

THE KEY INSTRUMENT IN THE OBERHOLZER-GEE/STRUMPF FILE-SHARING PAPER IS DEFECTIVE

**STAN LIEBOWITZ
SCHOOL OF MANAGEMENT
UNIVERSITY OF TEXAS AT DALLAS
APRIL 2010**

I examine the key instrument (German kids on vacation) used by Professors Oberholzer-Gee and Strumpf. Their measured relationship between the instrument and the variable that it is instrumenting for, American downloading, is seen to have outlandish implications, indicating an important error. The instrument is also shown to be related to American record sales, contrary to the requirements of their analysis. The data set used by O/S is biased, considerably overstating the share of German files. Finally, I demonstrate that the instrument must have a de minimus impact on American downloading (and thus American record sales) negating its potential usefulness.

Who are you going to believe, me or your own eyes?
Chico Marx in Duck Soup

In February of 2007 the Journal of Political Economy published an influential article by Professors Oberholzer-Gee and Strumpf (O/S) examining whether file-sharing by Americans reduced sound recording sales in the United States.

Professors Oberholzer-Gee and Strumpf (O/S) compare sales of individual sound recordings to the unauthorized downloads of those sound recordings. It is clear, as O/S acknowledge, that comparing the sales of songs to the number of unauthorized downloads of those songs suffers from a severe simultaneity problem because the popularity of a song influences both its sales and downloads. In an attempt to overcome this simultaneity problem, O/S chose a dazzlingly recondite instrument for filesharing—the number of students on school holidays each week in Germany (or Number of German Students Vacationing: NGSV) during the 17-week period (the last third of 2002) of their analysis. They refer to this as their “most important instrument” because it allows them to use the time element in the panel data analysis. Using this instrument, O/S conclude that file-sharing has no impact on record sales.

What rationale do Professors Oberholzer-Gee and Strumpf provide for their use of NGSV? First, they note that if file-sharing influences sound recording sales then a shock to the amount of American file-sharing will induce a change in American sound recording sales. Next, they claim that the amount of American file-sharing responds to changes in the number of computers using file-sharing services. Finally, they look for a shock to the number of computers on file-sharing networks and claim to have found one in NGSV. O/S assume, although without any supporting evidence, that German students are more likely to keep their computers attached to file-sharing networks during school vacations than on days when they attend school.

O/S further argue that Germany is the second most common source for American downloads, behind only the US itself (although, as I show below, their measurement of American use of German files is much too high). Thus, German school holidays are presumed to have a sufficiently large impact on American file-sharing activity that American record sales would be measurably influenced if file-sharing had an impact on record sales.

Notice that the O/S thesis absolutely requires that computers not be continuously hooked up to file-sharing services because otherwise German vacations would have no impact on German file-sharing usage. It is the fundamental assumption of the entire approach chosen by O/S and I will expound on the implications of this assumption at various portions of the exposition.

In this comment I examine various aspects of NGSV, a variable which is based, fortunately, on publicly available data.¹ The plan of this paper proceeds as follows. First, I examine the construction of the NGSV variable which is essential to understanding issues that arise later in the paper. This examination also shows that German school holidays are related to American record sales in perfectly predictable ways, contrary to assertions of O/S. Next, I present what is essentially a “smoking gun” demonstrating that O/S’ analysis is faulty: their key empirical result, the relationship between NGSV and American downloading, is so unreasonably large that it must be wrong. Third, I examine the German share of files downloaded by Americans which is a key statistic used later. I also discover that the data set used by O/S considerably overstates the share of German files. Finally, I provide an analysis indicating that it is not possible for the impact of NGSV to rise above background noise because the effect of the German school holidays on American downloading (and thus American record sales) must be vanishingly small.

1. The construction of NGSV and its relationship to American Sales

I begin with an examination of NGSV. Figure 1 portrays the relationship between German school holidays and record sales in the US during the last third of 2002, the timeframe of the O/S analysis. The line with triangle markers illustrates the pattern of vacations and the share of students on vacation in a given week.² Each year, the combination of end-of-summer vacations and fall-breaks occur almost exclusively during September and October. These holidays are staggered across German States to avoid peak-load vacation congestion. These holidays, in the

¹ The data from the small file-sharing network used in their main regressions is not publicly available, making replication of those regressions impossible. In a previous comment (Liebowitz, 2007) I argued that many factual claims made by O/S were incorrect and that my attempted replication of their ancillary results (which were based on publicly available data) revealed that their conclusions from their additional empirical tests were invalid. But because O/S did not make their downloading data available to other researchers I could not gauge the validity of their main regressions. This comment focuses on their main regression results.


