The Effect of File Sharing on Record Sales
An Empirical Analysis*

Felix Oberholzer-Gee
Harvard Business School
`foberholzer@hbs.edu`

Koleman Strumpf
University of Kansas School of Business
`cigar@ku.edu`

October 2006

Abstract

For industries ranging from software to pharmaceuticals and entertainment, there is an intense debate about the appropriate level of protection for intellectual property. The Internet provides a natural crucible to assess the implications of reduced protection because it drastically lowers the cost of copying information. In this paper, we analyze whether file sharing has reduced the legal sales of music. While this question is receiving considerable attention in academia, industry and in Congress, we are the first to study the phenomenon employing data on actual downloads of music files. We match an extensive sample of downloads to U.S. sales data for a large number of albums. To establish causality, we instrument for downloads using data on international school holidays. Downloads have an effect on sales which is statistically indistinguishable from zero. Our estimates are inconsistent with claims that file sharing is the primary reason for the decline in music sales during our study period.

* We thank Bharat Anand, Gary Becker, Shane Greenstein, Austan Goolsbee, Alan Krueger, Steven Levitt, Tom Mroz, Alan Sorensen, Joel Waldfogel, Steven Wildman, and two anonymous referees. We also received valuable comments from participants at the 2004 AEA meeting and seminar participants at Cornell University, Duke University, the Federal Communications Commission, the Federal Trade Commission, George Mason, Harvard Business School, Kenan-Flagler Business School, Marshall School of Business, the NBER University Research Conference, SUNY-Buffalo, Syracuse, Tulane, the University of Arizona, the University of Chicago Graduate School of Business, the University of Florida, the University of North Carolina, the University of Michigan, the University of Pennsylvania Law School, and the University of Zurich. This project would not have been possible without the assistance of several individuals and organizations. MixMasterFlame and the FlameNap network shared P2P data with us, and BigChampagne LLC, the CMJ Network, Nathaniel Leibowitz, and Nevil Brownlee generously provided auxiliary data. We thank Keith Ross and David Weekly for assistance in understanding the KaZaA, OpenNap, and WinMX search protocols and database indices. Sarah Woolverton's tireless efforts to improve the quality of our song matching algorithm and Christina Hsiung Chen's research assistance is also appreciated. The financial support of the George F. Baker Foundation (Oberholzer-Gee) and the Kenan Faculty Fund (Strumpf) is gratefully acknowledged. Aural support from Massive Attack, Sigur Ros and The Mountain Goats is appreciated.

I. Introduction

File sharing is now one of the most common online activities. U.S. households swap more than 300 million files each month, a figure that has grown by over 50% in the last two years (Karagiannis et al., 2004; Billboard, 2006). Sharing files is largely non-rivalrous because the original owner retains his copy of a downloaded file. The low cost of sharing and significant network externalities are key reasons for the dramatic growth in file-sharing. While few participated prior to 1999, the founding year of Napster, in 2006 there were about ten million simultaneous users on the major peer-to-peer (P2P) networks (BigChampagne, 2006). Because physical distance is largely irrelevant in file sharing, individuals from virtually every country in the world participate.

There is great interest in understanding the economic effects of file sharing, in part because the music industry was quick to blame the phenomenon for the recent decline in sales. Between 2000 and 2005, the number of CDs shipped in the United States fell by 25% to 705 million units (RIAA, 2006). Claiming that file sharing was the culprit, the recording industry started suing thousands of individuals who share files. The industry also asked the Supreme Court to rule on the legality of file-sharing services, a question which critically hinges on the "market harm" caused by the new technology. Congress is currently considering a number of measures designed to counter the perceived threat of file sharing.

While concerns about P2P are widespread, the theoretical effect of file sharing on record sales and industry profits is ambiguous (Bakos et al, 1999; Takeyama, 1997; Varian, 2000). Participants could substitute downloads for legal purchases, thus reducing sales. The inferior sound quality of downloads and the lack of features such as liner notes or cover art perhaps limit such substitution. Alternatively, file sharing allows users to learn about music they would not otherwise be exposed to. In the file sharing community, it is common practice to browse the files of others and discuss music in file server chat rooms. This learning may promote new sales. Other mechanisms proposed in the theoretical literature have unclear effects on sales. Individuals can use file sharing to sample music, which will increase or decrease sales depending on whether users like what they hear (Shapiro and Varian, 1999). The availability of file sharing could also change the willingness to pay for music – it could either decrease it due to the ever present option of downloading, or it could increase it through network effects and the greater ease of sharing (Takeyama, 1994; Liebowitz, 1985). Finally, it is possible there is little effect on sales. File sharing lowers the price of music, which draws in low-valuation individuals who would otherwise not have purchased albums. Rob and Waldfogel (2006) find in a recent survey that college students value albums they purchased in the store at $15.91. In contrast, respondents' willingness to pay for albums they downloaded was only $10.66, a value below the average purchase price of a CD.

With no clear theoretical prediction, the effect of file sharing on sales is an empirical question.[1] Most of what we know about the effects of file sharing is based on surveys. The evidence is mixed. File sharers generally acknowledge both sales displacement and learning effects, and it is unclear if either effect dominates. Rather than relying on surveys, this study is the first to use observations of actual file-sharing behavior of a large population to assess the impact of downloads on sales. Our dataset includes 0.01% of the world's downloads (1.75 million file transfers) from the last third of 2002, a period of rapid growth in file sharing. We match audio downloads of users in the United States to a representative set of commercially relevant albums for which we have concurrent weekly sales, resulting in a database of over ten thousand album-weeks. This allows us to directly study the relationship between downloads and sales. To establish causality, we instrument for downloads using international school holidays, a supply shock that is plausibly exogenous to sales.

We find that file sharing has only had a limited effect on record sales. After instrumenting for downloads, the estimated effect of file sharing on sales is not statistically distinguishable from zero. The economic effect of the point estimates is also small. When considering the policy implications of these results, it is important to take into account the precision of our estimates. Based on all specifications presented in this paper, even our least precise results, we can reject the hypothesis that file sharing cost the industry more than 24.1 million albums annually (3% of

---

[1]The entertainment industry's opposition to file sharing is not a priori evidence that file sharing imposes economic damages. The industry has often blocked new technologies which later become sources of profit. For example, Motion Picture Association of America President Jack Valenti argued that "the VCR is to the American film producer as the Boston strangler is to the woman home alone" (Congressional Hearings on Home Recording, 12 April 1982). By 2004, 72% of domestic industry revenues came from VHS and DVD rentals or sales (DEG, 2005; MPAA, 2005). Other examples include the record industry's initial opposition to radio in the 1920s and 1930s and to home taping in the 1980s.

sales and less than one third of the observed decline in 2002). Models that consider the dynamics of file sharing allow us to make more precise statements. For example, if we account for the growth in file sharing during our study period we can reject a null that P2P displaced more than 12.4 million in CD sales or about 15% of the 2002 decline. These results continue to hold after permitting downloads to influence sales with a lag, omitting data from the holiday shopping season, and allowing the impact of downloads to vary by album popularity. In total the estimates indicate that the sales decline over 2000-2002 was not primarily due to file sharing. While downloads occur on a vast scale, most users are likely individuals who in the absence of file sharing would not have bought the music they downloaded.

Our conclusion is supported by other data and methods of analysis. For instance, in the most recent U.S. consumer expenditure survey (2004), households without a computer, who seem unlikely to engage in file sharing, report that they reduced their spending on CDs by 43% since 1999. Quasi-experimental evidence on the long-term effect of P2P on music sales also lead to similar results. For example, we document that the share of sales during the summer months when fewer students have access to high-speed campus Internet connections did not change as a result of P2P. Similarly, sales did not decline more precipitously in the Eastern Time Zone of the United States where P2P users can more conveniently download files provided Europeans. Using several years of data, we also show that the number of P2P users is not correlated with album sales. Finally we document that the recording industry often experiences sales reductions, including a recent episode with a sharper reduction than the current period. These experiments are an important complement to our micro-data results. While the main estimates focus on high-frequency variation over several months, the experiments focus on long-term trends using data spanning ten years.

Our results have broader implications beyond the specific case of file sharing. A longstanding question in economics concerns the level of protection for intellectual property that is necessary to ensure innovation (Posner, 2005). Economic research on the role of patents and copyrights likely began with the critique in Plant (1934) and continues today in the debate between Boldrin and Levine (2003) and Klein, et al. (2002). We provide specific evidence on the impact of weaker property rights for the case of a single industry, recorded music. The file-sharing technology available in 2002 had markedly lowered the protection that copyrighted music recordings enjoyed, so it is interesting to analyze to what extent this reduced protection adversely affected sales. For our study period, we do not detect a significant impact. The paper also contributes to a growing literature which studies the interactions between the Internet and brick and mortar economies, which follows the seminal empirical work of Goolsbee (2000).

The outline of the paper is as follows. The next section provides an overview of the empirical literature. Section III describes the mechanics of file sharing, and we discuss our data in Section IV. Next we describe the econometric approach. Section VI presents the results, and the last section discusses the implications of this study.

II. The Literature

Empirical research on file sharing and record sales has been limited and inconclusive, primarily, we believe, due to shortcomings with the data. Most of what we know about the effect of file sharing on sales is based on phone surveys. There are numerous industry studies which arrive at a diverse range of conclusions. For instance, Forrester Research (2002) and Jupiter Media Metrix (2002) find neutral or positive effects, while the International Federation of the Phonographic Industry (2002), Edison Media Research (2003) and Forrester Research (2004) document a sales displacement. A general difficulty with these studies is that they compare the purchases of individuals who download files with the purchases of those who do not. While downloaders may in fact buy fewer records, this could simply reflect a selection effect. File sharing is attractive to those who are time-rich but cash-poor, and these individuals would purchase fewer CDs even in the absence of P2P networks.

A handful of academic studies rely on micro data to address the issue of unobserved heterogeneity among file sharers.[2] Rob and Waldfogel (2006) study the survey responses of a convenience sample of U.S. college students. For hit albums which sold more than 2 million copies since 1999, they find no relationship between downloading and sales. Expanding the set of albums to include all music the students acquired in 2003, downloading five albums displaces

[2] The Journal of Law and Economics published additional papers in a symposium on file sharing in 2006. Oberholzer-Gee and Strumpf (2005) discusses these studies and additional work.

the sale of one CD. This difference is interesting. One interpretation is that piracy does not affect hit albums but hurts smaller artists. It is also possible that file sharing had less of an effect on sales in earlier years. After instrumenting for downloads with the school the students attend – everyone at Penn has broadband access while this is not true for the other schools – the resulting estimates are too imprecise to draw any firm conclusions. Zentner (2006) employs European survey data to study the relation between file sharing and sales. Using measures of Internet sophistication and access to broadband as instruments, Zentner finds some displacement. Unfortunately, neither the Rob and Waldfogel study nor Zentner's work allow inferences about the total impact of file sharing on record sales because neither paper studies a representative sample of file sharers. Zentner also lacks information about the number of downloads and CD purchases. Our approach differs from the current literature in that we directly observe file sharing. Our results are based on a large and representative sample of downloads, and individuals are generally unaware that their actions are being recorded.

III. File sharing Networks

File sharing relies on computers forming networks which allow the transfer of data. Each computer (or node) may agree to share some files and has the ability to search for and download files from other computers in the network. Individual nodes are referred to as clients if they request information, servers if they fulfill requests, and peers if they do both. Our data come from the OpenNap network, an open-source descendant of Napster. OpenNap is an example of a centralized P2P network which has individual clients log into a central server.

