liebowitz, S. J. (2008). Testing file-sharing's impact on music album sales in cities. *Management Science,* 54(4): 852–859. http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1070.0833

Liebowitz, S. J. (2016a). How much of the decline in sound recording sales is due to file-sharing? *Journal of Cultural Economics*, 40(1): 13–28. https://link.springer.com/article/10.1007/s10824-014-9233-2

Liebowitz, S. J. (2016b). Why the Oberholzer Strumpf file-sharing article is not credible. *Econ Journal Watch*, 13(3): 373–396. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2887151

Liebowitz, S. J. (2017). Responding to Oberholzer-Gee and Strumpf's attempted defense of their piracy paper *Econ Journal Watch*, 14(2).

Oberholzer-Gee, F. and Strumpf K. (2007). The effect of file sharing on record sales: An empirical analysis. *Journal of Political Economy*, 115(1): 1–42. http://www.journals.uchicago.edu/doi/abs/10.1086/511995

Oberholzer-Gee, F., Strumpf K. (2016). The effect of file sharing on record sales, revisited, *Information Economics and Policy*. 37: 61–66. http://www.sciencedirect.com/science/article/pii/S0167624516301226

Appendix

The following materials are numbered to match the sections of the paper.

3 Monthly piracy and record sales

Here are the full regressions, including the monthly dummies, that lead to Table 3:

Table A: Dependent Variable is monthly unit sales of albums (including singles)

	Coef	T-stat	Coef	T-stat
# of filesharers	–0.6378	–0.96	–2.54	–2.16
Unemployment rate			–6326399	–1.87
December	6.16e+07	11.9	5.67e+07	9.2
November	1.23e+07	3.54	7.64e+06	1.76
October	1.18e+06	0.26	–4.58e+06	–0.76
September	–875,156	–0.18	–6.01e+06	–1.01
August	–1.15e+06	–0.24	–5.13e+06	–0.99
July	–1.20e+06	–0.34	–3.27e+06	–0.89
June	4.21e+06	0.74	2.98e+06	0.54
May	–2.24e+06	–0.68	–6.73e+06	–1.42
April	–519,282	–0.11	–4.39e+06	–0.87
March	4.09e+06	0.79	3.48e+06	0.78
February	2.81e+06	0.87	3.14e+06	0.99
Constant	5.23e+07	11	9.94e+07	3.89
Observations	46		46	
R-squared	0.911		0.917	

Since Table 3 only shows the size of the impact for 2004, here are the predicted piracy-induced reduction in sales for 2003 and 2005:

Table B: Implied Impact of Piracy on Sales for Additional Years

Additional implied impact of piracy on sales	OS Reported	Direct Replication	Add Unemployment
Implied impact on sales, 2003	–19,078,501	–28,497,114	–113,488,038
Implied Percentage Drop from 2000 to 2003	–2.4%	–3.6%	–14.5%
Share of drop due to piracy, 2003	–14.8%	–22.1%	–88.1%
implied impact on sales, 2005	–33,470,308	–49,993,824	–199,097,384
Implied Percentage Drop from 2000 to 2005	–4.3%	–6.4%	–25.4%
Share of drop due to piracy, 2005	–25.6%	–38.3%	–152.4%

4 Genre piracy and genre sales

Piracy intensity measure

The genre piracy intensity measure, for our purposes, should reflect of number of recorded pirated songs listened to as a share of all recorded songs being listened to by those consumers who would have listened to those songs if piracy did not exist. If piracy has no effect on sales, this ratio should be close to zero since pirated songs would only be listened to in order to decide whether to purchase the song or not, and that decision should not take too many repeated plays. Pirate listeners who would not have purchased the song without piracy should be excluded from these calculations since their behavior does not directly affect record sales.

For obvious reasons, we cannot construct an ideal measure of piracy intensity. OS have information from their sample of pirates on the number of pirated downloaded songs and the national sales for their sample of albums, classified by genre. NPD has data on the share of users in their sample who listen to pirated songs and the share who listen to purchased songs.

These variables allow the creation of rough measures of piracy intensity. OS use "number of downloads" [number of pirated files] and appear to divide it by the number of album *titles* in their sample of albums in a genre [footnote 23 discussed why this appears to be the case] although it is not clear that this was their intention. I create a variable which I believe better fits their variable name "downloads per album" by using album sales in the denominator, not album titles.