² The share of German students on vacation in my Figure 1 is very similar to that reported in Figure 1 in O/S (2007). But there are two important differences between their statistics and the statistics I use. Although O/S are mum on how they determined the number of students on vacation, their numbers indicate that they have included all grades, starting with kindergarten, whereas my statistics only include secondary school, starting with grade 5 which is still probably too early a grade because most fifth graders are not yet purchasing music. I also exclude vocational students since students in those schools are usually working part time, usually attend school part time, and usually go to work on school holidays. Whereas O/S include 12.5 million students, my more limited sample includes a mere 6.3 million. If I limit the number of students to those beyond grade 7, which seems to better match the age at which they are likely to become interested in purchasing music, there are only 4.4 million students, or about a third of the number used by O/S. The statistics I used come from an online database run by the German government available at: <https://www-genesis.destatis.de/genesis/online>. I used code selection 21111-0003, Schüler: Bundesländer, Schuljahr, Geschlecht, Schulart.

main, do not apply to Universities or the workplace (except for Christmas). After the first week of November no more school vacations occur until the Christmas break. This means that the vacations occur before the Christmas selling season in the US begins.

O/S claim that German school holidays are not related to US record sales:

For German vacations to be a valid instrument, they must not be directly related to U.S. music demand. This seems likely because the vacation variable varies over time for reasons that are specific to Germany...If file sharing were eliminated tomorrow, German school holidays would have no relation to U.S. record sales.³

The problem with this claim is that it is contrary to the facts—German school holidays are related to US record sales (the line with the circle markers) in a consistent and predictable way.



The high level of American record sales during the pre-Christmas shopping season of November and December occurs when NGSV is at its lowest level. The same pattern of these two variables will be found every year, notwithstanding O/S' claims to the contrary. These facts, if we were to

exclude the Christmas week, impose a negative relationship between US sound recording sales and NGSV of $-.48$ in 2002. Therefore, an important assumption made by O/S is clearly false. Because the Christmas week has both high American record sales and high (all) German students on vacation, including that week turns the overall correlation very slightly positive and the correlation for the last half of the period strongly positive ($.32$). Since O/S use panel data where record sales change every week, and some records are likely to sell mainly early in the period and others mainly later in the period, any strong correlations during portions of their period are potential problems for their analysis.

³ Page 15, O/S, 2007.

2. A Crucial Error: The Smoking Gun

An extremely important error is revealed in O/S' claim about the impact of changes in NGSV on American downloads in the first stage of the instrumented regression results found in their Table 7. If the first stage results are inaccurate, then their final results in the second stage will be inaccurate as well. In the first stage, O/S regress a group of variables, including the key instrument NGSV, on American downloads (O/S measure American downloads from a small file-sharing service that provided them raw data). A correlation between NGSV and American downloads is required, and the stronger the better, if NGSV is to be deemed a useful instrument for American downloads. The relationship they find is all that and more, as described by O/S:

The first-stage estimates imply that a one-standard-deviation increase in the number of children on vacation [NGSV] boosts [American] weekly album downloads by slightly more than one-half of their mean, an effect that is statistically significant and economically meaningful. (O/S, 2007, page 23)

O/S draw the correct inferences from their statistics. Multiplying the coefficient on NGSV, .67 (which varies only slightly between their regression specifications 2 through 4) by the standard deviation of NGSV, 3.6 (found in their Table 5), does provide a value for the change in American downloads, 2.4, that is slightly more (54%) than half the average value of American downloads reported in their Table 5 (4.4).

This purported finding deserves extremely careful attention because its implications can easily slip by. The statement says that American file-sharing activity increases by 50% (from 4.4 downloads to 6.6 downloads) when NGSV increases by one standard deviation (3.6 million students according to their Table 5). This enormous change in American filesharing represents the full panoply of American downloads from the entire world, not just the sliver of American file-sharing activity based on downloaded German files.⁴

We can further investigate the implications of this NGSV coefficient by starting with the average value of NGSV (9.8 million) found in O/S (2007) Table 5 (page 17). If the average value of NGSV were 9.8 million, as reported, then the regression coefficient on NGSV implies that American downloads would fall to zero (it actually predicts *negative* American downloads, although it is not clear there is any meaning to that concept) in those weeks when NGSV drops to zero (seven