During our study period in the fall of 2002, P2P networks were already quite large. FastTrack (which includes the popular KaZaA service) had grown to 3.5 million simultaneous users by December 2002. The second largest network was WinMX, which had about 1.5 million simultaneous users in 2002. Even the smaller networks were fairly large. OpenNap had at least 25,000 simultaneous users sharing over 10 million files. Napster no longer operated in the fall of 2002.

IV. Data

We use two main data sources for this study. Logs for two OpenNap servers allow us to observe what files users download. Weekly album-level sales data come from Nielsen SoundScan (2005). SoundScan tracks music purchases at over 14,000 retail, mass merchant and online stores in the United States. Nielsen SoundScan data are the source for the well-known Billboard music charts. To develop our instruments, we rely on a large number of additional data sources which we discuss in the next section.

A. File Sharing Data

Our data were collected from two OpenNap servers, which operated continuously for seventeen weeks from 8 September to 31 December 2002. The servers were connected to T-3 lines which provided actual Internet transmission speeds of several megabits per second. The information on file transfers is collected as part of the log files which the servers generate. An excerpt of a typical log file is:

```
[2:53:35 PM]: User evnormski "(XNap 2.2-pre3, 80.225.XX.XX)" logged in
[2:55:31 PM]: Search: evnormski "(XNap 2.2-pre3)": FILENAME CONTAINS "kid rock devil"
    MAX_RESULTS 200 BITRATE "EQUAL TO" "192" SIZE "EQUAL TO" "4600602" "(3 results)"
[3:02:15 PM]: Transfer: "C:\Program Files\KaZaA\My Shared Folder\Kid Rock –Devil
    Without A Cause.mp3" (evnormski from bobo-joe)
```

The last two lines in the log file show user "evnormski" downloading the song "Devil Without a Cause" by Kid Rock from user "bobo-joe". Information on downloads are the building blocks of our analysis. We focus on downloads because these are the files users actually obtain and they can potentially displace sales. Over the sample period we observe 1.75 million file downloads.[3] We restrict the analysis to audio files by clients in the U.S. The server logs include the I.P. address for each client which we use to identify our users' home country.

An important question is whether our sample is representative of data on all P2P networks.[4] While we are unaware of any database spanning the universe of music downloads, we were able to compare the data from our servers with a sample of more than 25,000 downloads from

[3] At the end of 2003, roughly one billion songs are downloaded per week (*Wall Street Journal*, 19 November 2003). During February 2001, at Napster's peak, about half a billion songs were downloaded per week (Romer, 2002).
[4] At the time of our study, the OpenNap network represented about one percent of the total number of P2P users. A more comprehensive discussion of this point is in Appendix A of Oberholzer-Gee and Strumpf (2005).

FastTrack/KaZaA, the leading network at the time. We find that the availability of titles is highly correlated on the two networks. Using a standard homogeneity test based on 1,789 unique tracks, we cannot reject a null that the two download samples are drawn from the same population (Pearson χ^2 statistic is 1824.1). The resemblance of files is not surprising. Individuals in our data are similar to those on the most popular networks because the user experience is quite similar and many individuals employ software which allows them to simultaneously participate on several networks. For example, roughly one third of OpenNap participants uses the WinMX software, which simultaneously accesses WinMX and FastTrack/KaZaA, the two largest networks during our study period. We also find that users on these larger networks and those on our servers have access to a comparable number of files and that network size has little effect on the *distribution* of downloads. Based on these tests, we conclude that our sample is representative of the file transfers on the major P2P networks during our study period.

B. Sales Data and Album Sample

In this study, we focus on a *sample* of albums sold in U.S. stores in the second half of 2002. The sample is representative of all commercially relevant albums, allowing us to draw meaningful inferences about P2P's impact on overall music sales.[5] The sample is drawn from a *population* of albums on 11 charts produced by Nielsen SoundScan (2005): Alternative Albums (a chart with 50 positions), Hard Music Top Overall (100), Jazz Current (100), Latin Overall (50), R&B Current Albums (200), Rap Current Albums (100), Top Country Albums (75), Top Soundtracks (100), Top Current (200), New Artists (150), and Catalogue Albums (200). The charts are published on a weekly basis, and we include an album in the population if it appears on any chart in any week during the second half of 2002. The original population is extensive (2,282 albums) and includes many poorer-selling albums. For instance, our data include two albums which sold fewer than 100 copies during our study period, and the 25th percentile of sales in our data is only 12,493 copies.[6] While we study the commercially most relevant music, it would be incorrect to think of our population as a set of superstar albums. From this population, we draw a genre-

[5]The genre charts we sample from made up 81.8% of all CD sales in the United States in the last third of 2002. This is virtually identical to the 2002 share of 83.6% for the Big Five record companies, and 97% of the albums on the annual version of these charts were released on RIAA-associated labels.
[6]A typical measure of album success is gold certification which occurs at sales of half a million copies.

based, stratified random sample of 680 releases. To reflect the popularity of different music styles, we set the sample share of a genre equal to its fraction of CD sales in 2002.[7] Within each genre, we randomly select individual titles.

The average album in the resulting sample sold 143,096 copies during our study period. Table 1 reports sales statistics for the full sample and for individual categories. Across all categories, 44% of population sales are represented in the sample. A two-sample Kolmogorov-Smirnov test comparing the distribution of sales on the original charts and in our sample is unable to reject the null that sample sales are representative of the population of all albums (p=0.991). We also reject this null comparing each of our 11 original charts with the sample sales for that particular chart (p>0.539 for all 11 charts.)

In order to compare sales and downloads, we match the 260,889 audio files which U.S. users successfully transferred during our study period to the 10,271 songs on the 680 albums in our sample. The matching procedure is hierarchical in that we first parse each transfer line, identifying text strings that could be artist names. These text strings are then compared to the artist names in our set of albums. The list of artists contains the name on the cover and up to two other performing artists or producers that are associated with a particular song. For example, the track "Dog" on the B2K album "Pandemonium" is performed by Jhene featuring the rapping of Lil Fizz. For "Dog," B2K, Jhene and Lil Fizz are recognized as artists. Once an artist is identified, the program then matches strings of text to the set of songs associated with that particular artist. Using this algorithm, we match 47,709 downloads in the server log files to our list of songs, a matching rate of about 18%.

There are two reasons why this rate is less than 100%. A first is that an observed download is a transfer of a file that is not in our sample. These transfers are not of any concern, they simply reflect the fact that we are working with a sample. A second reason for a match rate of less than 100% could be that our matching algorithm fails to recognize songs. To investigate this possibility, we hand-checked a file with 2,000 randomly chosen unmatched transfers, comparing these downloads against our sample. Only five of the unmatched songs were in our sample. As

[7]Albums can appear on more than one chart because some charts (e.g., New Artists, Top Current) comprise many musical styles. For sampling purposes, we grouped all albums by style; a Rap album on the Top Current list is grouped with all other Rap albums during the sampling process. In the descriptive statistics, we classify albums by their original charts.

a result, we believe that the 18% match rate mostly reflects transfers of songs that are not in our sample.

C. *Descriptive Statistics*

As this is one of the few data sets that allow us to directly observe P2P users, we describe our data in some detail to convey a sense of what individuals do on P2P networks. A first stylized fact is that file sharing is truly global in nature. While over ninety percent of users are in developed countries, a total of 150 countries are represented in the data. U.S. users make up 31% of the sample. Table 2 shows the top countries for users and downloads. As the data indicate, there is only a loose correlation between user share and other country covariates such as Internet use or the software piracy rate. Column 3 in Table 2 confirms that interactions among file sharers transcend geography and language. U.S. users download only 45.1% of their files from other U.S. users, with the remainder coming from a diverse range of countries including Germany (16.5%), Canada (6.9%) and Italy (6.1%).

While file sharing activities are dispersed geographically, only a limited number of songs are transferred with any frequency. Table 3 shows the average song is downloaded 4.6 times over the study period, but the median number of downloads is zero.[8] Although our sample is representative of all commercially relevant music in the second half of 2002, it is striking to see that more than 60% of the songs in our sample are never downloaded. Aggregated up to the album level, users made 70 downloads from the average album in our sample. The most popular album among file sharers (and the second-best seller) has 1799 downloads, while the median number of downloads per album is 16, the 75th percentile is 63, the 90th percentile is 195, and the 95th percentile is 328. Both downloads and sales closely follow a power-law (pareto) distribution.

File sharing is limited to a select number of songs and most of these songs come from just a few charts. Tracks on the Top Current chart ("Billboard 200") are most frequently downloaded. Downloads from this chart alone make up 48% of all file transfers. Another 25% come from the "Alternative" category. The remaining 9 charts are not particularly popular among file sharers (see Table 3). In view of the low cost of sharing and sampling music on P2P, one could expect

[8] The 75th percentile of downloads per song is 2, the 90th percentile is 11, and the 95th percentile is 22.

users to seek out a great variety of songs representing many musical styles. But this is not the case. P2P downloads closely resemble the play lists of Top 40 radio stations. As a result, it is not surprising that songs from higher-selling albums are downloaded more frequently (Table 4). In the top quartile of sales, albums average 200 downloads. In the bottom category, the mean number of downloads is only 11. The fact that common factors drive downloads and sales is a key concern for the development of our empirical strategy.

V. Empirical Strategy

A. Econometrics

Our goal is to measure the effect of file sharing on sales. We observe sales and downloads for seventeen weeks which allows us to estimate a model with album fixed effects,

$$S_{it} = X_{it}\beta + \gamma D_{it} + \sum_{s}\omega_s t^s + \nu_i + \mu_{it}\,. \tag{1}$$

i indicates the album, S_{it} is observed sales, X_{it} is a vector of time-varying album characteristics that includes a measure of the title's popularity in the U.S., D_{it} is the number of downloads, t denotes time in weeks, and the summation allows for a flexible time effect. The key concern in our empirical work is that the number of downloads is likely to be correlated with unobservable album-level heterogeneity. As the descriptive statistics suggest, the popularity of an album is likely to drive both file sharing and sales, implying the parameter of interest γ will be estimated with a positive bias. The album fixed effects ν_i control for some aspects of popularity, but only imperfectly so because the popularity of many releases in our sample changes quite dramatically during the study period.

We address this issue by instrumenting for D_{it} in a 2SLS model. Valid instruments Z_{it} predict file sharing but are uncorrelated with the second-stage error μ_{it}. As in the differentiated products literature, where the problem is correlation between prices and unobserved product quality, we use cost shifters to break the link between unobserved popularity, downloads and sales (Berry, 1994; Bresnahan, et al. 1997). An advantage of our instruments, which we discuss below, is that they do not rely on the common but potentially problematic assumption that product characteristics are exogenous (Nevo, 2001). In the Appendix, we presents a formal model of purchase and download behavior which is the foundation for our econometric approach. In particular the Appendix shows why we can use linear demand equations rather than the more complicated transformations which are typical in this literature (Berry, 1994).

B. Instruments

Our most important instrument is the number of German secondary school kids who are on vacation in a given week. German users provide about one out of every six U.S. downloads, making Germany the most important foreign supplier of songs. School vacations produce an increase in the supply of files because German teens, who do most of their file sharing at home, can spend more time trading music when they are not in school, and they can stay up later during school vacations, allowing them to engage in file sharing during the peak U.S. trading hours (early evening, EST) (Niesyto, 2002). Supporting this intuition, we find that the number of German kids on vacation is a significant predictor of the number of files uploaded from Germany to the United States (p=0.011). The effect is particularly large for music genres that are popular in Germany.

