EJW
ECON JOURNAL WATCH
Scholarly Comments on
Academic Economics

ECON JOURNAL WATCH 14(2)
May 2017: 174–195

Responding to Oberholzer-Gee and Strumpf's Attempted Defense of Their Piracy Paper

Stan J. Liebowitz[1]

LINK TO ABSTRACT

In September of 2016, I published an article (Liebowitz 2016b) in *Econ Journal Watch* (henceforth referred to as EJW) that critically examined a well-known 2007 article by Felix Oberholzer-Gee and Koleman Strumpf (henceforth OS), who had found piracy to have essentially no impact on record sales (OS 2007). Although OS were invited by the editors of EJW to publish a response contemporaneously with my article, they instead submitted their response to *Information Economics and Policy* (IEP) several weeks after EJW published my article. That response by OS (2016) also included their reactions to several other criticisms I had made about a set of four "quasi-experiments" found in their 2007 article. My criticism of those quasi-experiments is found in Liebowitz (2017).

This paper responds to the claims made by OS in IEP.[2] Comparing their IEP article to my original EJW article reveals that their IEP article often did not respond to my actual criticisms but instead responded, in a cursorily plausible manner, to straw men of their own creation. Further, they made numerous factual assertions that are clearly refuted by the data, when tested.

1. University of Texas at Dallas, Richardson, TX 75080. I would like to thank Steve Margolis and Alejandro Zentner for comments, and the Center for the Analysis of Property Rights and Innovation for financial support.

2. I attempted to respond to the OS comment in IEP, but my submission was rejected. The editor of IEP stated "that continued debate on this one paper [OS 2007] is not helpful to researchers actively working in this area, or to IEP readers." I should note that Oberholzer-Gee has been, for that editor, a reference, coauthor, and dissertation committee member.

I must also note that OS's reply is incomplete. For example, I noted in my EJW article that OS reported the average number of German students on vacation—their key instrument—to be 9.855 million (out of a reported 12.491 million students), although these values implied that the average *share* of students on vacation was a remarkably high 79 percent (Liebowitz 2016b, 381). After examining the German school calendar for that period, I showed that this average share of students on vacation was too high by a factor of about four. Either the 9.855 and 12.491 values were misreported, or there are important errors in their main data set, which remains publicly unavailable. I accordingly referred to this large error affecting a crucial variable as a "canary in the OS coal mine." I made plain that I considered the variable to be "mismeasured" and that OS's major results thus might be entirely unreliable:

> It appears, therefore, that there is something very amiss with the OS NGSV [students on vacation] numbers that were used in their regressions. With the NGSV values used in the regressions being too large, on average by a factor of four, there is no telling what the impact might be on the regression estimates since we have no idea how these measurements correlate with the correct values. (Liebowitz 2016b, 382)

But in their IEP article, OS do not acknowledge or address this important issue at all.

Another 'canary in the coal mine' that I discovered working on this rebuttal has to do with a heretofore unreported but important error in OS (2007). OS mismatched their weekly download data and school vacation data, with the two sources being off by a week. The error corrupts their reported regression results based on their key instrument. This new material is discussed in detail below.

OS (2016) also fail to respond to some other matters; for example, they do not attempt to defend two of their four quasi-experiments that I have argued to be seriously flawed. But, as we will see, OS did mount responses to a number of points I made in EJW that, by comparison to the charge that their key instrument was seriously mismeasured, were often far less important. In a similar vein, they also 'responded' to several arguments that they claimed or implied I had made in my EJW article, when in fact I had not made those arguments.

To make it easier to follow the material below, the numbers at the end of my section headings indicate the section numbering found in OS's IEP article.

The excessive variability of OS's download data (2.1)

The key novelty of OS's 2007 article was its authors' access to download information from two pirate servers during the final 17 weeks of 2002. Their main results depended entirely on the reliability of this download data.

In EJW I examined the variability of the weekly OS pirate download data, concluding that it was both remarkably higher than economic intuition would suggest and also much higher than the variability of other piracy data. Due to its abnormal week-to-week variability, I concluded that the OS download data appeared to be of dubious reliability. In their IEP response, OS misrepresent my analysis and provide an error-ridden rebuttal.

Note that OS's download data initially contained 260,889 American pirate downloads (covering 17 weeks) that were reduced to 47,709 after OS matched these downloads to the sample of music albums used in their regression analysis (see OS 2007, 9). When I examined their download data in EJW, I expressly used the initial, unreduced sample of downloads, precisely to avoid any possible concerns that the album-matching process might have played a role in causing the unusually high weekly variability. Here is what I wrote in EJW, in the main text:

> The OS download data in the charts above [showing extreme variability] are taken from OS's tables listing the *full* American download numbers. In other words, these are the American download numbers *prior* to OS's attempt to match American downloads to American sound recordings. (Liebowitz 2016b, 381, emphasis in original)

Nevertheless, in their IEP article OS overlook what I wrote and blame the seemingly unreasonable variability in downloads on the very album-matching process that I expressly circumvented:

> The difference [higher variability in OS's download data than other data sets] is not difficult to understand. Our analysis omits album-weeks during which an album has not yet been released… The omission is empirically important… [T]he effective [matched] 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. (OS 2016, 62)

Obviously, their 'explanation' has nothing to do with my demonstration of the unusually high variability in their data, since I used the "total weekly number of downloads," not the matched downloads.