⁴ In their "model 5" regression in their Table 7, the coefficient on NGSV is three times as large and thus the implications are even more implausible by a factor of three.

out of the seventeen weeks in their data).⁵ [The average value (9.8) is approximately three times the standard deviation mentioned in the above quote and thus the 50% change becomes a 150% decline.] Think about that for a few seconds and let it sink in. The results do not pass what Kennedy (2003) refers to as the “laugh test.”⁶

Although this example nicely illustrates the implausibility of this coefficient, the average value of NGSV reported in O/S’s Table 5 is incorrect and is very different than the true average value of NGSV, nullifying the specifics of the above example, although the basic logic continues to hold.⁷ Using an average value for NGSV that is consistent with the underlying data, say 2.2 million students (not 9.8), would imply that American downloads would drop by 34% (by 1.5 downloads per album) to 2.9 downloads per album when all German students were in school. On the other hand, American downloading would rise by 156% (by 6.9 downloads per album) to a level of 11.3 (=6.9+4.4) when all German students went on vacation (weekends, summer or Christmas). So, if this first stage regression result were correct, American filesharing would whipsaw from 2.9 downloads per album when German students were in school to 9.1 downloads per album when German students were not in school, such as during the weekend. These prognostications are not out-of-sample extrapolations and thus they should provide realistic results. But they do not.

These predictions are easily falsified by looking at data on American downloads, and by plain common sense. Further, the number of German students supposedly causing this enormous impact on American downloads is amazingly small. As I report in footnote 2, O/S include all ages of students in NGSV, including students in kindergarten and elementary school, which approximately triples the number of students relative to the number of students who are more

⁵ The coefficient in the first stage is approximately .67 and when multiplied by 9.8 implies a change in American downloading of -6.6, which is (absolutely) larger than the average American downloading value of 4.4.

⁶ See Kennedy’s “Ten Commandments of Applied Econometrics” (page 393): “Apply the “laugh” test—if the findings were explained to a layperson, could that person avoid laughing?”

⁷ O/S claim in their Table 5 that the NGSV variable (German kids on vacation using O/S nomenclature) has a mean of 9.855 million and a maximum of 12.491 million (when all students are on Christmas vacation). These values, if true, would imply that, on average, 79% (9.855/12.491) of German schoolkids are on vacation each week during the 17 weeks of observations. Examination of Figure 1 found in Section 1 (or Figure 1 in O/S 2007) makes it clear that it is mathematically impossible for the average to be so high since in no week except Christmas are more than 50% of students on vacation and in 7 of the 17 weeks no German schoolkids are on vacation at all. My calculations indicate that the correct average share of schoolkids on vacation is about 18%, not the 79% reported. Nevertheless, the reported standard deviation for NGSV seems reasonable. I have found that care needs to be taken before accepting any number reported by O/S. For example, in Table 3, the reported “mean number of downloads” for “all genres” does not match either a weighted or unweighted mean of the numbers reported for individual genres and Liebowitz (2007) cites many more similar examples.

likely to be interested in music and file-sharing.⁸ Dividing by three to approximate the true number of students that are in school and of the age to engage in file-sharing,⁹ this O/S regression result implies that, from average levels, it would take a mere 2 million Germans in this cohort turning off their computers to completely eliminate US file-sharing even though approximately 100 million other file-sharers are still online with their files available to American file-sharers.¹⁰ It predicts that a power failure in a portion of Germany, or any event that caused German students to turn off their computers, would completely eliminate American file-sharing. How realistic is that?

Making the size of this coefficient even more incredible is that fact that, as we will see below, the German school day takes place when Americans are sleeping and the repertoire of music in which Germans are interested is quite different than the music that interests Americans. The magnitude of the predicted changes in American downloading are so implausible that it would not be possible to believe in a value even one tenth of the size claimed.¹¹

This demonstrates that there is something very wrong with the O/S regression results. Economists, and everyone else, must reject regression results when they are unreasonable.