Downloads per album title is not a good choice for a pirate intensity measure because the number of album titles may not reflect the number of albums sold. Here is a simple example to illustrate the point:

> Genre A has 1 million albums sold, 1 million total pirated albums, and 1,000 titles (each selling 1000 units, on average). Genre B has 1 million albums sold, 1 million total pirated albums, and 2 titles (each selling half a million on average). Each genre has the same number of albums sold. Each genre has the same number of pirated albums. Each genre, therefore, has the same piracy intensity, about 50%. Yet, the OS measure of piracy intensity would indicate that Genre B has 500 times the piracy intensity as Genre A because the number of pirated files per album title is 500 times greater.

Finally, these measures take the form a/b, where, if we could do it correctly, a is a measure of piracy (on the part of those who would otherwise purchase the music) and b is a measure of sales. A more correct variable would take the form a/(a+b) [where a+b equals the size of the potential market], and this ratio would be linearly related to piracy intensity as it would be expected to influence sales, assuming that piracy reduces sales. The variable actually used, a/b, grows more rapidly as piracy grows than the variable a/(a+b), if piracy reduces sales (b). In other words, piracy increases not only increase the numerator of a/b, but also decrease the denominator. Thus, the measures of piracy being used will show too high a level of piracy for genres with high levels of piracy and the coefficients relating sales changes to piracy intensity will be *understated* relative to the more appropriate variable a/(a+b). This understatement of the size of the results discussed in the text merely lends more credence to the claim that this evidence supports the hypothesis that piracy decreases sales.

The history of Napster

This figure, taken from my 2006 article, demonstrates the growth of Napster. Napster was not even tracked by ComScore until February of 2000, when it had 1.3 million users. It peaked at just about ten times that level in February of 2001 (after which pirates moved to other pirate services). Obviously, the number of users when Napster started in 1999 would be much less than in 2000. Time Magazine put Napster on its cover in October of 2000. The RIAA brough legal action against Napster in *December* of 1999 in order to nip its piracy problem in the bud.

Figure: American Napster Users at Home (000s)

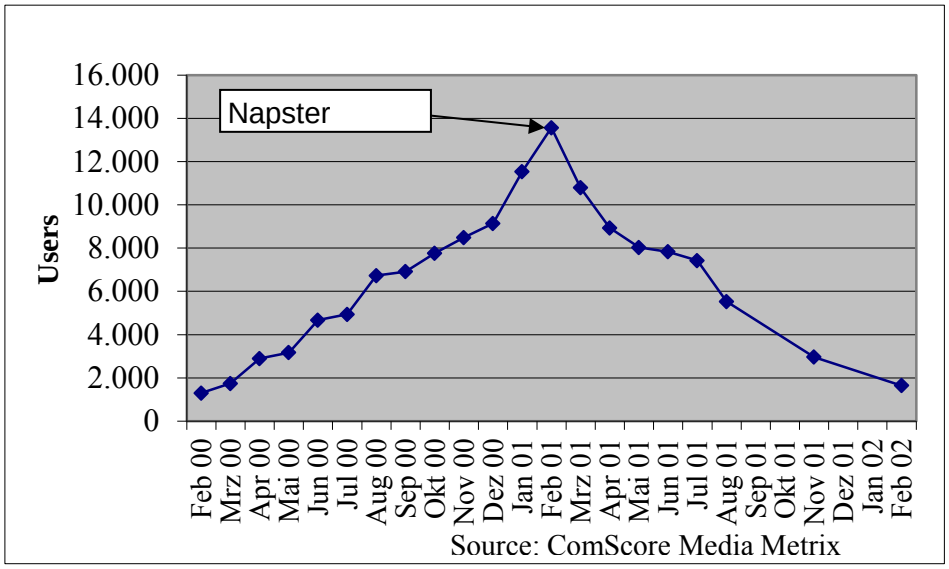

Radio Station genres

Here is a list of radio station genres and sub genres according to Arbitron, the leading organization measuring the size of radio audiences. It should make clear how difficult it can be trying to match a record category, of which there are not as many, to the correct radio genre, since the two have different terminologies:

Adult Contemporary
- A/C
- HOT A/C
- MOD A/C
- SOFTA/C

Adult Hits
- ADLT HT

Alternative
- AAA
- ALTERN
- NW ROCK

Contemporary Hits Radio
- CHR
- POP CHR
- RHY CHR

Classical
- CLASSCL

Country
- CLCNTRY
- COUNTRY
- NWCNTRY

Adult Standards
- EZ LIST
- MOR
- NOSTLGA
- VARIETY

New AC/Smooth Jazz
- JAZZ
- NAC

Oldies
- 70 HITS
- 80 HITS
- CL HITS
- OLDIES
- RHY OLD

Rock
- 70SROCK
- ACTROCK
- AOR
- CL ROCK

Spanish
- LAT URB
- MEX REG
- SP ADUL
- SP CONT
- SP LANG
- SP NEWS
- SP OLD
- SP REG
- SP RELG
- SP TROP
- SP VRTY
- TEJANO

Urban
- URB A/C
- URB OLD
- URBAN

News/Talk/Information
- ALLNEWS
- ALLSPRT
- NWSTALK
- TALK

Remaining Formats
- CHLDRAD
- EDUCAT
- ETHNIC
- OTHER
- R&B

Religious
- CHRSTN
- CO INSP
- GOSPEL
- RELGOUS
- SGOSPEL

5 Data claims made by OS

OS claim to have found leading markets with increases in sales over the period 2000–2005. In Table C are listed the nominal and real sales changes for the ten leading markets over the period 2000–2005.

Table C: 2000–2005 Change in Revenue

	Real	Nominal
USA	–29.9%	–12.6%
Japan	–10.7%	–13.6%
UK	–14.4%	1.6%
Germany	–42.8%	–32.5%
France	–26.2%	–14.5%
Canada	–36.7%	–21.7%
Australia	–29.7%	–8.3%
Italy	–41.0%	–25.8%
Spain	–56.2%	–34.6%
Netherlands	–46.1%	–30.2%
Switzerland	–19.5%	–19.5%

None of the top 10 markets have an increase in real retail revenue over this period. These data come from the IFPI annual publication "Recording Industry in Numbers."
OS also report that music sales rose in 4 of top 5 markets in 2005. The table below, using the same IFPI data on retail sales, reveals that 1 of the top 5 markets had a real increase in revenue and 2 of the top 5 had a nominal increase in revenue. An additional five countries (ranked by size) are shown to indicate none of the other top 10 markets had real sales increases that year although Italy had a nominal increase.

Table D: 2004–2005 Change in Revenue

	Real Revenue	Nominal Revenue
USA	–5.30%	–1.90%
Japan	1.94%	1.64%
UK	–5.46%	–2.66%
Germany	–1.56%	0.44%
France	–4.24%	–2.54%
Canada	–4.32%	–2.12%
Australia	–12.33%	–9.63%
Italy	–1.73%	0.07%
Spain	–8.47%	–5.07%
Netherlands	–18.27%	–16.57%

Finally, the claim that market shares only fell at a single record company is refuted by the following data. Market shares are from the news story in footnote 34 which references SoundScan data for the market shares of leading record companies. Sales figures for each

company are derived from Nielsen's SoundScan's measure of U.S. sales. I multiplied the industry yearly sales by the market share statistics for firms to derive each firms' sales. The last column calculates the change in sales, which is negative for all firms except for Universal (UMG).

2005 Sales Declines Supposedly at only a Single U.S. Record Company

	Market Shares		Sales (000s)		
	2004	2005	2004	2005	Sales Change
UMG	29.59%	31.71%	201,475	207,449	5,974
SONYBMG	28.46%	25.61%	193,781	167,542	–26,239
WMG	14.68%	15.00%	99,954	98,131	–1,824
EMI	9.91%	9.55%	67,476	62,477	–4,999
OTHERS	17.36%	13.18%	118,202	86,224	–31,978
Industry Sales	680,889	654,206			

Conflict of Interest Disclosure Form

Please state the sources of financial support for the research.

Center for the Analysis of Property Rights and Innovation (at UTD)

Please identify each interested party from whom you (or a close relative or partner) have received financial support summing' to at least U.S.$10,000 in the past three years:

Law and Economics Consulting Associates Ltd, Australia

British Academy of Songwriters, Composers and Authors; Musicians' Union; and UK Music 2009 Limited (joint plaintiffs)

Please disclose any paid or unpaid positions as officer, director, board member of relevant nonprofit organizations of profit-making entities held by you (or by a close relative or partner):

Director, Center for the Analysis of Property Rights and Innovation

Has another party had the right to review the paper prior to its circulation?

No

Signature:

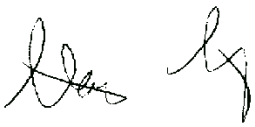

Stan Liebowitz

May 11, 2017