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. The sixteen German Bundesländer (states) start their academic year at different points in time to smooth the demand for the German tourism industry and avoid traffic jams (Kultusministerkonferenz, 2002). For example, Bavarian students were still on summer vacation during the first week of our study period while Rheinland-Pfälzer kids were already back in school (see Figure 1). A second difference to a typical U.S. vacation schedule is that many, but not all Bundesländer grant their students one or two weeks of fall vacation. In Rheinland-Pfalz, this happened in weeks 4 and 5. Bavaria, in contrast, did not schedule a longer fall recess. These länder-specific holidays move from year to year. A Bundesland with early summer vacations in one year is given a later slot in the following year (Agentur Lindner, 2004). As we explain in greater detail below, there is good reason to believe this variable is exogenous.

If file sharing were eliminated tomorrow, German school holidays would have no relation to U.S. record sales.

We create three additional instruments by interacting the German-kids-on-vacation variable with album-specific characteristics. These instruments are particularly useful because they vary across both time and albums and provide identification even if a full set of week and album fixed effects is included.

German-kids-on-vacation × band is on tour in Germany: Tours spur local interest and sales of an album, and they are likely to create a positive supply shock of downloadable files. This instrument is not directly related to U.S. sales because the promotional effect of tours will not spill across the Atlantic and because the timing of fall and winter concerts in Germany typically reflects idiosyncratic features like venue availability and weather. We expect the effect of German vacations to be even larger if an artist happens to be on tour in Germany that week.

German-kids-on-vacation × indicator for misspellings in song titles: To download a song, a user's search query must match a shared file. At the time of our study, file sharing programs were rather rigid in determining matches.[10] Unless both the searcher and sharer agree on the naming convention, no match will occur. This two-sided search problem suggests that songs with unconventionally spelled titles may be more difficult to find. We use MS Word's spell checker to determine if an album has any songs with unconventional spelling. We expect misspellings to reduce the size of the positive supply shock coming from German vacations.

German-kids-on-vacation × rank of album on German charts: Albums that are popular in Germany are easier to download because the supply of these files is larger. Our measure for German popularity is the rank of the album on the weekly German Top 100 chart (Musikmarkt, 2002). Obviously, there is a concern that these chart positions might also measure U.S. popularity. However, the instrument is included along with album fixed effects, so it is the timing of the chart rankings in Germany that identifies downloads. There are important

[10]For example, "lose yourself," the name of a popular song, would typically return over a thousand results, but mistyping even one character (such as "lose yourse;f") or omitting part of a word ("lose yours") returned zero results.

differences in the dynamics of song popularity in the two countries due to taste differences and differences in release dates.

For all our instruments, we provide additional evidence for their exogeneity in the following sections. Summary statistics for the instruments are in Table 5. Each measure exhibits noticeable variation.

C. Mechanisms Underlying the Main Instruments

A key presumption in our analyses is that each instrument provides an exogenous shift in the cost of downloading files. We test this idea by analyzing how long it takes to download a file. For a subset of observations (weeks 3 through 6), we have more detailed server log files which allow us to calculate the actual difficulty of obtaining files. We construct five measures: the time between a download request and the successful initiation of the download (C_1), the time between a search request and a download request (C_2), the time between the initiation of the download and its successful completion (C_3), the ratio of search requests to the number of successful downloads (C_4), and the percentage of failed or canceled download requests (C_5). Each C_i term captures aspects of delay or frustration which a U.S. downloader might experience. The measures are aggregated up to the album-week. For example, C_1 is the average time until download initiation among all observed requests for that album in a particular week.

Mean C_i values are presented in the last row of Table 6. The first three columns show that the typical file takes twenty minutes to download, starting from the initial file search until the transfer is complete.[11] The long transfer times suggest there is an ubiquitous scarcity of supply, even of popular albums. While slow download speeds are the norm in our data, the estimates in Table 6 show that searching and downloading audio files in the U.S. is considerably easier when a larger number of German school children are on vacation. This reduction is even larger when the artist is on tour and when the album is highly ranked on the German charts.[12] The misspellings interaction significantly increases the time between a search and a download request as well as the number of unfulfilled downloads (C_2, C_4, C_5), but it has little effect on the time it

[11] Gummadi et al, 2003 independently document these long download times.

[12] Note that the German tour and singles chart variable parameters are identified using only within album variation since fixed effects are included. This mitigates concerns that album popularity in the U.S. is driving the parameter estimates.

takes to transfer a file (C_1, C_3). This is consistent with the argument that misspellings create confusion, though they do not slow down the file transfer itself. The estimated effects on download times are economically significant. For example, a one standard deviation increase in the vacation variable implies a 1.25 minute reduction in the time for a download to begin (C_1), which is an eighth of the typical delay..

These results are meaningful only if the cost of downloading influences the number of file transfers. This is not obviously true because P2P users can engage in other activities while files are being downloaded, which could mean they are insensitive to the time cost of file sharing. To check if the variation in download time that is due to our instruments has a significant impact on the number of transfers, we estimate the system

$$
\begin{aligned}
C_{it} &= Z_{it}\delta + v_i + \mu_{it} \\
D_{it} &= C_{it} + v_i + \varepsilon_{it}
\end{aligned} \quad , \tag{2}
$$

where Z_{it} is the full list of instruments and C_{it} denotes total download time (C_1+C_2+C_3). The sixth column of Table 6 shows that P2P users are fairly sensitive to the time cost of file sharing: a one standard deviation increase in download time reduces downloads by almost half of their mean. We find similar effects when we separately estimate (2) for each of the five C_i terms. These estimates confirm our initial claims. German vacations influence the cost of downloading, and this effect has an important impact on the number of downloads in the U.S.[13]

D. Specific Concerns

In this section, we discuss more specific concerns related to individual instruments.

German-kids-on-vacation: A potential difficulty with the vacation variable is that it might be correlated with time-varying album popularity in the U.S. We perform a number of tests to see if this is the case. First, we check if German vacations happen to coincide with official U.S.

[13] A different approach to show that German vacations influence downloading activity is to look at international data. We find that school holidays have an important effect only in countries whose time zones are complementary to Germany's.

holidays. We find that there is little overlap.[14] A second possibility is that German school vacations proxy for American vacations which are likely to have a direct impact on music sales. As there is no centralized data on holidays for all 14,000 U.S. school districts, we collect information on the number college students who are out of school during our study period. The sample includes all schools in the top two tiers of U.S. News and World Report's 2002 ranking. Information on school breaks is available for 157 schools, leaving us with data for 2.17 million students, almost a quarter of all U.S. college students. Figure 1 compares the vacation patterns in Germany and the U.S. There are marked differences. When some German kids are off in early fall, U.S. students are mostly in school. During the Thanksgiving break in the U.S., German kids are in school. Both populations are off during the Christmas break, although the break starts earlier for U.S. students. To test more formally if the number of German kids on vacation proxies for the number of U.S. kids, we include the latter in the first stage of equation (1). We find no evidence that the measured effect of German vacations on American music downloads is mediated by U.S. vacations.[15]

In a final test, we check more directly if the German vacation variable is in fact uncorrelated with U.S. demand for music albums. We do this by interacting the instrument with an album's rank on the U.S. MTV charts.[16] MTV rankings have the advantage that videos are often shown prior to the release of a CD, at a time when songs from a forthcoming album first appear on file-sharing networks.[17] This interaction is included in both stages of equation (1).

$$\begin{aligned} D_{it} &= X_{it}\beta + Z_{it}\delta + \varphi_1 Gkids_t \times MTV_{it} + \sum_s \omega_s t^s + \nu_i + \varepsilon_{it} \\ S_{it} &= X_{it}\beta + \gamma\, \hat{D}_{it} + \varphi_2 Gkids_t \times MTV_{it} + \sum_s \omega_s t^s + \nu_i + \mu_{it} \end{aligned}, \qquad (3)$$

where Z_{it} is our full set of instruments. As required under our assumptions, φ_1 is positive: German vacations have a larger effect for files that are more popular in the U.S. In the second stage, however, φ_2 is economically small and statistically insignificant. When an album becomes

[14] Estimates over our 17 week observation period yield: US Holidays$_t$ = 1.148 (1.61) - 0.182 (0.16) ×German Kids, where US Holidays is the number of official American holidays (such as Columbus Day or Thanksgiving) in week t and German Kids is the German holiday instrument.

[15] Controlling for the entire set of instruments, the estimated effect of German vacations on downloads changes from 0.667 (0.054) without the U.S. students-on-break variable to 0.643 (0.057) with this variable.

[16] We thank one of our referees for this suggestion.

[17] We also used the Billboard Airplay ranking to explore these effects, with similar results.

more popular in the U.S., this boost in popularity is not directly related to German vacations, supporting our claim that the holiday shocks are exogenous.

A second concern is that Germans supply only a narrow slice of music that is of interest to U.S. file sharers. If those who like the type of music that Germans make available substitute downloads for purchases in an atypical fashion, we measure a local average treatment effect, not a true population effect (Imbens and Angrist, 1994). Fortunately, there is substantial overlap between American and German musical tastes. Of the albums that entered our sample via the Billboard 200, 62.65% are also on the top 100 German charts. More generally, we study Amazon rankings to compare sales ranks in the two countries (Goolsbee and Chevalier, 2003). With the exception of Latin and Country music, Wilcoxon matched-pairs signed-ranks tests cannot reject the null of equal distributions for the eleven genres in our sample. In the robustness section of the paper, we test if the undersupply of Latin and Country music affects our estimates. We show that this is not the case, suggesting the measured effect of downloads on sales is likely to be a good estimate of the average population effect.

German-kids-on-vacation × indicator for misspellings in song titles: Because misspellings appear to be more likely in some genres than in others, one might argue that this indicator is likely to proxy for album popularity. In our application, this concern is not valid for two reasons. First, as an empirical matter, we find that misspellings are not correlated with sales, even in models without album or genre fixed effects.[18] Second, all our specifications presented in the results section include album fixed effects which control for an album's time-invariant popularity.