It is always possible that the coefficients and standard errors were incorrectly reported by O/S in their tables. This would fit the pattern of errors found in some other tables (see footnote 7) and would be consistent with misreported results for other of their tests, as discussed in Liebowitz (2007). There is no way for me to check this possibility, however, since I have been unable to gain access to their data. But it is not possible that fixing a mere typo or clerical error could restore their conclusions because it is unreasonable to think that NGSV could have a measurable impact on American file-sharing or American record sales, as I demonstrate in Section 4. Before the calculations in Section 4 can be undertaken, however, we need to

⁸ Although I am sure O/S had impeccable reasons for including first graders in the sample, they did not explain this choice. Their use of an inflated-size group shouldn't bias their results as long as the share of students in different grades is similar across the German States, which seems likely. But their coefficient (and standard error) will be watered down by the extra non-file-sharing students being counted. We would expect a more accurate but smaller student population to generate, in the otherwise same data set, a larger coefficient for the smaller music-interested vacation-impacted cohort of students.

⁹ These are students beyond grade 7, as discussed in footnote 2.

¹⁰ Although measures of the number of file-sharers are imprecise at best (see Liebowitz, 2006) data from OECD (2004) indicate that Americans made up 55% of worldwide file-sharers (see Table 1 below) and that there were 60 million American file-sharers (page 89 OECD) leading, very roughly, to 100 million file-sharers.

¹¹ If all Europeans turned off their computers, which is far more than ten times the number of computers discussed in the text, we know that American file-sharing would not drop to zero.

examine the share of German files downloaded by Americans, a statistic required for the calculations in Section 4.

3. Germans are Overrepresented in the O/S database

An important concern with the O/S methodology is whether their sample of American downloading behavior is representative of overall downloading behavior by Americans. The central and most basic question would seem to be whether the impact of NGSV that O/S might measure, based on the network for which they have data, would be representative of NGSV's impact for file-sharing networks more generally. For example, if their sample had a much larger share of German file-sharers than found on other file-sharing networks then O/S might find that NGSV has a much larger impact on the Americans in their sample than it would have on typical American file-sharers. In this case their overall methodology might provide misleading results because American record sales are influenced by the full set of American downloaders, not just the small sample of American downloaders in the O/S data set.

Surprisingly, although O/S exert considerable effort examining whether their sample is representative of the full population of songs, they somehow overlooked whether their sample was representative of the full geographic population of users. Nevertheless, they claim that their dataset is representative in a general sense:

An important question is whether our sample is representative of data on all P2P networks...On the basis of these tests, we conclude that our sample is representative of the file transfers on the major P2P networks during our study period. (O/S, 2007, p. 7)

But their conclusion is refuted by their own statistics. The first two columns of Table 1 are from Table 2 of their 2007 paper.

	[1] Share of World Internet Users CIA (O/S, 2007, Table 2)	[2] [O/S dataset] Share of World File- Sharers (O/S 2007, Table 2)	[3] 2003 Share of World Filesharing Users (OECD 2004, p. 190)	[4] [O/S dataset] Countries US downloaded from (O/S 2007, Table 2)	[5] 2002 Broadband Share of Internet Users (OECD 2009)	[6] Share of domestic repertoire 2002; (IFPI)
United States	27.4%	30.9%	55.4%	45.1%	12.8%	92%
Germany	5.3%	13.5%	10.2%	16.5%	8.5%	45%
Ratio US/Germany	5.2	2.3	5.4	2.7	1.5	

The first column repeats information that O/S take from the CIA World Factbook showing that there were more than 5 times as many American Internet users as German Internet users. Yet, the data reported from the limited O/S file-sharing dataset, found in column 2, has only 2.3 times as many American file-sharers as German file-sharers. It is possible that German Internet users might be more likely to engage in file-sharing than are American Internet users, but this is refuted by data in the third column which come from, to quote O/S, “the authoritative BigChampagne database” as reported by the OECD.¹² The Big Champagne data reveal that there are 5.4 times as many American file-sharers as German file-sharers so that the ratio of Americans to Germans in the full population of file-sharing activities is 2.4 times as high as the ratio in the sample that is the O/S dataset.

Although O/S claim that German computers have an unusually large share of high speed connections which, if true, could increase the downloading of German files by Americans beyond the German share of file-sharers, this O/S claim is not supported by any evidence, not even the evidence they cite¹³ and is contrary to the broadband statistics in column 5 of Table 1.

Therefore, the file-sharing dataset used by O/S has an unrepresentatively high share of German users and the scaling factor measuring the degree to which Germans are overrepresented is 2.4. This implies that if the O/S statistics on the share of German files downloaded by Americans, 16.5%, were correct *for their data set*, we would need to scale it down by 2.4 to make it applicable to the overall population, implying that the full population of American file-sharers download approximately 7% of their files from Germans.