Figure 1a. All U.S. downloads in OS dataset

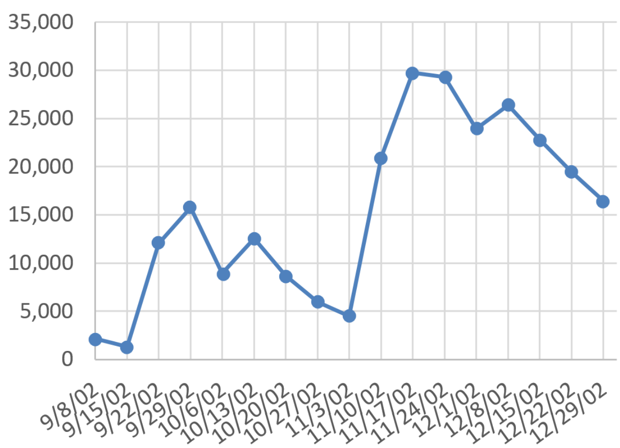

Figure 1b. Matched U.S. downloads in OS dataset

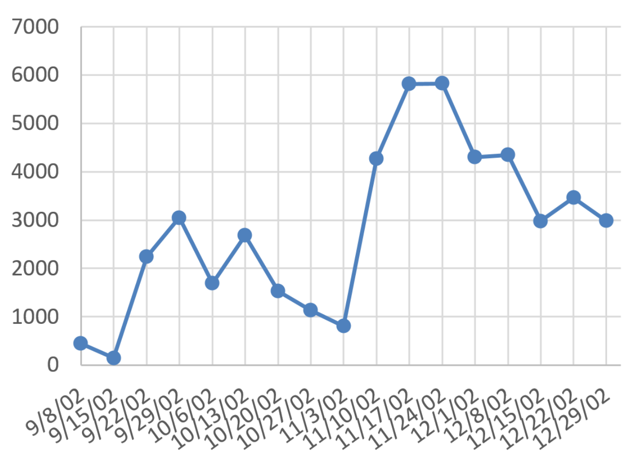

But even had I used the matched data, their 'explanation' would still be wrong, because their claims about differences in the matched and total download data—claims for which they offer no evidence—are wrong. For example, their assertion that album releases in their matched data cause "empirically important" increased variability is false. We can examine differences in the data by comparing

Figures 1a and 1b, which show the total downloads and matched downloads, respectively (note the different values on the vertical axes). If the matched-album data set were to have significantly increased weekly variability, as OS claim, then the weekly variability should be clearly greater in Figure 1b than in Figure 1a. But clearly, the same excessive weekly gyrations affect both samples, and this is confirmed by the correlation of 0.974 between the two measures. Any difference in variability is trivial. The additional OS assertion, that the difference between matched and unmatched data would be greatest in the early weeks, is evidently wrong as well, since the main noticeable differences are in the later weeks and the correlation is higher in the first half of the period (0.99) than the second half (0.87).[3]

After incorrectly suggesting that the excess variability in their pirate download data was due to their matching downloads to albums, OS change direction and argue that the temporal variability of their weekly download data is reasonable because their data show weekly trends similar to a different and purportedly reliable data set. After declaring that the leading data sets measuring downloads, BigChampagne and NPD, were of "poor quality," OS assert that a data set from Internet2 contained a "better" measure of pirate downloads. OS also assert that Internet2 weekly data during 2002 had a positive correlation of 0.49 with their weekly data.

These assertions, too, are erroneous. Most importantly, I find that the Internet2 data are *not* positively correlated with the OS data. Instead, the correlation of Internet2 values with the OS album-matched download data is −0.68 for the full download data set, which is the more appropriate comparison, and −0.62 for the matched data. I realize that my calculated correlations are remarkably different than the +0.49 correlation reported by OS, but all I can do is make my raw numbers and calculations available to everyone,[4] allowing any reader to check the raw data and results for themselves, and hope that OS will provide their calculations and data so as to help ascertain why their results are different from mine.

If OS were correct that Internet2 provided the best download data, then the strong negative correlation of this data with the OS data would actually further validate my claim that OS's weekly download data appear to be unreliable. Al-

3. The OS download data (both pre- and post-matching) can be found in Table 3 on page 38 of the March 2004 working paper version of OS's "The Effect of File Sharing on Record Sales: An Empirical Analysis" (**link**).

4. Here are the 34 numbers, for each of 17 weeks. OS download data (taken from the source in footnote 3) for weeks of September 8 to December 29, 2002: 2164, 1347, 12051, 15742, 8922, 12534, 8688, 5967, 4468, 20936, 29755, 29284, 23914, 26404, 22820, 19428, 16465; Internet2 data can be found in Table 6 (**link**) in each weekly report, measuring packets in Gigabytes for the "File Sharing" category, weeks of September 9 to December 30, 2002: 168.3, 180.6, 175.6, 84.41, 141.6, 116.3, 82.28, 96.23, 84.39, 76.0, 70.56, 68.37, 63.43, 52.35, 24.73, 6.74, 9.714. My full analysis, including the raw data from which the above numbers are taken, and the Excel sheet calculating the correlations, can be found in the online appendix (**link**).

though I would like to present that negative correlation between Internet2 and OS data as a powerful piece of evidence that the OS data are unreliable, I cannot do so—because, contrary to OS's claim, Internet2 is anything but a superior data source, for our purposes. Internet2 is a specialized research and education network. OS know this since they elsewhere defined Internet2 as "the U.S. high-speed network which primarily connects universities" (OS 2010, 28). Internet2 data, therefore, are narrowly based and not representative of overall Internet piracy, particularly at times when students go home, such as during the Christmas or summer breaks.

I can only conclude that the claims made by OS in their IEP article actually enhance my original arguments that OS's download data, the key novelty of their analysis, is highly suspect.

A newly discovered error in OS (2007)

OS state, with regard to their download (piracy) data: "Our data were collected from two OpenNap servers, which operated continuously for 17 weeks from September 8 to December 31, 2002" (OS 2007, 6). September 8, 2002, was a Sunday. December 31 was a Tuesday. This is not 17 weeks of data—it is 16 weeks and three days. How OS handled the partial week is unclear.

Nevertheless, it is reasonable to expect, and in fact the logic of their analysis requires, that their data on German schoolkids would be based on the same 17 weeks as the download data. But it appears that the school data are mismatched by a week from their download data.

In Figure 2, I have reproduced from an earlier version of their paper (**link**) a chart representing the share of German students on vacation each week that is more legible than the one found in Figure 1 in OS (2007). The aggregated German students are represented by the blue line with circle markers, although the circles are covered in later weeks by yellow triangles and green squares representing students in the German states of Bavaria and Rheinland-Pfalz respectively. Note that during the first two weeks, all the students in Bavaria were reported to be on vacation whereas for the first four weeks, none of the students from Rheinland-Pfalz were reported to be on vacation.