A second difficulty with the misspelling instrument could be that misspellings cause our song matching algorithm to fail. This would result in a negative relationship between misspellings and measured downloads, even if misspellings had no effect on actual downloads. More importantly, the second-stage estimates would be attenuated towards zero, since the variation in fitted downloads would be largely due to noise. Several pieces of evidence suggest this is not true. First, the estimates in the last sub-section show that misspellings do in fact have real effects

[18] The effect of misspellings on sales is statistically insignificant and economically small. A one-standard-deviation increase in misspellings raises sales by a mere 11,000 copies during our entire study period. This increase is about 3% of a standard deviation in sales.

on transfer times and user behavior. Second, we can check for misspellings in unmatched downloads. If the criticism is correct, there should be more misspellings in the unmatched than in the matched sample. This is not the case.[19]

German-kids-on-vacation × rank of album on German charts: The idea underlying this instrument is that vacation periods in Germany will boost downloads in the U.S. more when many German users make a particular file available. Because the instrument is included along with album fixed effects, it is the timing of the chart rankings in Germany that identify downloads. However, if U.S. popularity shocks happen to coincide with high German chart positions, we would measure the effect of downloads on sales with a positive bias. We can test for this spurious correlation in two ways. First, assuming that the German vacation variable is a valid instrument, we can perform overidentification tests for this and the other interactions that we use as instruments. These tests, reported in the results section of the paper, provide no indication that any of our instruments are invalid. A second and more direct test is to see whether shocks in U.S. demand are correlated with German popularity.[20] Under our hypotheses, U.S. demand shocks must not get magnified when albums become more popular in Germany. For example, we expect U.S. vacations to increase P2P activity, but this increase must not vary with German popularity. The model is

$$\begin{aligned} D_{it} &= Z_{it}\delta + \varphi_1 Ukids_t + \varphi_2 Ukids_t \times Gcharts_{it} + \varphi_3 Ukids_t \times MTV_{it} + \varphi_4 Gkids_t \times MTV_{it} \\ &\quad + \sum_s \omega_s t^s + \nu_i + \varepsilon_{it} \end{aligned} \tag{4}$$

$Ukids_t$ denotes the number of U.S. college students on break (our measure of U.S. demand shocks), $Gcharts_{it}$ is a title's rank on the German charts, and MTV_{it} is the position on the MTV chart (our measure of U.S. popularity). The effect of interest in this specification, φ_2, shows whether a shock in demand in the U.S. is mediated by German popularity. This is not the case: φ_2 is -0.0008 with a standard error of 0.0134, and this effect is only one tenth of the size of the German kids × German chart interaction in our later specifications. The data show that relative popularity in Germany interacts with German but not with U.S. vacations.

[19] The rates are 0.041 (N=35614) and 0.038 (N=7163), respectively. The Pearson $\chi 2$ statistic is 1.402.
[20] We thank one of our referees for this suggestion.

VI. Results

Before turning to the estimates, it is instructive to graph some of the data.. Figure 2 shows the weekly time series of sales and purchases for two albums. The "Superstar" album was largely ignored in file sharing networks until it became available for sale in week ten of our sample. This suggests it is the publicity associated with an official release which drives downloads as well as sales. Notice also the rapid but non-monotone decay in sales and downloads, which highlights the importance of using high-frequency data.

A. Panel Analysis

In Table 7 we report results for model (1). The unit of observation is the album-week. The models include a control for time-varying U.S. popularity, the album's position on the American MTV charts, and a polynomial time trend of degree six. As expected, a simple OLS specification yields a large positive effect of 1.093 with a standard error of 0.023. A model which adds album fixed effects is given in column (I). While we continue to find a positive effect of downloads on sales, the relationship is now much weaker. The remaining estimates in Table 7 instrument for downloads. We begin by using the number of German kids on school vacation (column II). The first-stage estimates imply that a one standard deviation increase in the number of children on vacation boosts weekly album downloads by slightly less than one fifth of a standard deviation in downloads, an effect that is statistically significant and economically meaningful. Once we instrument for downloads, the estimated effect of file sharing on sales is small and statistically indistinguishable from zero.

We next consider specifications in which we add the band-on-tour-in-Germany interaction and the remaining time-varying instruments (columns III and IV). The tour and the German-chart interactions are of particular interest since they vary across albums as well as over time and provide an additional source of identification. The instruments have the expected first-stage signs. Tours and better chart positions magnify the effect of German students on vacation. The reverse is true for misspellings, which make it more difficult to search for files. Sargan overidentification tests are reported at the bottom of the table. In these richer models downloads continue to have economically small and statistically insignificant effects on sales.

To help improve the precision of our second-stage estimates, in column (V), we allow the effect of the German vacation instrument to vary by album. The logic for including these interactions follows from the same arguments used for the other instruments. When German kids spend more time on P2P networks, the resulting supply shock will vary across albums because the students supply the files that happen to be popular in Germany at the time of the shock. As before, we face a potential problem with using this type of variation: If it so happens that the exogenous German shock is spuriously correlated with album-specific surges in popularity in the U.S., our estimates would be biased. The specification in column (V) addresses this issue in four ways. As before, we include album fixed effects to make sure it is the timing of the supply shocks that identify downloads. Second, we introduce album-specific U.S. popularity effects at both stages of the model by interacting the MTV variable with the album fixed effects. The model thus controls for changes in the U.S. popularity of a release. Third, relying on the assumption that the number of German kids on vacation is a valid instrument, we conduct overidentification tests in a specification that includes only two instruments: the vacation variable and one of the vacation × album-fixed-effect interactions. There are 680 such tests. To err on the side of caution, we exclude from the final specification all interactions whose overidentification tests cannot reject the null at a significance level of greater than 0.20. There are 21 such interactions. Fourth, we estimate a variant of model (3), now with German kids × album fixed effect × U.S. MTV interactions. In the sales equation, these interactions are individually and collectively not different from zero.

Column (V) in Table 7 reports these results. Our instruments retain their statistical significance.[21] The mean of the coefficients on the vacation-album-fixed-effect interactions is -1.143, leaving the average effect of vacations on downloads almost unchanged from the earlier specifications. Model (V) shows that the German supply shock is largest for the type of music that is popular in Germany. For example, the mean of the coefficients on the vacation-album-fixed-effect interactions is -0.91 for Rock. Not surprisingly, German vacations have the largest

[21] The vacations × misspellings interaction is collinear with the vacations × album fixed effects and cannot be included in this specification.

effect on the number of international songs that are downloaded in the U.S. (mean is -0.71). In contrast, the effect of vacations is much smaller, but still positive, for genres that are less popular in Germany (-1.52 for Latin music, -1.54 for Country, and -1.57 for holiday music.) At the second stage, the estimated effect of downloads on sales is virtually unchanged in this specification, but the standard error drops considerably.

To see if our results are driven by our modeling choice for the time trend in downloads and sales, we replace the polynomial with week fixed effects in columns (VI) and (VII) of Table 7. In these specifications, we lose the German-kids-on-vacation instrument because it does not vary across releases. The results remain similar, with more precise second-stage estimates when we allow the effect of vacations to vary by release (column VII).

Table 7 suggests file sharing had a surprisingly small effect on sales that is statistically indistinguishable from zero. The instrumented point estimates fall within a very narrow range and suggest that file sharing did not heavily impact the music industry as a whole. If file sharing were to be eliminated, the most negative estimate (column VI) implies industry sales for all of 2002 would increase by 6.5 million albums. Using the most positive estimate (column VII) industry sales would fall by 8.9 million copies.[22] In 2002, the industry sold 803 million CDs. The robustness of these results extends to specifications not reported in Table 7. For example, we arrive at the same conclusions if we omit the misspelling or the German rank instrument.

B. Dynamic Analysis

The models in Table 7 only allow for a contemporaneous effect of downloads on sales, but it is quite possible that downloads influence sales at a later point in time. For example, users might sample music which they consider buying in the future. In Table 8, we address this issue by studying the effect of several weeks of downloads on sales and by estimating Generalized Methods of Moments (GMM) models. Downloads are highly correlated across time which prevents us from including downloads in past weeks as individual covariates. Instead, we study the effect of a weighted sum of current and past downloads on current sales. Downloads are

[22]The impact is the difference between predicted sales and the fitted value when downloads are set at zero. Using equation (1), the summed impact for our album sample and for our 17 week observation period is $\sum_t\sum_i S_{it}(D_{it})-S_{it}(0) = \gamma\times\sum_t\sum_i D_{it} \equiv \gamma\times 240\text{m}$. We multiply this number by a scaling factor to get the annual impact for the entire music industry (this is described below Table 11).

instrumented using the core set of instruments (specification IV in Table 7) or the extended set (specification V). Our formal measure is the weighted stock of current and previous weekly downloads, $D_t^{Stock} = \sum_{s\geq 0}\delta_s \times D_{t-s}$. The weights δ_s are chosen in a grid search that minimizes the unexplained fraction of the variance in our sales equation subject to $\delta_s \geq \delta_{s+1}$. The optimal weights $(\delta_0,\ldots, \delta_T)$ are (1,0.1,0.1). It is interesting that the weights which best fit our data give much importance to downloads in the current week, while downloads further back in the past do not heavily influence sales.[23]

In these models, we continue to find small and statistically insignificant effects for the weighted sum of three weeks of downloads, both in specifications with a polynomial time trend (Table 8, I&II) and with week fixed effects (III&IV). As in the panel results, standard errors drop significantly with the extended set of instruments (II&IV). We also constructed stock variables for the sum of downloads during the past four and six weeks and found no evidence of a sales crowd-out in these models.

Models (V) and (VI) in Table 8 use the GMM estimator developed by Arellano and Bond (1991). The GMM models are more general than the previous specifications in the sense that we do not need to make any assumptions about the appropriate lag structure. The lag of sales that is included on the right-hand side accounts for any effect that past downloads might have had on current sales. The model is estimated in first differences. We instrument for past sales using suitable lags of their own levels and our core set of first-differenced instruments.[24] For this type of model, the two-step estimates of the standard errors tend to be downward biased (Blundell and Bond, 1998). We correct standard errors using the two-step covariance matrix derived by Windmeijer (2000). Arellano-Bond tests for autocorrelation are applied to the first-difference equation residuals. Second-order autocorrelation would indicate that some lags of the dependent variable which are used as instruments are endogenous, but the tests reveal no such problem.

---

[23] Our working paper presents additional results showing that file sharers are impatient. These findings are consistent with those of Einav (2004) for movie consumption.

[24] The formal model is,

$$S_{it} = \alpha S_{i,t-1} + X_{it}\beta + \gamma D_{it} + \sum_s \omega_s t_s + \nu_i + \mu_{it}.$$

The lagged sales term soaks up any delayed effect of downloads, regardless of how far in the past they occurred (taking a Koyck transformation yields a specification with infinite lags of downloads on the right hand side). Estimating in first differences purges the fixed effects. We instrument for the first-differenced $S_{i,t-1}$ which are now endogenous.

The results of these models, with a polynomial time trend as in (V) or with week fixed effects as in (VI), are similar to our previous findings. The estimates are fairly precise, making these GMM models an alternative to using our extended set of instruments.

C. "Drop-out" Hypothesis

A possible explanation for our inability to find a statistically significant relationship between file sharing and sales is that file sharers and consumers who purchase music are in fact two separate groups. According to this hypothesis, growth in file sharing does displace sales but we cannot identify this effect because our data do not reflect the increasing number of file sharers.

There are three responses to this conjecture. First, it is inconsistent with what we know about consumer behavior. The premise underlying the "drop-out" hypothesis is that file sharers no longer buy CDs. However, every survey we are aware of, including the industry studies listed in the literature section, indicates that downloaders, even heavy ones, continue to purchase legal CDs. We corroborated these findings with our own survey of individuals who were engaged in file sharing (Oberholzer-Gee and Strumpf, 2004). Ninety percent reported that they recently purchased a CD, a value reaching one hundred percent among the most active downloaders.

Secondly, we can test the "drop-out" hypothesis directly by controlling for the increasing number of users. An implication of the hypothesis is that our download sampling rate declines over time because the servers for which we have data handle a limited number of users. Growth in file sharing, however, is managed by additional server capacity which we do not observe. If we accounted for this growth, the hypothesis suggests, we would find a displacement effect because the "drop-outs" are replacing purchases with transfers. We address this issue by scaling up the number of downloads in our sample to reflect the growth in file sharing. We use the number of FastTrack/KaZaA users as a proxy for the rate of growth.[25] Because the number of users increased by over a third over our observation period, we should be able to detect a drop-out effect if it exists. Table 9 reports these estimates for three panel models, three models using a stock of previous downloads, and for two GMM models. In all these specifications, downloads still do not have a significant effect on sales. A third approach to testing the drop-out hypothesis

[25] We use 22 data points on the number of KaZaA users in the period from 9/9/2002 to 2/4/2003 to fit a fractional polynomial trend in the number of users. The model explains 85% of the variation.

is to compare the long-run sales growth of individual genres of music. We return to this point in Section VII.