Nevertheless, it is difficult to believe that even this lower figure is not too high and that there is not something else amiss with the O/S dataset, statistics, or the basic assumption about computers being unattached from file-sharing networks when the operators are not using them. For example, the numbers in column 2 indicate that there are 2.3 times as many Americans as Germans in the O/S database while the numbers in column 4 indicate that Americans download 2.7 times as many files from Americans as from Germans according to the

¹² See O/S (2007), footnote 9.

¹³ O/S apparently misread the Nielsen NetRatings article (http://www.nielsen-online.com/pr/pr_021021_germany.pdf) upon which they rely (it provided statistics on the portion of *Internet users* who have various forms access whereas O/S report the statistics as if they represented the portion of the *entire population*). Also, the article included relatively slow ISDN connections as high speed connections although ISDN is almost never classified as high speed because it is barely faster than dial-up (as confirmed by the pages and time per person statistics in the report). Finally, O/S’ own statistics indicate that Germans do not have above normal downloads (relative to the number of file-sharers) as would be expected if Germany had a large share of high speed connections whereas the US does have above normal downloads relative to its number of file-sharers, consistent with its higher broadband rate.

O/S dataset. These statistics indicate that Americans are only slightly more likely to download files from other Americans than they are to download files from Germans. Given the time zone differences (Germans are asleep during prime American file-sharing times as discussed in Section 4c), the important repertoire differences between Germany and the US, (discussed in Section 4d) and the fact that the US has a broadband penetration rate that is 50% higher than Germany's rate (column 5 of Table 1¹⁴), it is hard to believe that Americans would not download a much higher share of their files from other Americans if everyone turns off file-sharing programs when not using them, as is required by the O/S methodology.¹⁵

But there is an even more fundamental problem with their analysis, to which we now turn.

4. The Irrelevance of NGSV to American Downloads or Why the Plan of their Paper could not Work.

If NGSV is to be a practical instrument for American downloading, the impact of German school holidays on the number of files available to Americans and the resulting change in American downloading must be large enough to rise above the background noise. Indeed, O/S justify their use of German school holidays by noting the large (Internet) population of Germany. Their claim that Germany provided more files to American file-sharers than any other foreign country, although we now know it to be biased, nevertheless helped make plausible the possibility that a relationship between German school holidays and American downloads would rise above the background statistical noise. That is presumably why German school holidays were examined and not, say, school holidays in the much less populous Luxembourg.

But the share of German files controlled by German secondary school students is in fact quite small. Further, the typical school holiday affects only a small portion of German students. Additionally, only a portion of the files made available to Americans by Germans are of value to Americans due to both time zone differences and repertoire differences. The net result of these

¹⁴ I created the numbers in column 5 by combining broadband penetration over the entire population with Internet penetration for the population to calculate broadband penetration of the Internet population. The 2002 numbers found here: 4th quarter <http://www.oecd.org/dataoecd/63/53/41551452.xls> "Broadband penetration, historical time series (June 2009)" and the Internet penetration numbers are found here: www.oecd.org/dataoecd/19/46/34083096.xls "Households with access to the Internet (1), 2000-08".

¹⁵ Additionally, Canadians and Americans are in the same time zones, have very similar musical tastes and Canadians have almost triple the German broadband penetration. These facts make the low share of Canadian files downloaded by Americans from in the O/S database suspicious.

multiple factors is that the school holiday variable used by O/S has a vanishingly small impact on the number of files available to Americans.

a. Student File-sharers as a Fraction of German File-sharers

O/S present no evidence and entertain no examination of the share of German files that can be attributed to German students, although that would seem to be a crucial factor in the analysis they propose.¹⁶ Although I was unable to find statistics for the file-sharing population distribution Germany, I was able to find estimates of the file-sharing population distributions from the US and France, which I assume to be somewhat similar. These statistics, found in Table 2, indicate that youths aged 12-17 make up a fairly small share of the total file-sharing population.

The numbers in Table 2 clearly indicate that secondary students are not the only users, or even the main group of users of file-sharing networks. In the US, individuals between 18 and 29 years of age are just as likely to use these networks, and there are more individuals in this age bracket than there are in the in the younger age bracket. In France, individuals aged 18-39 are more likely to use file-sharing networks than are secondary students and there are many more individuals in these older age groups than in the 12-17 age group.