In the Appendix I reproduce a web page providing the school holidays for every state in Germany during the 2002/2003 school year. It shows that the week of September 8 is a vacation week for some German states. In particular, Bavaria (Bayern) had its summer vacation from August 1 through September 16, a Monday. Thus, during the week of September 16, one of five days was a vacation. We can see this this 20 percent value for Bavaria in the third week on OS's chart. Similarly, the school calendar for Rheinland-Pfalz indicates their vacations took place from Monday September 30 to Friday October 11. OS's chart shows Rheinland-Pfalz having their fall vacation in the fifth and sixth week.

Figure 2. Colored version of OS's Figure 1 (2007)

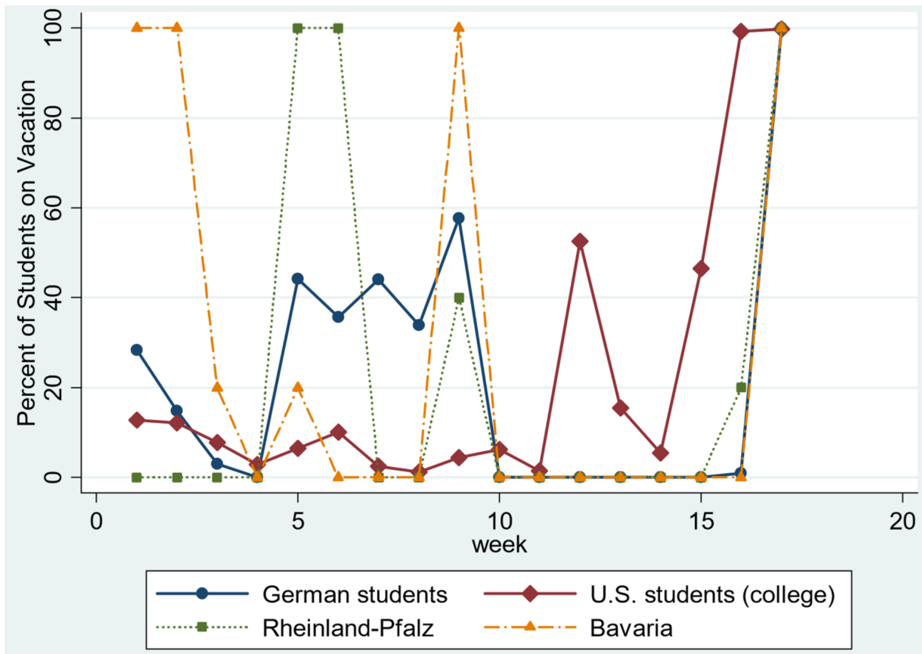

Source: OS, "The Effect of File Sharing on Record Sales: An Empirical Analysis," December 2006, page 55 (**link**).

But, if the first week of download data begins on Sunday September 8, as OS state, then the school week of September 16 (where Bavaria had 20 percent vacations) was the *second week*, not the third week reported by OS. Similarly, when Rheinland-Pfalz went on vacation on Monday September 30, that is the fourth week of download data, but OS count it as the fifth week.

We can see this at the end of the data range as well. The 16th week begins on Sunday December 22, and the 17th week of downloads begins on Sunday December 29. The 2002 school-year holiday calendar shows that the Christmas holiday (*Weihnachten*) begins on December 23 (a Monday) in every German state. So, the last two weeks have all students on vacation—yet the OS data incorrectly show virtually no German students (just the students in Rheinland-Pfalz) on vacation during the 16th week.

This means that the OS analysis has mismatched the download and vacation data. The download data begin on September 8, but the vacation data actually begins the week of September 1.[5] Each of these two variables is mismatched, relative to each other, by one week during the entire seventeen weeks of observations. OS are comparing this week's U.S. piracy with last week's German school vacations, although the logic of the relationship—i.e., American piracy increases because German students on vacation increase the files available to American pirates—is clearly one that requires direct concurrency. By itself, this is sufficient to call into question all the results of OS's analysis based on these two variables, which is almost their entire set of regression results.

Could American pirates be affected by German school holidays?

Using their download data, OS (2007) claimed that American pirates downloaded a surprisingly large share of their files, 16.5 percent, from German pirates. Because of the importance of German files to American pirates, OS argued that their key instrument, the number of German students on school vacations, would also have an important impact on the availability of pirate files to American pirates, assuming (without evidence) that German students shut off their computers while at school but keep them on during school vacations.

In EJW I analyzed this claim by making five adjustments required to estimate what share of files downloaded by Americans were likely to be influenced by German school holidays. I performed this analysis in a conservative fashion, making numerous assumptions favorable to the OS hypothesis. My conclusion was that German school holidays were unlikely to influence more than 0.14 percent of files available to American pirates. This number is so low that any impact of German school holidays would likely be swamped by statistical noise. Further, even if this number understates the true impact by a considerable margin, the true impact would still be trivial in size.

OS make several claims in IEP to dispute my conclusions about the small impact of German school holidays on American pirates. I address those claims below.

5. Figure 3 below gives a corrected version of OS's Figure 1, listed as Figure 2 here. It is also worth noting that some of the values found in OS's chart are wrong. For example, they indicate that the share of students reaches almost 60 percent during the ninth week. In fact, the share of students on vacation never reaches above 48 percent, except for the Christmas weeks.

Germans overrepresented in the OS database (2.4)

In EJW, I concluded that OS's download data overstated the importance of German files to American pirates. American pirates were only 2.3 times as numerous as German pirates in the OS data, but as I documented, more authoritative sources indicated that 2.3 was far too low a number. Consider:

1. The U.S. population under age 40 was 4 times that of Germany. (OECD)
2. The U.S. had 5 times as many pirates as Germany. (BigChampagne)
3. The U.S. Internet penetration rate was 1.1 times that of Germany. (OECD)
4. The U.S. broadband penetration rate was 1.7 times that of Germany. (OECD)
5. The musical taste of American pirates and the time of day they downloaded files was more closely matched to that of other Americans than to Germans (who were in far-off time zones listening to German and American music), implying that American pirates would download American files at a rate beyond what the first four factors alone would imply.