D. Robustness Tests

To further corroborate our results, we perform a large number of robustness checks, some of which we report in Table 10.[26] The tests fall in three broad categories: models for subsets of our sample, alternative econometric specifications, and models that allow the effect of file sharing on sales to vary by popularity. We first investigate the importance of the holiday season when many consumers purchase CDs as gifts. It is possible that downloads are less substitutable for sales during this period due to the reluctance to give downloaded music as a present. Note that this is also an argument against the idea that file sharing is the main cause of the sales decline, since purchases are heavily concentrated in the holiday season. Still, it is straightforward to test for this effect. In Table 10, we exclude the December data from our sample. We report these results for specifications IV, VI and VII in Table 7. Even without the December data, there is no statistically significant effect of file sharing on sales. In a second test, we omit albums that are not downloaded during our study period. These less popular releases might have little sales even in the absence of file sharing, making the effect of P2P on sales miniscule by definition. Omitting these albums, however, does not change our conclusions. The same holds if we restrict our sample to better-selling albums.

We next test if the undersupply of Latin and Country music influences our estimates. Recall that differences in the popularity of musical genres is only an issue if the substitutability of downloads and CD purchases varies across groups of listeners. The last specification in the first panel of Table 10 re-estimates our models without Latin or Country releases. As expected, this increases the effect of vacations on downloads, from a coefficient estimate of 0.667 in model IV of Table 7 to 0.744 in this model. However, the measured effect of downloads on sales remains similar, a finding that is consistent with the idea that the substitutability of downloads and purchases is roughly similar across genres.

---

[26]We thank our referees for suggesting several of these points. Many additional robustness tests can be found in Oberholzer-Gee and Strumpf (2005). This working paper also presents pooled specifications utilizing only cross-album variation, and these estimates also show file sharing has little impact on sales.

In the second panel in Table 10, we explore two alternative specifications. To reduce the importance of outlier albums with a large number of sales, we use log(sales) as the dependent variable. The impact on sales continues to be insignificant in all three specifications. In the next model, we first-difference both sales and downloads and express them as percentage changes. An advantage of this model is that it nicely captures album-specific trends in popularity. Unfortunately, this advantage comes at the cost of a reduced number of observations due to the first-differencing and the weeks with zero downloads or sales. Using our core set of instruments, we now find a statistically significant but economically small effect of downloads on sales. However, the estimated coefficient drops considerably when we introduce week fixed effects.

The previous models constrained the effect of downloads on sales to be identical for all releases. In the bottom panel of Table 10, we relax this assumption. We first explore the idea that the effect varies by artist popularity. We do this by interacting the download variable with two measures of popularity: an artist's last and his best-ever Billboard ranking. The rankings themselves are subsumed in the album fixed effects, but the interaction term varies by week. To make it easier to interpret the results, Billboard ranks are coded as [201 − actual rank] so that larger numbers indicate greater popularity.[27] We estimate these models using specification IV in Table 7. There is no indication that more popular artists are affected differentially. Neither the interaction terms nor the joint effect of the main and interaction terms are statistically significant.

From a welfare point of view, it is particularly interesting to study variations in the effect of file sharing across younger and older artists because such differences might influence their decision to start and continue a career in music. Interacting downloads with the number of albums an artist produced, we find no significant differences across more or less experienced performers. Finally, we investigate whether the effect of downloads on sales varies with the number of popular songs on a CD. As documented earlier, most file sharers obtain just a few tracks from a CD. One might suspect that P2P is a fairly good substitute for CDs with only one or two popular songs. We calculate a Herfindahl index for each album-week as a measure of concentration of downloads. The index is included in both the first and the second stage. There is no evidence that albums with more concentrated downloads suffer disproportionately from file sharing.

[27]More precisely, the term is a three-way interaction: [downloads × band had Billboard ranking × (201−Billboard rank)].

VII. Quasi-experimental Evidence

Our data also allow us to study the impact of P2P on sales in a quasi-experimental context. In particular we can examine how album sales respond to exogenous variation in file sharing intensity during certain times of the year, in particular geographic areas, across music genres, and from secular growth. One of the advantages of this approach is that we can utilize several years of data, which allows us to investigate the long-term impact of file sharing. In all cases we continue to use sales data from Nielsen SoundScan (2005).

The first experiment involves variation over time. The number of file sharing users in the U.S. drops twelve percent over the summer (estimated from BigChampagne, 2006) because college students are away from their high-speed campus Internet connections. If downloads crowd out sales, we should observe that the share of albums sold in the summer increases following the advent of file-sharing. We consider a differences-in-differences approach and compare the share of summer sales in the period prior to file sharing (the control group) with sales following the introduction of file sharing (the treatment group). We calculate the share of album sales occurring in the May to September period using weekly SoundScan data. We find that the introduction of widespread file-sharing has had virtually no impact on summer sales. In the four years (1995-1998) preceding the introduction of Napster, the average share of summer sales was 37.0% with a range of 36.4-37.8%. During the more recent period of extensive file-sharing (1999-2005), the average share of summer sales was 37.2% with a range of 35.9-37.8%.

A second experiment considers spatial variation. Recall that U.S. users download over a third of their music files from Western European countries such as Germany and Italy. Due to time zone differences, such transfers are easier for East rather than West Coast users. This is because the peak file-sharing period (7pm to 3am) overlaps between Western Europe and the East Coast, which have a six hour time difference, but not between Europe and the West Coast, which have a nine hour difference. So East Coast users can draw on a larger base of files from international users than West Coast users. Consistent with these differences, we find that there is more file sharing on the East Coast than on the West Coast.[28] If file sharing had a large negative effect on record sales, then sales during the file sharing era should decrease more on the East Coast than on the West Coast. For the period 1998-2002, we obtained total album sales for the one hundred one largest "Designated Market Areas" from SoundScan. Despite the differences in the availability of files, sales have not noticeably varied across the country. In 1998, the last year in the pre-P2P period, the share of album sales in the Eastern Time Zone was 43.9%. This share has hardly moved since then. In 1999-2002, the mean was 43.5% and the range was 42.7-44.0%. This suggests some common national factors, rather than file-sharing, are driving sales trends.

A third experiment, which also provides a test of the "drop-out" hypothesis, is to see whether download intensity influences long-run sales growth after explicitly controlling for trends in music format popularity. The model for the period 1999-2005 is,

$$\text{(5)} \qquad \text{Sales Growth}_g = \alpha + \gamma\times\text{Downloads}_g + \lambda\times\text{Listenership}_g + e_g$$

where g indicates genre, Sales Growth_g is the percentage growth in sales over 1999-2005, Downloads_g are measures of genre-specific download intensity from our data, and Listenership_g is the genre-specific radio listenership growth rate (Arbitron, 2006) which controls for trends in popularity. Since downloading is relatively concentrated across genres (Table 3), the "drop-out" hypothesis predicts a greater sales reduction for genres which are popular on file sharing networks. The estimated γ is not statistically significant using either download levels or downloads relative to purchases. For example, using mean downloads per album and controlling for genre sales levels, the estimated γ is 0.05 with a standard error of 0.52 (the mean for downloads is 61.2, and for sales growth it is -5.8).

Finally, we consider whether growth in file sharing can be linked to changes in total album sales. The key question is whether periods of particularly rapid growth in the user-base are linked to sharper sales reductions. A simple test is to consider annual sales since the advent of widespread file sharing in 1999. According to SoundScan, album sales increased in three of the seven years over this period, in contrast to movie ticket sales which rose in only two years. It is worth stressing that extended sales declines frequently occurred prior to file sharing. While real

[28] Unfortunately, IP addresses can only be matched imperfectly to locations, so this finding is merely suggestive.

revenues have fallen 28% over 1999-2005, real revenue fell 35% during the collapse of disco music in 1978-1983. Real sales also dropped 6% over 1994-1997.[29] More direct evidence comes from regressing total album sales, including paid digital downloads, on the average number of simultaneous file sharing users in the U.S. (BigChampagne, 2006),

$$\text{(6)} \qquad Sales_t = \gamma \times Users_t + \nu_m + \mu_t$$

where t indicates a month, and ν_m are monthly fixed effects which account for seasonality. Using monthly data from August 2002-May 2006 (N=46) and defining Sales and Users in millions (with respective sample means of 56.0m and 5.0m), the estimated γ=-0.427 with a robust standard error of 0.33. There is little evidence that growth in the number of users has had a statistically or economically significant effect on sales.[30] The estimates remain insignificant if equation (6) is estimated in first differences.

The results of these quasi experiments are consistent with our earlier findings. Looking at variation in downloading intensity that is due to geography, seasonality, the genre of music, or secular growth, we find no evidence that the advent of P2P technologies is the primary cause of the recent slump in music sales.

VIII. Conclusions

Using detailed records of transfers of digital music files, we find that file sharing has had no statistically significant effect on purchases of the average album in our sample. Even our most negative point estimate (Table 7, model VI), implies that a one standard deviation increase in file-sharing reduces an album's weekly sales by a mere 368 copies, an effect that is too small to be statistically distinguishable from zero. Because our sample was constructed to be representative of the population of commercially relevant albums, we can use our estimates to test hypotheses about the impact of P2P on the entire industry. Using ninety-five percent confidence bands, these tests are presented in Table 11. Taking into account all our (instrumented) estimates including the least precise results in Tables 7-9, we can reject a null that

[29] These are calculated from nominal RIAA revenues listed in Lesk (2003) and RIAA (1998; 2006).

[30] If file sharing were eliminated, the point estimates imply monthly sales would only increase by 2.1m.

P2P caused a sales decline greater than 24.1 million CDs. For reference, the music industry sold 803m CDs in 2002, which was a loss of 80m from the previous year (RIAA, 2004). Our estimates become more precise if we relax the assumption that file sharing exclusively affects this week's sales via this week's downloads and if we allow for growth in the number of file sharers. For example, the scaled models in Table 9 reject a null of losses greater than 12.4 million. Relying on our five most precise estimates, we conclude that the impact could not have been larger than 8.9 million CDs. While file sharers downloaded billions of files in 2002, the consequences for the industry amounted to no more than 1.1% of sales.

If file sharing is not the culprit, what other factors can explain the decline in music sales? Several plausible candidates exist. A first reason is the change in how music is distributed. Between 1999 and 2003, a fifth of music sales shifted from record stores to more efficient discount retailers such as Wal-Mart. About half of the RIAA's reported decline in CD shipments can be linked to the resulting reduction in store inventories. A second factor is the ending of a period of atypically high sales, when consumers replaced older music formats with CDs. Perhaps more important than these developments is the growing competition from other forms of entertainment. A shift in entertainment spending towards recorded movies alone can largely explain the reduction in sales. The sales of DVDs and VHS tapes increased by over $5 billion between 1999 and 2003. This figure more than offsets the $2.6 billion reduction in album sales since 1999. Consumers also spent more on video games, where spending increased by 40%, or $3 billion, between 1999 and 2003, and on cell phones. Teen cell phone use alone tripled between 1999 and 2003.