Table 2: File-sharing Usage by Age Group			
US Oct 2002; % of population *		France June 2003, % of Internet Users**	
12-17	41% ***	12-17	31%
18-29	41%	18-24	47%
30-49	21%	25-39	31%
50-64	8%	40-59	22%
65+	3%	60+	11%
*Comes from Pew file at: http://www.pewinternet.org/~media/Infographics/Trend%20Data/January%202009%20updates/UsageOverTime_Mar_11_2009%20-%20Read%20Only.zip			
**Table 5.11, p 195, OECD Info Tech Outlook 2004			
*** p 194. OECD Info Tech Outlook. Pew numbers usually only include those above 18 years of age. The OECD document provides a 2001 Pew value for those aged 12-17 as a share of Internet users which I adjusted to share of overall population.			

With a few simple assumptions these figures can be used to estimate the share of all file-sharers who are between 12 and 17 years of age. I assume that the overall age distribution is uniform (in Germany this would overcount younger Germans relative to the actual

¹⁶ Remember that these holidays do not affect university students.

distribution¹⁷) and that the file-sharing population ends at 60 or 64 years of age. After applying these assumptions I find that adolescents between 12 and 17 years of age represent about 18% of file-sharers in the US and 13% in France. I will take the average of these two numbers to represent the share in Germany: 15.5%, or approximately one out of seven German file-sharers.

By focusing on student vacations, O/S ignore the vast majority of German file-sharers who are unaffected by school holidays. The potential shock from changes to NGSV affects a much smaller number of files than is implied by the statistics put forth by O/S purporting to measure the importance of German files to American downloaders.

Another important fact ignored by O/S is that a majority of German students go into the less academic tracks known as *Hauptschule* or *Realschule* and these students leave school at grades 9 and 10 respectively. This would mean that most youths of 16-19 years of age are not affected by school vacations. Nevertheless, I will ignore this factor to be as generous as possible in judging the O/S methodology.

One other very simple factor ignored by O/S is that students go to school only on weekdays. File-sharing that occurs on the weekend cannot be affected by school vacations. Thus, if file-sharing were uniform during the week, school holidays would affect only 71% (5/7) of possible file-sharing days.

In sum, school holidays affect the 15.5% slice of the file-sharing population representing secondary students only 71% of the days in a week, leading to a potential reduction in German files due to these holidays of 11%, or about 1/9th of the German files available to Americans in a week.

But even this figure is an overstatement of the share of German files impacted by school holidays.

b. How large is the typical supply “shock” due to school holidays

Figure 1 demonstrated that for the most part the share of kids on vacation was on the order of 20%-40%, with an average of 30% for all weeks with non-zero vacations.¹⁸ Since the vacation

¹⁷ The German population is skewed toward older individuals. See <http://www.nationmaster.com/country/gm-germany/Age-distribution> .

¹⁸ This calculation includes several weeks with very low shares as well as the Christmas week with a share of 1. If all weeks are included, the average share of students on vacation is 18%. If weeks with a share of students on vacation below 3% are dropped and Christmas is dropped the average is 28%. If only the very small (<3%) weeks are dropped the average is 37%.

times are adjusted by the German government to avoid peak-load travel problems in national vacations, as noted by O/S, it is not surprising that the share of students on vacation in any one week is not very large. But this also means that the ‘shocks’ brought about by the NGSV variable are only about 1/3 of what you would expect if all the students went on vacation at the same time.

If, as reported in the previous section, only 1 out of 9 German files could be affected if all students went on vacation in a week, and if only 1/3 of the students are affected in an average week when vacations take place, then the size of the typical vacation shock on German files would be 1/27, or about 3.5%, of the number of German files typically available on file-sharing networks. I should also note that since the purpose of these stuttered vacations is to smooth vacation travel, we can presume that some students are expected to travel. If so, they (and their families) are likely to have their computers turned off the entire time, lowering the files available to Americans, quite the opposite of what O/S assume.

c. Are the extra German files available when Americans are likely to download them?

O/S tell us that files from German students on vacation are particularly accessible to American downloaders: “These [German] kids can stay up later when out of school allowing them to engage in file sharing during the peak U.S. hours (early evening, est [sic]). Vacations also provide more overall time for file sharing.”¹⁹ O/S do not actually spell out the details behind this claim. Let’s take a look.

Table 3 shows the correspondence (using 24 hour, or what Americans refer to as “military time”) between the time in Germany and the time in the population center of the US, the Central Time Zone.²⁰ The prime download time in the US is represented by the shaded boxes in the upper portion of Table 3 (assumed to be 7pm-11pm). The lower shaded portion of Table 3 indicates when German students would be in school and therefore represents the time periods

¹⁹ See page 14, O/S 2007.