These facts make a compelling case that American pirates would be considerably more than 2.3 times as plentiful as German pirates and thus that the OS data overcounted Germans relative to Americans. I concluded that it would be conservative to raise the number from 2.3 to 4.6, thereby halving the share of Germans and the importance of German files to Americans. This seemed conservative since 4.6 is less than the simple comparison of the number of pirates (point 2 above), and ignores the time-zone, language, and broadband advantages of American files to American pirates.

In IEP, OS provide two defenses of their data and its 2.3 ratio of American to German pirates. First, they criticize BigChampagne's measure of the number of pirates (point 2 above). They argue that BigChampagne's methodology was flawed because it measured the "number of users who share files, not the number of users who download content" (OS 2016, 63). OS apparently fail to understand that BigChampagne is measuring exactly what should be measured because American pirates get their files from other pirates *making files available*, not from pirates who do not make their files available. Additionally, OS claim that BigChampagne failed to properly measure the number of pirates in Japan (ibid.)—but it is only the comparison of Germany relative to the U.S. that matters, making this claim irrelevant.

Second, OS (2016, 63) assert that the importance of German files to Americans is supported by "better" contemporaneous evidence from Expand Networks. I should note that the Expand Networks paper cited by OS does not say anything about the importance of German files to American pirates. The referenced paper, however, does say that Expand Networks collected their data by monitoring "the line that connects some 10,000 local [Israeli] ADSL and cable users to the USA" (N. Leibowitz et al. 2002, 2). It seems fanciful to believe that this setup could have provided an accurate measurement of the importance of German files to American pirates. Thus, even if there were some data from Expand networks showing a large share of German users, a claim unsubstantiated by OS, it seems incorrect for OS to claim that Expand Networks had better data than BigChampagne, which, unlike Expand Networks, had piracy monitoring as its main source of revenue.

How many children are affected by typical school holidays? (2.5)

At the end of their section 2.5, OS dispute my claim that the typical German school holiday affects only a small share of German students at a time. My EJW paper calculated (in a manner generous to OS) that 25% of German students were affected by a typical school vacation. This would have been appropriate if vacations were evenly spread out with non-vacation weeks, but the vacation weeks and non-vacation weeks tend to be bunched together, as seen in Figure 3.[6] OS (2016, 64) correctly point out that it would be better to have used the average weekly *variations* in the number of students on vacation.

But they also implied, erroneously, that my usage of a 25 percent average share of students on vacation understated a value calculated using weekly variability:

> [T]he 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. … [T]he claim that…there is little variation in this type of supply is incorrect. (OS 2016, 64)

For some reason, OS never calculate the average weekly variation that they propose, although it is easy to do and could presumably have substantiated their assertion that the weekly vacation variability was greater than the value I had used.

6. This figure differs from the share of German kids on vacation found in Figure 4 of my earlier EJW article because the figure in my prior article matched OS's analysis by beginning on September 2 and ending on December 23, whereas this figure begins and ends a week later so as to match the weeks used in OS's download data (2007) as discussed above.

Performing that calculation, the average week-to-week fluctuation is found to be 13 percent—48 percent *lower* than the 25 percent value that I used in EJW, even though it includes the "two consecutive weeks where no kids and then all kids are on vacation."[7] This method, therefore, leads to an impact of German vacations on American file availability that is considerably *smaller* than the one I used in EJW. If this measure of the impact of school vacations is used in place of my original estimate, the estimated overall impact of German school holidays on files available to American pirates would be 0.07 percent, not the 0.14 percent I reported, making the impact of school vacations even less consequential to American pirates than what I reported in EJW.

Figure 3. Share of German schoolchildren on vacation, 2002

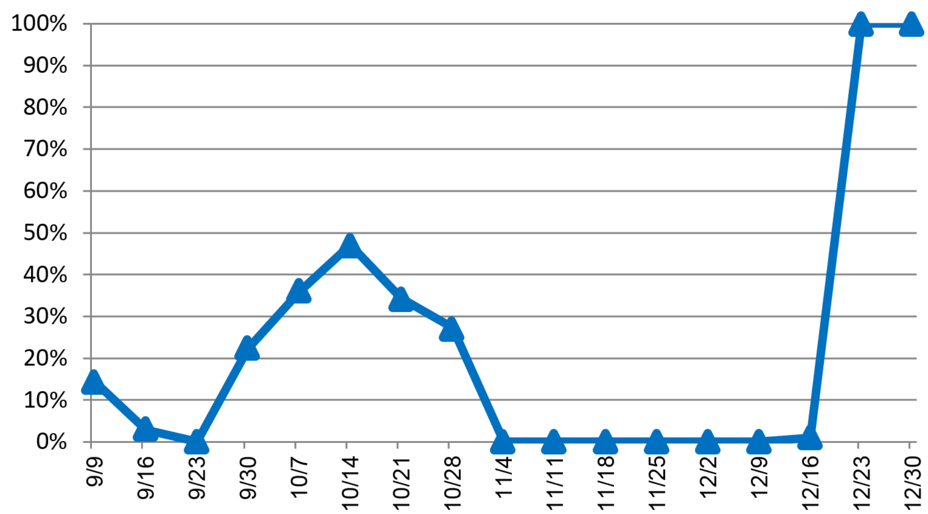

Schoolchildren as a fraction of German file sharers (2.5)

The importance of German school holidays depends in part on the share of German pirate files that is controlled by schoolchildren. To estimate this share, I relied on surveys measuring pirates in the population by age group. I used data from France and the U.S., because I could not find similar German data. These surveys included only individuals more than 11 years of age, since younger children were

7. If we remove the weeks with no weekly change and Christmas, as I did in my original calculation, the average value changes slightly to 13.6 percent. If we use the original, off-by-a-week OS data, the average value would be 14.7 percent, still well below the generous 25 percent value I used in the calculations.

thought to be unlikely to pirate much music. The data indicated that piracy rates for individuals aged 18–40 were of similar magnitudes to those aged 12–17, and there were many more of these older individuals. Using these surveys, and providing extra weighting to age groups with higher piracy rates (to pick up likely intensity differences), I concluded that the share of all German pirate files controlled by students in Germany aged 12–17 was likely to be in the range of 15 percent.