An interesting question is whether our results continue to hold in more recent years. Since the time of our study, P2P technology has become more efficient, broadband access is much more widespread, and the number of file sharers has doubled. While a full analysis is outside the scope of this paper, there are several trends that are inconsistent with the view that P2P now displaces sales on a large scale. First, our natural experiments, for which we have data up to 2005, give no indication that file sharing has caused a sales decline in more recent years. Second, music sales have been flat or even rising in major markets with a quickly growing file-sharing population. For example, music purchases in Germany were unchanged in 2005, and sales in France rose by 7.5% in the same year. Third, in the United States, where sales continued to decline from 681 million in 2004 to 654 million in 2005, the entire drop is due to losses at a single firm, the recently merged Sony-BMG, which has experienced severe post-merger integration difficulties. If file sharing were responsible for the observed sales decline in the U.S., we would not expect this activity to only affect the products of a single firm.

The advent of the new P2P technologies can be considered in a broader context. A key question is how social welfare changes with weaker property rights for information goods. To make such a calculation, we would need to know how the production of music responds to the presence of file sharing. Based on our results, we do not believe file sharing had a significant effect on the supply of recorded music. For artists who produce commercially relevant products, the effects documented in this study are simply too small to change the number or quality of recordings that they release. And for new bands that are about to launch their career, the probability of success is so low as to make the expected income from producing music virtually zero, so file sharing will not change the relevant incentives. If we are correct in arguing that downloading has had little effect on the incentives to produce music, we agree with Rob and Waldfogel (2006) who find that file sharing likely increased aggregate welfare. The limited shifts from sales to downloads are simply transfers between firms and consumers. But the sheer magnitude of P2P activity, the billions of tracks downloaded each year, suggests the added social welfare from file sharing is likely to be high.

Appendix A – Model

A. Setup

Consider a stylized model of downloading and purchase behavior. Suppose that each individual values music but faces some acquisition costs. There is population heterogeneity in these values and costs. Individuals first decide whether to download and then later whether to purchase.

In particular, let:

- $V_{ij} \geq 0$ be the value of purchased album $i = 1,...,N$ for individual $j \in \mathbb{R}^+$.
- $D_{ij} \equiv \gamma V_{ij}$ be the value of downloaded album i for individual j. Presumably $0 \leq \gamma \leq 1$ since downloads are inferior to the original album (lower sound quality, no liner notes, and perhaps remorse at not compensating the artist) though all that is needed is $\gamma \geq 0$.
- $p > 0$ be the cost of a purchased album (presumed to be constant since album prices rarely vary)
- $q_{ij} > 0$ be the monetized cost of downloading album i for individual j. This cost stems from time spent searching for and downloading the album. q_{ij} varies across individuals (due to different value of time or the speed of Internet connection) and albums (since some albums are longer and hence take more time to download).

Preferences are assumed to be separable over the goods. Given a single outside good which serves as the numeraire, after substituting the budget constraint the utility function of individual j is,

(A1) $U_j = \sum_i \mathbb{1}_{ij}(\text{purchase}) \cdot (V_{ij} - p) + \mathbb{1}_{ij}(\text{download}) \cdot (\gamma V_{ij} - q_{ij})$

where $\mathbb{1}_{ij}(.)$ is an indicator that the individual bought or downloaded album i.

Individuals face a sequence of discrete choices. First they must decide whether to download any of the albums, and then whether to purchase any of them (the discount factor is near unity since these decisions occur at nearly the same time). These are discrete choices in that each album can be downloaded or purchased once or not at all.

We presume the values of the albums and the costs of downloads are independent. The population density of values for album i is $V_i \sim f(V_i, \alpha_{Vi})$ and the population distribution is $F(V_i, \alpha_{Vi})$. The population density of costs for album i is $q_i \sim g(q_i, \alpha_{qi})$ and the population distribution is $G(q_i, \alpha_{qi})$. The $\alpha_.$ terms parameterize the distributions. α_{Vi} measures the <u>popularity of an album</u> which is viewed in terms of first order stochastic dominance: $F(V, \alpha_{VA}) \leq F(V, \alpha_{VB})$ (with a strict inequality for at least one V) when $\alpha_{VA} > \alpha_{VB}$. That is, albums with higher values of α_{Vi} are more valuable in aggregate or equivalently their population distribution is shifted to the right. α_{qi} measures the <u>cost of downloading an album</u> and is defined analogously: $G(q, \alpha_{qA}) \leq G(q, \alpha_{qB})$ (with a strict inequality for at least one q) when $\alpha_{qA} > \alpha_{qB}$.

B. Preliminary Result

To fix ideas, we first consider the case where preferences are independent across downloads and purchases. That is, we ignore the possibility of crowd-out or learning. From (A1) an individual purchases *iff* $V_{ij} > p$ and downloads *iff* $\gamma V_{ij} > q_{ij}$, and so aggregate values are, (A2) Total Purchases of album $i \equiv \int_{q>0}(1\text{-}F(p,\alpha_{Vi}))g(q,\alpha_{qi})dq = 1\text{-}F(p,\alpha_{Vi})$

(A3) Total Downloads of album $i \equiv \int_{q>0}(1\text{-}F(q/\gamma,\alpha_{Vi}))g(q,\alpha_{qi})dq$

These equations yield the first result.

> **Result 1.** *More popular albums have higher total downloads and total purchases, even if there is no feedback between purchases and downloads.*

Proof:
Consider album A and a less popular album B, $\alpha_{VA}>\alpha_{VB}$, which both have the same cost distribution, $\alpha_{qA}=\alpha_{qB}\equiv\alpha_q$. From (A2),

(A4) Purchases(A) – Purchases(B) = $F(p,\alpha_{VB})\text{-}F(p,\alpha_{VA}) \geq 0$

where the inequality follows from first order stochastic dominance. From (A3),

(A5) Downloads(A) – Downloads(B) = $\int_{q>0}(F(q/\gamma,\alpha_{VB})\text{-}F(q/\gamma,\alpha_{VA}))g(q,\alpha_q)dq > 0$

where the inequality again follows from first order stochastic dominance.

□

This highlights the problem with simply regressing downloads on purchases: both are endogenously determined by popularity, so OLS will yield a spurious positive relationship.

C. Main Model
More generally downloads should influence purchases (we continue to presume there is no spillover between albums). The effect of downloads is modeled as a shift in the α_{Vi}:

(A6) $\alpha'_{Vi} \equiv \alpha_{Vi}$ following a download = $\phi(\alpha_{Vi})$

where $\phi(.)$ is a weakly monotone increasing function, $\alpha_{VA}>(<)\alpha_{VB} \rightarrow \phi(\alpha_{VA})>(<)\phi(\alpha_{VB})$. (A6) allows downloads to increase or decrease the popularity of an album (and hence purchases), and for this effect to vary by the ex ante popularity: $\alpha'_{Vi}\geq\alpha_{Vi}$ or $\alpha'_{Vi}\leq\alpha_{Vi}$ and this relationship may vary with the level of α_{Vi}. The only restriction is that downloading does not change the ranking of album popularity, e.g. $\phi(.)$ is an order-preserving function.

A <u>modified definition of album popularity</u> is also used: when $\alpha_{VA}>\alpha_{VB}$, then we presume $f(V,\alpha_{VA})\geq f(V,\alpha_{VB})$ (with a strict inequality for at least one V) $\forall V\geq p$. That is, a more popular album (with a higher α_{Vi}) has a greater mass of individuals at every value which could lead to purchases. More popular albums have a thicker right tail in their density of values. This is typically a stronger condition on the density than stochastic dominance.

We presume individuals download myopically. That is, they do not take into account the potential for learning (the shift from α_{Vi} to α'_{Vi}) when making their downloading decision.

The positive correlation of purchases and downloads from Result 1 still holds in this more general framework. For example consider albums A and B with $\alpha_{VA}>\alpha_{VB}$ and $\alpha_{qA}=\alpha_{qB}\equiv\alpha_q$ The change in download equation (A5) in the proof of Result 1 is unaffected. The change in purchases equation is, (A7) Purchases(A) – Purchases(B) | Downloads have feedback

$$= \int_{V>p}((f(V,\phi(\alpha_{VA}))-f(V,\phi(\alpha_{VB})))G(\gamma V,\alpha_q)+(f(V,\alpha_{VA})-f(V,\alpha_{VB}))(1-G(\gamma V,\alpha_q)))dV > 0$$

where the first term is for individuals who download ($\gamma V_{ij}>q_{ij}$) and the second is for those who do not download ($\gamma V_{ij}<q_{ij}$). The inequality follows from the modified definition of popularity and the monotonicity of $\phi(.)$. Again the intuition is that album popularity drives both downloads and purchases.

The main objective of the paper is to understand the shape of $\phi(\alpha_{Vi})$, which shapes the effect of downloads on purchases. This cannot be measured from simply regressing downloads on purchases due to the positive correlation result. Instead it suggests using instruments, variables which shift downloads but have no direct effect on purchases. A natural instrument is the download costs parameter, α_{qi}.

Result 2. *Download costs influence purchases only though their effect on downloads. Download costs reduce album downloads.*

Proof:
Consider album A and a less costly to download album B, $\alpha_{qA}>\alpha_{qB}$, which both have the same popularity distribution, $\alpha_{VA}=\alpha_{VB}\equiv\alpha_V$. From (A3),

(A8) Downloads(A) – Downloads(B)

$$= \int_{q>0}(g(q,\alpha_{qB})-g(q,\alpha_{qA}))F(q/\gamma,\alpha_V)dq$$
$$= -\gamma^{-1}\int_{q>0}(G(q,\alpha_{qB})-G(q,\alpha_{qA}))f(q/\gamma,\alpha_V)dq < 0$$

where the second equality is from integration by parts and the inequality again follows from first order stochastic dominance. After separately integrating the downloading and non-downloading populations, the change in purchases equation is,

(A9) Purchases(A) – Purchases(B) | Downloads have feedback

$$= \int_{V>p}(G(\gamma V,\alpha_{qA})-G(\gamma V,\alpha_{qB}))(f(V,\phi(\alpha_V))-f(V,\alpha_V))dV$$

In the absence of feedback effects, $\phi(\alpha_V)=\alpha_V$, purchases are identical for the two albums (or simply see (A2)).

□

Asides:

- While the proof compares two albums, the equations can equivalently be interpreted as a comparison of the same album at two moments in time when its cost of downloading differ.
- After allowing for feedback, higher download costs increases (decreases) purchases *iff* downloading decreases (increases) album sales. That is, (A9) is positive *iff* $\phi(\alpha_V)<\alpha_V$ (this follows since costs are increased--so the first term in the integral is negative—and an application of the modified popularity definition—so the second term is negative when $\phi(\alpha_V)<\alpha_V$).

Result 2 shows download cost shifters are appropriate instruments. A cost drop increases downloads and (*iff* the feedback effect from downloads is positive) increases purchases. The opposite holds for a cost hike. With enough data we can ascertain the shape of $\phi(\alpha_V)$ for a wide range of popularity levels.