²⁰ The contiguous 48 US states span four time zones—Eastern, Central, Mountain, and Pacific. Obviously different American time zones provide slightly different correspondences with German time. The population-weighted average time zone for Americans, according to the US Census, is the Central Time zone <http://www.census.gov/geo/www/cenpop/meanctr.pdf>.

that are directly affected by vacations. I assume that the German school day runs from 8am local time to 4pm local time although in many instances students finish by 1pm (13:00).²¹

The German school day corresponds to the period 1am-9am in the United States. Most American downloaders will be sleeping during the first two thirds of this period since few

Table 3: Correspondence of US (Central Time Zone) and German Time

	US Time	German Time
	17:00	0:00
	18:00	1:00
Key US Download Time	19:00	2:00
	20:00	3:00
	21:00	4:00
	22:00	5:00
	23:00	6:00
	0:00	7:00
German Kids in School or on Vacation	1:00	8:00
	2:00	9:00
	3:00	10:00
	4:00	11:00
	5:00	12:00
	6:00	13:00
	7:00	14:00
	8:00	15:00
	9:00	16:00

Americans stay up past 1am on weeknights (remember that only weekdays are affected by German school holidays) and it is estimated that less than 2% of the population is awake during the night due to insomnia.²²

Most Americans will be rising and going to school or work during the last portion of this period. It is obvious that American downloading behavior cannot be seriously affected during periods of time when almost no Americans are using their computers (computers are presumed shut off when not directly used, as required by the fundamental O/S assumption).

O/S, in the above quote, claim that German students will stay up later when on vacation and that the hours prior to their later bedtime coincide with the prime download time in the US, although no supporting evidence is provided. There actually is some support for a version of this story. According to a survey of German children (Loessl et. al., 2008), the average German secondary

student goes to bed at about 10:30 pm on weeknights (3:30 pm in the US) and about 12:30 am on weekends (5:30 pm in US), leading to a change in German bedtimes of about 2 hours. Assuming that the bedtimes on school holidays are like those on weekends, the vacations would cause a two hour delay in the bedtimes of German students relative to school nights.

Note that even this change in bedtimes (from 3:30pm to 5:30 pm in the US) occurs prior to the prime American download time. But let's stack the deck strongly in favor of the O/S story and

²¹ This information comes from 'typical schedules' found Wikipedia's discussion of the German School Day under the entry "Education in Germany" searched on April 22, 2009.

²² It is estimated that about 13% of adults have insomnia and average 1 hour a night less sleep (12% less sleep) than non-insomniacs. This works out to an average of 8 minutes a night for the population, which works out to an average of about 1.7% of the American population awake during this 8 hour period. See <http://www.healthcures.com/sleeping-dissorders/insomnia-what-is-it-and-how-many-of-us-does-it-affect/>.

assume that the entire two-hour delay in German school kids' bedtimes during vacations occurs during the key American download period. This would mean that American downloaders would, during their main 4 hour download period, encounter additional files from German students during half of that period. We can make an additional very generous (to the O/S thesis) assumption that Americans download files only during this 4 hour window, which would mean that this later German bedtime increases by 50% the files from German students available to Americans.

Results in Sections 4a and 4b indicated that school kids on vacation could increase German files available to Americans by about 3.5%. Now we find that at most, due to time zone differences, the extra files available to Americans are only half the maximum potential. The net effect of all three impacts (share of population, size of holiday shock, and the time zone differences) is to limit the increase the Germans files available to Americans due to school holidays by approximately 1.8% (=50%*3.5%).

d. Differences in the German and American repertoires.

Germans do not listen to the same music as Americans.²³ Often, they listen to music from German performers performing in German. The simplest proxy for this difference is to look at the share of domestically produced sound recordings in the market. The IFPI reports that German domestic repertoire was 45% of total sales in 2002 (as found in Table 1).

Examining 2 years of top-10 albums charts supports this view that a large portion of the music that Germans purchase is not the same as the music that Americans purchase. In the same IFPI document we find for 2004 and 2005 (the only years reported) that only 5 of the 20 top US albums (two years of top-10) also appeared in the top-10 lists of German albums, and 7 of the 20 top German albums were in the German language. Thus, it is probably fair to state that somewhat less than half of German songs are not of interest to Americans, and this is presumably also true for German songs on file-sharing networks. This reduces the 1.8% figure to approximately 1%, representing the share of all German files that constitute vacation-influenced German files of interest and availability to Americans.