OS do not provide any data on Internet piracy rate by age to contradict my calculations. Instead, they criticize my calculation by making several claims, some of them misleading. They state that my "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" (OS 2016, 63).

Contrary to their first assertion, I did try to account for students younger than age 12. As I wrote in EJW, in my calculations I removed all pirates over the age of 50 in the U.S. and all those over the age of 60 in France, in order to compensate for those (presumably small in number) pirates younger than 12 who were not being measured in the surveys (Liebowitz 2016b, 390). OS's second assertion mischaracterizes my statement about German students not speaking English; that statement was about German students *below the age of 10*, 80 percent of whom had not yet taken any English classes (Liebowitz 2016a, 389–390 n.30).

OS then devote a lengthy paragraph to arguing there was "surprisingly little German music" listened to by Germans (2016, 63). The implication would seem to be that the degree to which Germans listened to German language music played a role in my calculations. *It did not.* Nevertheless, even though language differences are irrelevant to my EJW calculations, OS's assertion that there is a very strong overlap in musical tastes between the U.S. and Germany is inconsistent with the data. Data from the IFPI, the international organization representing record companies, indicate that 45 percent of sound recording sales in Germany in 2002 were of domestic origin, and German-language songs and albums frequently made lists of top ten sellers.[8]

To remain conservative in my calculations, I also made several assumptions likely to overstate the share of pirate files controlled by German students. For example, I assumed that Germany had a uniform age distribution, which overstated by almost 50 percent the relative number of individuals aged 15 compared to individuals aged 35, as I pointed out in EJW (Liebowitz 2016a, 390 n.32).

8. For example, the IFPI has a top ten singles list in 2004 including these songs: "Dragostea Din Tei," "Lebt denn dr alte Holzmichl noch?," and "Augen auf!." Additionally, these albums were in the top ten in 2005: *Noiz*, *Von Hier An Blind*, *Es Ist Juli*, *Willst Du Mit Mir Gehn*, and *Laut & Leise*. I do not have data for 2002 or 2003, but it is clear that a good deal of leading music in Germany at that time was in German and would be of little value to American pirates.

Finally, OS do not use any data to support their claim that half of German pirate files were controlled by schoolchildren. They merely assume that music piracy takes place only among people 24 years of age or less, with half of that group being impacted by school vacation days. They provide no evidence to support this claim, and ignore the plentiful evidence to the contrary. Nevertheless, even if we were to accept the OS claim that secondary students controlled half of the pirated files in Germany, it would increase my estimate by slightly more than three times, so that instead of German school holidays influencing 0.14 percent of American pirate files (or 0.07 percent, as seen in the last section), they would influence about 0.47 percent (or 0.23 percent), still a very small number that is consistent with my conclusions.

The implausibility of first-stage estimates (2.3)

OS tell us in their section 2.3:

> The first-stage estimates in our analysis—the influence of German vacations on downloads in the U.S.—seem large… (OS 2016, 62)

Their usage of the term "large," in this case, is a remarkable understatement.

The first-stage coefficient relating German school holidays and American downloading implies that when OS's average number of German students on vacation go back to school, American downloading would drop by 150 percent of its mean value, if that were possible. This indicates that when all German students are in school (seven of the 17 weeks of the study), American piracy would vanish. The terms *implausible* or *ludicrous* seem more accurate than *large*. Surely, the implication that American piracy ends when German kids return from vacation, when more than 99 percent of pirate files are still available (from pirates who are not returning German students), is absurd.

OS (2016, 62–63) try to explain away this implausible result by appealing to skewness in the data, ignoring my discussion of this topic in EJW (Liebowitz 2016b, 383–384). OS make the obvious point that removing observations with zero downloads—an exercise that I myself carried out and reported on in EJW—will raise the average number of downloads. They argue that after removing these observations, the implied impact on American piracy from kids returning from vacation "is smaller than the [new, higher] mean" and "seems quite plausible." Being smaller than the mean, however, merely implies that American piracy does not drop all the way to zero when an average number of German schoolchildren come back from vacation; in EJW I surmised that the piracy decline would be "still a very large 72 percent" (ibid., 384). So it seems OS are suggesting that in response to German students returning from vacation, any American piracy decline less than 100 percent would be "quite plausible," but this is nonsensical. With over 99 percent of pirate files still available to Americans from other suppliers, it would be far from "quite plausible" for the return of German schoolchildren to have any but a very small impact on American piracy.

Finally, OS argue that German holidays also cause a large decrease in the time it takes for American pirates to download files. But this additional data is not *independent* support because it comes from the same source, those 2002 pirate servers, as their dubious download data. Whatever is causing OS's implausibly large measured relationship between German vacations and American downloads is most likely also causing an implausibly large reported relationship between German vacations and the completion time of American downloads.

The OS instrument, U.S. record sales, and downloads (2.2)

An early version of my EJW article included a demonstration that the German school vacation variable was not independent of American record sales, although that discussion did not appear in the published version. OS (2016) try to counter this point despite its absence from the published EJW article. I did not include this material in the published article because it is *possible* (although I consider it doubtful) that the inclusion of other variables in the regressions might condition German vacations to become independent of American record sales.

What OS did not try to counter, however, was my discussion, included in both the early and published versions of my EJW article, of what appears to be a serious inconsistency in their analysis. OS hypothesized that German school vacations increased U.S. pirate downloads, and their first-stage regressions using disaggregated data (discussed in the previous section) indicated that German vacations enormously *increased* U.S. pirate downloads. But contrary to their very large positive measured relationship between American downloads and German vacations, these two variables, at the aggregate level, displayed a *negative* relationship (because, as seen in Figure 4, the weeks with no German vacations are also when downloads are highest).