D. Functional Form for the Estimation Equation

A final issue is the appropriate functional form for the estimates. We argue that a linear equation relating aggregate sales to downloads is appropriate. To see this, we first write the expressions for downloads and purchases of some album,

(A10) $\text{Downloads} = \int_{V>0} f(V,\alpha_V)G(\gamma V,\alpha_q)dV$

and,

(A11) $\text{Purchases} = (1-F(p,\alpha_V)) + \int_{V>p}(f(V,\phi(\alpha_V))-f(V,\alpha_V))G(\gamma V,\alpha_q)dV$

These can be can be combined to give,

(A12) Purchases

$$= (1-F(p,\alpha_V))+\int_{V>p} f(V,\phi(\alpha_V))G(\gamma V,\alpha_q)dV+\int_{0>V>p} f(V,\alpha_V)G(\gamma V,\alpha_q)dV-\int_{V>0} f(V,\alpha_V)G(\gamma V,\alpha_q)dV$$

$$\equiv \text{Purchases}_{\text{NoDownloads}}(p,\alpha_V) + \Psi(p,\gamma,\alpha_V,\phi(\alpha_V),\alpha_q) - \text{Downloads}(\gamma,\alpha_V,\alpha_q)$$

The first term on the bottom row measures total purchases in the absence of downloads, and is independent of the download cost parameter α_q. The remaining two terms reflect the effect of downloads. (A12) shows that it is roughly appropriate to use a linear specification in the estimates. It also highlights our instrument strategy. An exogenous shift in the distribution of download costs, as measured by α_q, influences downloads and, recalling the discussion after Result 2, will increase or decrease purchases based on the shape of $\phi(\alpha_V)$.

References

Agentur Lindner (2004). `http://www.agentur-lindner.de/special/schulferien/index.html`.

Arbitron (2006). *Format Trends Report*. http://wargod.arbitron.com/scripts/ndb/fmttrends2.asp.

Arellano, Manuel and Stephen Bond (1991). Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *The Review of Economic Studies* 58(2): 277-97.

Bakos, Yannis, Brynjolfsson, Erik and Lichtman, Douglas (1999). "Shared Information Goods." *Journal of Law and Economics*. 42: 117-156.

Belden Associates (2003). *Belden Sales and Site Survey Summary: Q1-2 2003 Report*. `http://www.beldenassociates.com`.

Berry, Steven (1994). "Estimating Discrete-Choice Models of Product Differentiation." *Rand Journal of Economics*. 25: 242-262.

BigChampagne (2006). "Average Simultaneous U.S. Users: August 2002-May 2006." personal correspondence.

Billboard (2006). Fuzzy Math. 1 July: 24-26.

Blundell, Richard and Stephen Bond (1998). Initial Conditions and Moment Restrictions in Dynamic Panel Data Models. *Journal of Econometrics* 87(1): 115-43.

Boldrin, Michele and David Levine (2003). "Perfectly Competitive Innovation." UCLA working paper.

Bresnahan, Timothy, Scott Stern, and Manuel Trajtenberg (1997). "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the late 1980s." *Rand Journal of Economics*. 28: S17-S44.

DEG (2005). "Industry Boosted by $21.2 Billion in Annual DVD Sales and Rentals." The Digital Entertainment Group. `http://www.dvdinformation.com`.

Edison Media Research (2003). *The National Record Buyers Study III*. Sponsored by Radio & Records. `http://www.edisonresearch.com`.

Einav, Liran (forthcoming). "Seasonality in the U.S. Motion Picture Industry". *Rand Journal of Economics*.

Forrester (2002). "Downloads Save the Music Business." `http://www.forrester.com`.

Forrester (2004)

Gentzkow, Matthew (2006). "Valuing New Goods in a Model with Complementarity: Online Newspapers." Chicago University working paper.

Goolsbee, Austan (2000). In a World Without Borders: The Impact of Taxes on Internet Commerce. *Quarterly Journal of Economics*. 115: 561-576.

Goolsbee, Austan and Judy Chevalier (2003). "Price Competition Online: Amazon Versus Barnes And Noble," Quantitative Marketing and Economics, 1(2), June 2003, 203-222

International Federation of the Phonographic Industry (2002). *Recording Industry in Numbers 2001*. International Federation of Phonographic Industry.

Imbens, Guido and Joshua Angrist (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica*. 62: 467-475.

Jupiter Media Metrix (2002). "File Sharing: To Preserve Market Value Look Beyond Easy Scapegoats." `http://www.jupiterresearch.com`.

Karagiannis, Thomas, Andre Broido, Nevil Brownlee, kc claffy, and Michalis Faloutsos (2004). "Is P2P dying or just hiding?" Presented at Globecom 2004 in November-December 2004. `http://www.caida.org/outreach/papers/2004/p2p-dying/`.

Klein, Benjamin, Andres Lerner, and Kevin Murphy (2002). "The Economics of Copyright 'Fair Use' in a Networked World." *American Economics Association: Papers and Proceedings*. 92: 205-208.

Krueger, Alan (2004). "The Economics of Real Superstars: The Market for Rock Concerts in the Material World." *Journal of Labor Economics*. 23: 1-30.

Kultusministerkonferenz, Statistische Veröffentlichungen (2002). Nummer 162 vom August.

Lesk, Michael (2003). "Chicken Little and the Recorded Music Crisis." IEEE Security & Privacy. (September). 73-75.

Liebowitz, Stan (1985). "Copying and Indirect Appropriability: Photocopying of Journals." *Journal of Political Economy*. 93: 945-957.

Liang, Jian, Rakesh Kumar, and Keith Ross (2004). "Understanding KaZaA." working paper.

MPAA (2005). "U.S. Entertainment Industry: 2004 MPA Market Statistics." Motion Picture Association: Worldwide Market Research.

Musikmarkt (2002). "Deutschland Single-Charts." `http://musikmarkt.lw-t1.thuecom-medien.de/content/charts/history.php3?jahr=2002`.

Nevo, Aviv (2001). "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*. 69: 307-342.

Nielsen SoundScan (2005). `http://home.soundscan.com/about.html`.

Niesyto, Horst (2002). *Digitale Spaltung - digitale Chancen: Medienbildung mit Jugendlichen aus benachteiligenden Verhältnissen*. Mimeo, Pädagogische Hochschule Ludwigsburg.

Nisenholtz, Martin (2002). "The Death of Free Content?" Speech at Jupiter Media Forum 2002. `http://nytdigital.com/learn/jupiter_200203181.pdf`.

Oberholzer-Gee, Felix and Koleman Strumpf (2004). "The Effect of File Sharing on Record Sales: Supplemental Details and Estimates." working paper.

Oberholzer-Gee, Felix and Koleman Strumpf (2005). "The Effect of File Sharing on Record Sales: An Empirical Analysis." working paper.

Pew Internet

Plant, Arnold (1934). "The Economic Aspects of Copyright in Books." *Economica*. 1: 167-195.

Posner, Richard (2005). "Intellectual Property: The Law and Economics Approach." *Journal of Economic Perspectives*. 19,2: 57-73.

RIAA (1998). *RIAA 1996 Statistical Overview.* Archived copy from the Internet archive, `http://web.archive.org/web/19980124173720/www.riaa.com/market/releases/statover.htm`.

RIAA (2004). *RIAA Market Data: The Cost of a CD*. Archived copy from the Internet archive, `http://web.archive.org/web/20030416004543/http://www.riaa.com/MD-U.S.-7.cfm`.

RIAA (2006). *The Recording Industry Association of America's 2005 Yearend Statistics*. `http://www.riaa.com`.

Rob, Rafael and Joel Waldfogel (2006). "Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students." *Journal of Law and Economics*. 49(1): 29-62.

Romer, Paul (2002). "When Should We Use Intellectual Property Rights?" *American Economic Review: Papers and Proceedings*. 92: 2. 213-216.

Shapiro, Carl and Hal Varian (1999). *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press.

Takeyama, Lisa (1994). "The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities." *The Journal of Industrial Economics*. 42: 155-166.

Takeyama, Lisa (1997). "The Intertemporal Consequences of Unauthorized Reproduction of Intellectual Property." *Journal of Law & Economics*. 40: 511-22.

Varian, Hal (2000). "Buying, Sharing and Renting Information Goods." *The Journal of Industrial Economics*. 48: 473-488.

Windmeijer, Frank (2000). "A finite sample correction for the variance of linear two-step GMM estimators." Institute for Fiscal Studies, IFS Working Papers: W00/19.

Zentner, Alejandro (2006). "Measuring the Effect of Music Downloads on Music Purchases." *Journal of Law & Economics* 49(1): 63-90.

Table 1 – Sample Sales by Category

	Obs	Mean sales	Std dev	Min	Max
Full sample	680	143,096	344,476	74	3,430,264
Catalogue	50	46,833	40,031	219	223,085
Current Alternative	117	118,599	130,257	9,210	785,747
Hard Music Top Overall	19	28,304	22,103	2,945	86,416
Jazz Current	21	21,940	62,522	86	290,026
Latin	21	27,590	35,840	3,143	153,209
New artists	50	15,816	13,635	319	61,673
R&B	144	46,512	67,050	2,151	457,338
Rap	76	39,307	61,278	1,069	324,426
Top Current ("Billboard 200")	83	744,022	710,054	4,092	3,430,264
Top Current Country	66	87,839	130,096	74	669,575
Top Soundtrack	33	44,920	79,264	1,788	318,538

Note: These figures only include sales over our seventeen week observation period. Most of the top-selling albums are classified as "Current" for the purposes of this table

Table 2 – The Geography of File Sharing (numbers in %)

Country	Share of users	Share of downloads	Users in U.S. download from (%)	Users in U.S. upload to (%)	Share World Population	Share World GDP	Share World Internet Users	Software Piracy Rate
United States	30.9	35.7	45.1	49.0	4.6	21.2	27.4	23
Germany	13.5	14.1	16.5	8.9	1.3	4.5	5.3	32
Italy	11.1	9.9	6.1	5.7	0.9	2.9	3.2	47
Japan	8.4	2.8	2.5	1.8	2.0	7.2	9.3	35
France	6.9	6.9	3.8	4.7	1.0	3.1	2.8	43
Canada	5.4	6.1	6.9	7.9	0.5	1.9	2.8	39
United Kingdom	4.1	4.0	4.2	4.2	1.0	3.1	5.7	26
Spain	2.5	2.6	1.8	2.0	0.6	1.7	1.3	47
Netherlands	2.1	2.1	1.9	1.6	0.3	0.9	1.6	36
Australia	1.6	1.9	0.8	2.2	0.3	1.1	1.8	32
Sweden	1.5	1.7	1.8	1.5	0.1	0.5	1.0	29
Switzerland	1.4	1.5	0.9	1.0	0.1	0.5	0.6	32
Brazil	1.3	1.4	1.2	1.3	2.9	2.7	2.3	55
Belgium	0.9	1.2	0.5	1.0	0.2	0.6	0.6	31
Austria	0.8	0.6	0.6	0.4	0.1	0.5	0.6	30
Poland	0.5	0.7	0.7	0.5	0.6	0.8	1.1	54

Notes on country covariates:
Shares of users and downloads is from the file sharing dataset described in the text. All other statistics are from *The CIA World Factbook* (2002, 2003), except the software piracy rates which are from the *Eighth Annual BSA Global Software Piracy Study* (2003). All values are world shares, except the piracy rates are the fractions of business application software installed without a license in the country. All non-file sharing data are for 2002 except population which is for 2003.