²³ This is not withstanding the O/S (2007) claim found on page 21 that "Fortunately, there is substantial overlap between American and German musical tastes."

e. How important is NGSV to American File-sharing?

Returning full circle, we can begin with the German share of files downloaded by Americans. Although O/S claimed a value of 16.5% in their sample, we have seen that 7% is a more reasonable although probably still overstated number.

We have also deduced that an overly optimistic calculation of the change in German files due to school holidays is to increase the number of German files available to Americans by 1%. The calculation for the percentage change in *all* files available to Americans caused by German school holidays is simple: 1% of 7%, or .07%, less than *one tenth of one percent*. Remember that these calculations were performed with numerous assumptions beneficial to the O/S thesis and that the calculation would be considerably smaller without the generous assumptions made throughout the analysis.

There is a more intuitive metric for this value of .07%. O/S claim that the average download time for a music file in their sample, counted from the download request to the completed download, is 1496 seconds or about 25 minutes.²⁴ If the extra German files increased the speed with which Americans can find and download songs in proportion to the increased quantity of files, the savings in time would be one second (1.04) out of the original 1496 seconds.²⁵ Is it even conceivable that American downloaders could perceive a change this small without the careful use of a stopwatch? Surely, common sense tells us that if the additional files from German school vacations cannot be noticed by American downloaders then it cannot affect their downloading behavior. Nor can it rise above the background statistical noise or impact American record sales.

5. Conclusion

Professors Oberholzer-Gee and Strumpf argue that NGSV is an appropriate instrument for their regression analysis. They find that NGSV has a powerful and significant impact on the number of American downloads in their first stage regressions. Then they fail to find any impact of file-sharing on record sales in the second stage.

Unfortunately for their analysis, NGSV is not an appropriate instrument since its relationship to American downloading must be vanishingly small. Their defense to this claim would likely have

²⁴ Calculated from O/S (2007) Table 6, sum of items in the last row, for the first three columns.

²⁵ This is in direct contradiction to a claim in O/S (2005, p. 19) that “a one standard deviation increase in the number of German kids on vacation reduces the time to download a song [by Americans] by about ten percent.”

been that if NGSV had a *de minimus* relationship to American downloads then why was it statistically significant in their first stage regressions? This reply fails, however, since their first stage regression results are clearly in error, predicting preposterously large changes in US file-sharing.

Without conducting a detailed autopsy on their paper it is not possible to know the causes of their erroneous estimates. Of course, I have shown that their approach was doomed to failure in the first place, so the source of their errors might not seem to matter. Nevertheless, although I have shown that their data set is not representative of the population they wish to represent and that their key instrument is related to American record sales for reasons outside their hypothesis, I would not want the reader to think that these problems are responsible for the unbelievable and untrustworthy estimates found in their paper.

The magnitude of the gulf between their coefficients and any remotely plausible results indicates a problem with their analysis that seems far more profound. It would seem that there must have been errors in creating their data, or errors in the analysis, or both. Such a conclusion implies that all their results and statistics should be dismissed as being unreliable.

References

IFPI The Recording Industry World Sales 2005. London: IFPI 2006.

Kennedy, Peter. A Guide to Econometrics, 5th Edition, MIT Press, Cambridge 2003.

Liebowitz, S. J. 2006. File-Sharing: Creative Destruction or Just Plain Destruction? *Journal of Law and Economics* 49(1) 1-28.

Liebowitz, Stan J., How Reliable is the Oberholzer-Gee and Strumpf Paper on File-Sharing? (September 2007). Available at SSRN: <http://ssrn.com/abstract=1014399>.

Loessl, B., G. Valerius, M. Kopasz, M. Hornyak, D. Riemann and U. Voderholzer, "Are adolescents chronically sleep-deprived? An investigation of sleep habits of adolescents in the Southwest of Germany", *Child: Care, Health and Development*, Volume 34 Issue 5, Pages 549 – 556, Published Online: 28 Jun 2008.

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

Oberholzer-Gee, Felix, and Koleman Strumpf. 2005. "The Effect of File Sharing on Sales: An Empirical Analysis." Manuscript, Harvard Bus. School and Univ. North Carolina at Chapel Hill.

OECD. 2004. OECD Information Technology Outlook 2004. Paris: Organisation Econ. Cooperation and Development.