Because OS never made their data or code public, their statistical analysis is a black box, and in order to test the accuracy and consistency of this black box, we must compare what went in with what came out (e.g., by comparing the incorrect average value of student vacations derived from the black box with the much lower value derived from aggregate statistics). Thus, to test the consistency of their analysis with regard to German vacations and American downloads, we need to use the weekly vacation data they fed into their black box, and not the vacation data corrected to make the weekly timing consistent. Using this data, weekly aggregate downloads and vacations have a correlation coefficient of −0.71 without the sui generis Christmas weeks, and −0.37 including the Christmas weeks,[9] and frequently have a statistically significant negative relationship in regressions with and without time trends.[10]

Figure 4. German school vacations and U.S. downloads, 2002

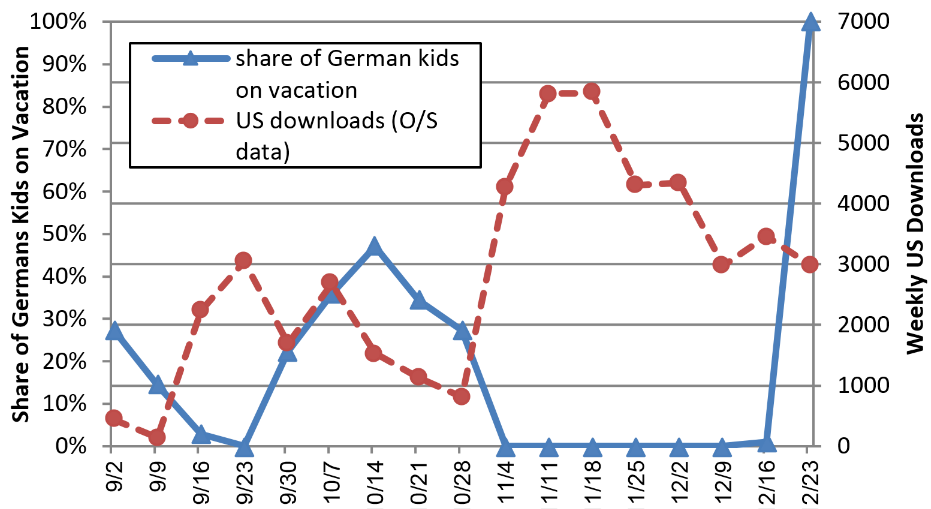

The aggregate weekly relationship, therefore, does not conform with OS's strongly positive measured relationship. The aggregate data, however, are just the sum of the individual album-based data used in OS's 2007 regressions. It does not seem possible for disaggregated downloads of individual albums to have such a strong positive relationship to German vacations while at the same time the relationship between German vacations and aggregated album downloads has a negative relationship (with no variables other than album fixed effects and time

9. These correlations, also reported in Liebowitz (2016b), are stronger than if we use the correctly matched download and vacation data, although the latter relationship is irrelevant for testing the consistency of OS's statistical analysis.

10. Regressing downloads on school holidays provides a negative relationship whether the Christmas week is included or not, but is significantly negative only excluding the Christmas week. Including a time trend (linear, quadratic, or sixth-degree polynomial), with or without the Christmas week, leads to a significant negative relationship in five of six cases, and an insignificant negative relationship in the sixth case.

trends in OS's regressions). This contradiction indicates a potentially serious problem with the black box of their data analysis, as described in greater detail in the section beginning on page 384 in my EJW article.

Quasi-experiments

The 2007 OS paper contained four quasi-experiments that they considered "an important complement" to their main analysis (2007, 4). Since these quasi-experiments used publicly available data, I attempted to replicate those experiments in (Liebowitz 2017). My replications did not support their conclusions in any of the four quasi-experiments. In IEP, OS attempt to defend their results in two of the four quasi-experiments, although they are silent about the other two quasi-experiments. I respond to their claims below.

Piracy and summer music sales (2.6)

As an intermediate step in their first quasi-experiment, OS (2007, 36) claimed that American piracy routinely fell in the summer because college students lost access to their high-speed Internet connections. They used 46 months of BigChampagne data, represented in Figure 5 (with summers circled), to conclude that piracy decreased every summer. This was a necessary result for OS to then conclude that piracy fell in other summers for which they did not have piracy data, a requirement for a second portion of the quasi-experiment where OS examined whether record sales rose in the summer due to the purported decrease in piracy, for all post-Napster years (2000–2005) relative to pre-Napster years.

OS tested whether piracy routinely fell in the summer by running a regression using this BigChampagne data set that contained three complete summers (2003–2005). Their regression, however, included only a single dummy variable to represent all three summers (and a time trend). This effectively made it impossible for them to know whether piracy routinely fell in *each* summer.

In my 2017 article, I demonstrated that piracy did not decrease every summer (Liebowitz 2017, 1–4). Using identical data, I ran the same regression except with separate dummies for each summer. The last of the three summers, 2005, had an economically and statistically significant *increase* in piracy, meaning that a decline in summer piracy was not a regular occurrence. This is the key finding that invalidates their quasi-experiment. But OS, in IEP, ignore my regression results and continue to mistakenly assert that "these summer months represent clear breaks from the growth trend in this period" (OS 2016, 64), when the evidence plainly shows that the third summer, at least, did not represent a clear break with the growth trend.

Instead of responding to the regression evidence I provided, OS focus their discussion on a conjecture I made about whether RIAA lawsuits against pirates might have caused the decline in piracy that took place in the first summer, 2003. The lawsuits provide a competing, idiosyncratic explanation for the decline in piracy taking place that summer, *possibly* invalidating the meaning of the regression results for 2003. Since the lawsuits were idiosyncratic, the 2003 decline in summer piracy, if it was due to the lawsuits, cannot be ascribed to the summers without piracy data, since those summers did not have the introduction of lawsuits. OS seemingly fail to understand this point, since they claim that the reason for the 2003 summer decline was irrelevant to their analysis. It is relevant to their quasi-experiment, however, because you cannot apply an idiosyncratic result to other years. If, by way of contrast, the quasi-experiment had been comparing the summer share of records in 2003, when measured piracy was low, to the summer share in 2005, when measured piracy was high, then the cause of the 2003 piracy decline would be irrelevant, as OS now wish to claim in IEP. But that was not the basis of their quasi-experiment published in 2007.

Figure 5. Number of simultaneous pirates (BigChampagne data)

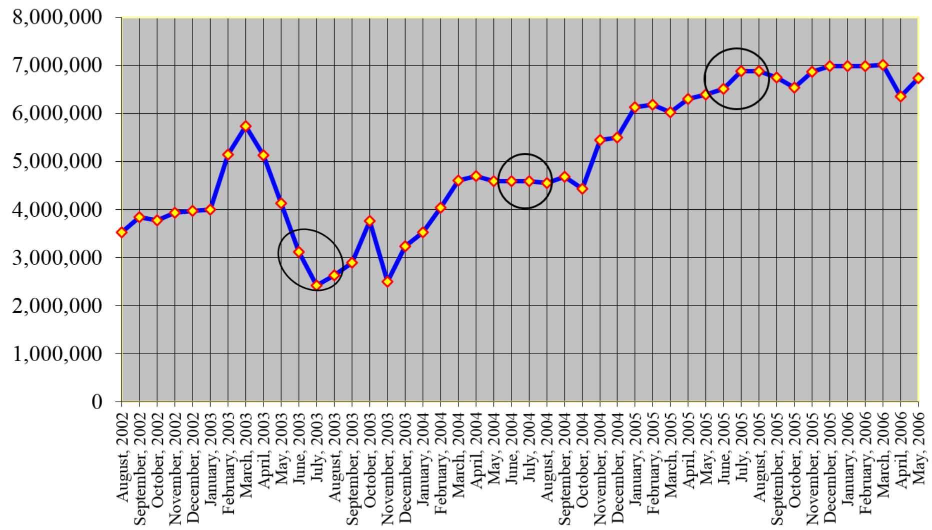

More crucially, however, OS ignore the fact that the 2005 increase in summer piracy, by itself, is sufficient to invalidate their quasi-experiment, and whether the lawsuits of 2003 further invalidate their results is only of secondary interest.

Monthly piracy and record sales (2.7)

In their third quasi-experiment, OS (2007, 36–37) ran a regression relating the monthly level of record sales to the monthly number of pirates, using the same BigChampagne data represented in Figure 5 in the previous section. Although OS found that piracy had a statistically insignificant negative impact on record sales, they claimed the effect was economically small. When I reran their regression (Liebowitz 2017, 6–8) with what was supposed to be identical data, the coefficient was 48 percent larger than theirs. When I included the unemployment rate in the regression, because the unemployed have less money to purchase albums and more time to pirate, the coefficient on piracy quadrupled and became statistically significant.

OS do not mention these results in IEP. Instead, and in my opinion, oddly, they criticize the nature of my very standard attempt to determine the sales decline caused by piracy, for individual years (OS 2016, 64). I ran the same regression they ran, using monthly data from mid-2002 through mid-2006: $Y = a + bX$, where Y is monthly record sales and X is monthly number of pirates.

I then determined how much of a decline the regression coefficient implied, in a given year, by multiplying the regression coefficient by the average value of X (the average 2004 level of piracy, for example) to determine a predicted impact of piracy on sales for that year. OS, in IEP, seem to be disputing this standard calculation. OS seem to be arguing that it is necessary to (go out of sample and) separately determine the decline in albums due to piracy, in each year after Napster (even when we do not have piracy data for those years), and then average those declines to determine an average yearly decline since Napster. Obviously, by incorporating the low assumed piracy values from early years into the 'average,' the average will be lower than using just a later year, such as 2004. Although one might be interested in following this methodology, it is not the way to answer the question in which I was interested: How much did piracy decrease record sales in 2004 (or 2003, or 2005), and how much of the overall sales decline was due to piracy?

How much did piracy hurt the recording industry? (3)

In their final section, OS return to some of their previous themes, minimizing the consequences of piracy for the recording industry.

First, they mix apples and oranges by comparing inconsistently defined "displacement rates," i.e., the extent to which a pirated song replaces a purchased song (OS 2016, 65). They used this same erroneous methodology in their 2010 article. I described those errors in detail, as well as providing a consistent methodology for comparing the "displacement rates" found in piracy studies, in a *Journal of Cultural Economics* article (Liebowitz 2016a), to which I recommend the interested reader turn for more detail. In answer to the question OS pose in their section heading—"How significant was the impact of file sharing, really?"—the correct answer, based on my literature review (ibid.), is 'very significant.'

Second, they claim that five studies, other than their own, using actual piracy data instead of piracy proxies, tended to find minuscule impacts of piracy. The problem with this claim is that OS misrepresent a majority of those studies in a way that exaggerates support for their thesis. One of those studies, by Sudip Bhattacharjee and coauthors (2007), did not investigate the impact of piracy on sales but rather "chart survival," which they found had declined. OS conflate chart survival with sales and then mischaracterize the decline as not being a decline. Remarkably, six years after Bhattacharjee publicly objected to OS's characterization of his study, saying it was "not correct," OS continue to mischaracterize it.[11] Similarly, Robert Hammond (2014) did not investigate the overall impact of piracy on sales, but merely the impact of piracy on individual artists. OS disregard Hammond's description of his own work.[12] Then, OS misrepresent Tatsuo Tanaka's 2004 rough draft as a finished study even though Tanaka himself referred to it as "version 0.1" on the cover page and stated in the conclusion "this research is very preliminary." The paper by Luis Aguiar and Bertin Martens (2016) uses proxies for both record sales and piracy, and should not be in OS's list of papers that are not reliant on proxies. OS also ignore David Blackburn's 2006 paper (which was completed and submitted, unlike Tanaka's 2004 paper), which used piracy data but found a strong negative impact of piracy.

Finally, OS claim that the very large decline in music sales since 1999 was nothing special because a similar decline occurred in the late 1970s to early 1980s. But the most recent decline is much worse than that of the late 1970s and early 1980s, as revealed by data and the memories of anyone who lived through both periods. Much of that earlier decline is an illusion, due to OS's reliance on RIAA reported music *revenues*. RIAA numbers are accurate for units, but not for revenues; revenue figures are derived by multiplying the measured units by the list price, not from measuring actual revenues. Since 1973, real list prices have remained in only two narrow bands, with the list price of an album ranging between $11 and

11. When David Glenn of *The Chronicle of Higher Education* asked Bhattacharjee whether OS's characterization of his paper as having a result similar to that of OS was correct, Bhattacharjee stated: "It is not correct to say that our work shows file sharing is unrelated to changes in sales" (quoted in Glenn 2010).

12. Hammond states: "showing that file sharing is not harmful to individual artists is not inconsistent with the well-documented fact that file sharing does great harm to the music industry" (2014, 406).

$12 (1983$) from 1973 through 1979 and then between $7.50 to $8.50 from 1982 until 2007.[13] The transition from the high price to the low price took place during the apparent revenue decline in the late 1970s and early 1980s. The sales decline beginning in the late 1970s, measured in units, is much smaller than the decline in revenues (18 percent versus 49 percent).

But even using the misleading revenue values, the 1980 decline is considerably less significant than the 2000 decline, per OS's Figure 3. The revenue decline that started in 1979 ended in 1982 (largely because the new lower list price had been reached), a four-year period, after which revenues rose for 17 years, surpassing its prior peak (even with lower list prices) within 10 years (of the trough). The decline that started in 2000 has lasted about four times as long, and 17 years later revenues still haven't begun to recover. Even though in 2015 the population was 40 percent larger and per capita income almost 60 percent larger than in 1982, real revenues had dropped to a level below the bottom level of the early 1980s.

It is also important to note that there is a simple hypothesis, neglected by OS, that is consistent with both the decline in the early 1980s as well as the most recent decline starting in 2000. The decline in the early 1980s coincides with the rise of the first format associated with piracy—cassette tapes, which became the dominant format by 1983. Therefore, the early 1980s sales decline coincided with a rise of piracy (which diminished as CDs became dominant), just as the decline starting in 2000 did.

Conclusion

I believe that the weight of the evidence, including my earlier examination of OS's data and methodology in EJW, my failed replication of their quasi-experiments (Liebowitz 2017), and, as just described, their unwillingness to explain some questionable results and their flawed and misleading responses when they have attempted to explain questionable results (OS 2016), clearly reveals serious problems with the original OS 2007 article. In many of the scientific and biological disciplines, if the original authors cannot provide a cogent explanation of questionable research results, journal editors often label those articles with 'expressions of concern' or retract the article outright. Although such editorial behavior appears to be far from the norm among economic journals, a quick perusal of the *Retraction Watch* website (**link**) should make clear how this works and how standard it has become in other research fields. I believe it would be prudent,

13. I do not know why the list price changed so dramatically at that point in time, but this price decline is the leading cause of the steep decline in RIAA 'revenues' in the period from 1978 until 1982.

and well within the norms of the broader research community, for the editors of the *Journal of Political Economy* to put some such marker on the "The Effect of File Sharing on Record Sales: An Empirical Analysis" so that future researchers and policy analysts will not be building their research programs or policy remedies on the flawed data, analyses, and conclusions found in that article.

Appendix

A document providing the raw data and calculations supporting results in this article is available for download (**link**). A spreadsheet containing the underlying data for the figures in the article is also available for download (**link**).

References

Aguiar, Luis, and Bertin Martens. 2016. Digital Music Consumption on the Internet: Evidence from Clickstream Data. *Information Economics and Policy* 34: 27–43.

Bhattacharjee, Sudip, Ram D. Gopal, Kaveepan Lertwachara, James R. Marsden, and Rahul Telang. 2007. The Effect of Digital Sharing Technologies on Music Markets: A Survival Analysis of Albums on Ranking Charts. *Management Science* 53(9): 1359–1374.

Blackburn, David. 2006. The Heterogeneous Effects of Copying: The Case of Recorded Music. Working paper. **Link**

Glenn, David. 2010. Dispute Over File Sharing's Harm to Music Sales Plays Again. *Wired Campus, Chronicle of Higher Education*, June 17. **Link**

Hammond, Robert G. 2014. Profit Leak? Pre-Release File Sharing and the Music Industry. *Southern Economic Journal* 81(2): 387–408.

Leibowitz, Nathaniel, Aviv Bergman, Roy Ben-Shaul, and Aviv Shavit. 2002. Are File Swapping Networks Cacheable? Characterizing P2P Traffic. Presented at the 7th International Workshop on Web Content Caching and Distribution (Boulder, Colo.), August. **Link**

Liebowitz, Stan J. 2016a. How Much of the Decline in Sound Recording Sales Is Due to File-Sharing? *Journal of Cultural Economics* 40(1): 13–28.

Liebowitz, Stan J. 2016b. Why the Oberholzer-Gee/Strumpf Article on File Sharing Is Not Credible. *Econ Journal Watch* 13(3): 373–396. **Link**

Liebowitz, Stan J. 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). **Link**

Oberholzer-Gee, Felix, and Koleman Strumpf. 2007. The Effect of File Sharing on Record Sales: An Empirical Analysis. *Journal of Political Economy* 115(1): 1–42.

Oberholzer-Gee, Felix, and Koleman Strumpf. 2010. File-Sharing and Copyright. In *Innovation Policy and the Economy*, vol. 10, eds. Josh Lerner and Scott Stern, 19–55. Chicago: University of Chicago Press.

Oberholzer-Gee, Felix, and Koleman Strumpf. 2016. The Effect of File Sharing on Record Sales, Revisited. *Information Economics and Policy* 37: 61–66.

Tanaka, Tatsuo. 2004. Does File Sharing Reduce Music CD Sales? A Case of Japan. Working paper. Link

About the Author

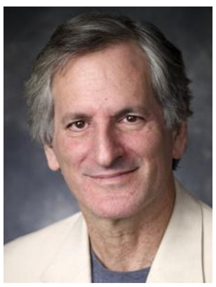

Stan Liebowitz is the Ashbel Smith Professor of Managerial Economics at the University of Texas at Dallas. He has authored many articles examining the influence of new technology on intellectual property, beginning with a study on photocopying's impact, for the Canadian government, and continuing with digital copying. He has also written numerous articles about network effects and a few articles about measuring research performance. Most of his research is available on the SSRN. His email is liebowit@utdallas.edu.

Go to archive of Comments section
Go to May 2017 issue

Discuss this article at Journaltalk:
http://journaltalk.net/articles/5940