Table 3 – Downloads by Genre

	# songs (# albums) in sample	Mean # of downloads	Std dev	Min	Max
			Song level		
All genres	10271	4.645	21.462	0	1258
Catalogue	714	4.361	10.370	0	152
Alternative	1707	7.021	18.153	0	312
Hard	270	4.830	8.684	0	52
Jazz	261	0.333	0.920	0	7
Latin	309	0.550	2.927	0	28
New artists	711	0.609	7.039	0	184
R&B	2249	1.635	7.680	0	159
Rap	1227	0.920	4.887	0	82
Current	1342	17.182	51.286	0	1258
Country	913	1.974	6.382	0	128
Soundtrack	568	1.673	5.301	0	61
			Album level		
All genres	680	70.162	158.628	0	1799
Catalogue	50	62.280	103.114	0	680
Alternative	117	102.436	122.794	0	674
Hard	19	68.632	82.899	0	264
Jazz	21	4.143	4.542	0	13
Latin	21	8.095	26.344	0	121
New artists	50	8.660	33.097	0	229
R&B	144	25.542	56.494	0	433
Rap	76	14.855	24.487	0	119
Current	83	277.807	333.935	2	1799
Country	66	27.303	51.649	0	344
Soundtrack	33	28.788	36.611	0	185

Table 4 – Downloads by Sales – Album Level

	Obs	Mean # of downloads	Std dev	Min	Max	Mann-Whitney
1st quartile: mean 7,235 copies [up to 12,493 copies]	170	11.358	38.472	0	402	- 14.067**
2nd quartile: mean 21,022 copies [up to 31,115 copies]	170	20.929	52.082	0	433	-12.431**
3rd quartile: mean 57.940 copies [up to 100,962 copies]	170	48.088	55.223	0	264	-8.187**
4th quartile: mean 486,184 copies [max 3,430,264 copies]	170	200.270	265.369	0	1799	

Mann Whitney test statistics are for the null that the 4th quartile with the highest sales comes from the same population as the other sales quartiles.
** 1% level of significance

Table 5 – Summary Statistics

	# obs	mean (std dev)	min	max
Sales (1,000s)	10093	9.580 (34.361)	0	874.137
Downloads	10093	4.360 (13.644)	0	368
German kids on Vacation (million)	10093	9.855 (3.576)	0	12.491
Band on tour in Germany	10093	0.003 (0.053)	0	1
Misspelling indicator	10093	0.062 (0.187)	0	1
Rank of single on German charts (calculated as 101 minus rank)	10093	1.576 (10.268)	0	100
Rank of single on MTV charts (calculated as 101 minus rank)	10093	2.158 (13.568)	0	100
Billboard rank previous album (calculated as 201 minus rank)	10093	61.136 (82.314)	0	200
Best Billboard rank ever (calculated as 201 minus rank)	10093	83.548 (89.994)	3	200
# previous releases	10093	6.718 (15.574)	0	194
HHI downloads	10093	2.460 (3.672)	0	10000

Table 6 – Download Times: Relation to Instruments and Impact on Number of Transfers

	(I)	(II)	(III)	(IV)	(V)	(VI) Impact of download time on download quantity	
	Time: Download Request to Initiation (sec) C_1	Time: Search Request to Download Request (sec) C_2	Time: Initiation Download to Completion (sec) C_3	Ratio: # Search Requests to # Downloads C_4	Percentage: Download Requests which are not completed C_5	Download Time (1st stage) $C_1+C_2+C_3$	Downloads (2nd stage) D_{it}
German kids on	-32.005	-4.336	-26.031	-0.453	-2.351	-62.420	
Vacation (million)	(5.51)**	(0.29)**	(2.69)**	(0.05)**	(0.10)**	(5.24)**	
German kids ×	-49.914	-3.966	-35.015	-0.480	-2.927	-89.010	
Band on tour	(20.31)*	(1.73)*	(13.35)**	(0.22)*	(0.51)**	(17.83)**	
German kids ×	22.494	6.157	8.609	0.672	1.963	7.302	
Misspellings	(33.66)	(2.182)**	(17.76)	(0.25)**	(0.58)**	(40.59)	
German kids ×	-0.347	-0.034	-0.471	-0.005	-0.024	-0.849	
rank German charts	(0.18)*	(0.02)	(0.16)*	(0.00)*	(0.01)*	(0.22)**	
Download time							-0.006
							(0.00)**
Album Fixed Effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1662	1952	1332	2164	1952	1332	1332
Mean for Dependent Variable	609.08	91.02	796.20	12.21	62.96	1491.18	7.25

Albums or album-weeks are omitted when the dependent variable is undefined (e.g. for C_1 when there are no successful album download initiations). Robust standard errors are in parentheses. These estimates are based on data from weeks 3-6 of our observation period.
** 1% level of significance * 5% level of significance

Table 7 – Panel Analysis - Downloads and Album Sales

	(I) Sales	(II) 1st stage down-loads	(II) 2nd stage Sales	(III) 1st stage down-loads	(III) 2nd stage sales	(IV) 1st stage down-loads	(IV) 2nd stage Sales	(V) 1st stage down-loads	(V) 2nd stage sales	(VI) 1st stage down-loads	(VI) 2nd stage Sales	(VII) 1st stage down-loads	(VII) 2nd stage sales
# downloads	0.277		0.003		0.024		-0.010		0.005		-0.027		0.037
	(0.025)**		(0.194)		(0.189)		(0.158)		(0.062)		(0.270)		(0.065)
German kids on vacation		0.671		0.670		0.667		1.818					
		(0.054)**		(0.054)**		(0.054)**		(0.125)**					
German kids × Band on tour				0.469		0.474		0.470		0.464		0.451	
				(0.168)**		(0.167)**		(0.161)**		(0.167)**		(0.161)**	
German kids × misspellings						-0.288				-0.290			
						(0.124)*				(0.124)*			
German kids × Germ charts						0.012		0.007		0.012		0.007	
						(0.001)**		(0.002)**		(0.001)**		(0.002)**	
U.S. MTV rank	0.079	0.036	0.089	0.037	0.088	0.035	0.089	0.058	-0.194	0.036	0.092	-0.042	-0.183
	(0.020)**	(0.008)**	(0.021)**	(0.008)**	(0.021)**	(0.008)**	(0.021)**	(0.103)	(0.256)	(0.008)**	(0.022)**	(0.102)	(0.255)
German kids × album FE	No	No	No	No	No	No	No	Yes	No	No	No	Yes	No
MTV × album FE	No	No	No	No	No	No	No	Yes	Yes	No	No	Yes	Yes
Polynomial time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No
Week FE	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Album FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10093	10093	10093	10093	10093	10093	10093	10093	10093	10093	10093	10093	10093
Prob χ^2>0 on excluded instruments		0.0000		0.0000		0.0000		0.0000		0.0000		0.0000	
Sargan test (p-value)					0.73		0.70		0.98		0.50		0.97
R-squared	0.75	0.74	0.76	0.74	0.76	0.73	0.76	0.74	0.79	0.82	0.77	0.85	0.79

Dependent variables are album sales (1,000s) and # downloads at the 1st stage. Robust standard errors are in parentheses. For the fixed-effects models, the reported *R*-squared is the sum of the explained within-variance and the fraction of the variance that is due to the fixed effects. Album-weeks prior to the release date are excluded from the sample.

** 1% level of significance * 5% level of significance

Table 8 – Dynamic Panel Analysis - Downloads and Lagged Album Sales

	(I) 2nd stage Sales	(II) 2nd stage sales	(III) 2nd stage sales	(IV) 2nd stage sales	(V) GMM Δ sales	(VI) GMM Δ sales
Weighted ∑ of three weeks of downloads (instrumented)	0.097	0.048	0.022	0.045		
	(0.115)	(0.039)	(0.170)	(0.041)		
Δ downloads					0.029	0.047
					(0.074)	(0.078)
U.S. MTV rank	0.092	-0.016	0.097	-0.022	0.085	0.041
	(0.015)**	(0.169)	(0.016)**	(0.168)	(0.091)	(0.080)
lagged sales					0.166	0.261
					(0.100))	(0.117)*
German kids × album FE in 1st stage	No	Yes	No	Yes	No	No
MTV × album FE	No	Yes	No	Yes	No	No
Polynomial time trend?	Yes	Yes	No	No	Yes	No
Week Fixed Effects?	No	No	Yes	Yes	No	Yes
Album Fixed Effects?	Yes	Yes	Yes	Yes	No	No
1st-stage specification is as in Table 7, model	IV	V	VI	VII		
Observations	8739	8739	8739	8739	8739	8739
Arellano-Bond test for AR(1) in first differences: Pr > z					0.302	0.204
Arellano-Bond test for AR(2) in first differences: Pr > z					0.638	0.522
R-squared	0.92	0.96	0.92	0.97		

The dependent variable is album sales (1,000s). The number of downloads is instrumented using the Table 7 specification listed in the fifth row from the bottom. The weighted sum of three weeks of downloads includes the current week. The weights are chosen in a grid search which minimizes the unexplained fraction of the variance in our models. Models (V) and (VI) use the Generalized Method of Moments estimator developed by Arellano and Bond (1991). In this model, the typical standard error estimator tends to be downwards biased (Blundell and Bond, 1998). Standard errors are corrected using the two-step covariance matrix derived by Windmeijer (2000). Arellano-Bond tests for autocorrelation are applied to the first-difference equation residuals. Second-order autocorrelation would indicate that some lags of the dependent variable which are used as instruments are endogenous. The tests reveal no such problem. Album-weeks prior to the release date are excluded from the sample.

** 1% level of significance * 5% level of significance

Table 9 – Robustness Check with Scaled Downloads – Testing the "Drop-out" Hypothesis

	(I)		(II)		(III)		(IV)	(V)	(VI)	(VII)	(VIII)
	1st stage downloads	2nd stage sales	1st stage downloads	2nd stage sales	1st stages downloads	2nd stage sales	2nd stage Sales	2nd stage Sales	2nd stage Sales	GMM Δ sales	GMM Δ sales
Scaled downloads		-0.009		0.022		0.029					
		(0.126)		(0.046)		(0.049)					
weighted ∑ of 3 weeks d'loads							0.078	0.038	0.037		
							(0.093)	(0.030)	(0.031)		
Δ downloads										0.072	0.123
										(0.053)	(0.072)
German kids on Vacation	0.856		2.608								
	(0.073)**		(0.171)**								
German kids × Band on tour	0.602		0.600		0.585						
	(0.225)**		(0.216)**		(0.216)**						
German kids × Misspellings	-0.377										
	(0.167)*										
German kids × German charts	0.014		0.008		0.008						
	(0.002)**		(0.002)**		(0.002)**						
U.S. MTV rank	0.036	0.089	-0.084	-0.198	-0.059	-0.182	0.093	0.139	-0.023	0.085	0.044
	(0.011)**	(0.020)**	(0.137)	(0.255)	(0.137)	(0.255)	(0.015)**	(0.158)	(0.168)	(0.097)	(0.077)
Lagged sales										0.166	0.261
										(0.101)	(0.118)*
German kids × album FE in 1st stage	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes	No	No
MTV × album FE	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes	No	No
Polynomial trend	Yes	Yes	Yes	Yes	No	No	Yes	Yes	No	Yes	No
Week FE	No	No	No	No	Yes	Yes	No	No	Yes	No	Yes
Album FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Specification as in Table (model)	7 (IV)	7 (IV)	7 (V)	7 (V)	7 (VII)	7 (VII)	8 (I)	8 (II)	8 (IV)	8 (V)	8 (VI)
Observations	10093	10093	10093	10093	10093	10093	8739	8739	8739	8739	
R-squared	0.74	0.76	0.85	0.79	0.87	0.79	0.82	0.86	0.87		
AB test for AR(1)										0.305	0.201
AB test for AR(2)										0.643	0.531

Dependent variables are album sales (1,000s) and scaled downloads at the 1st stage. Downloads are scaled to reflect the growth of KaZaA users over the sample period. For the fixed-effects models, the reported *R*-squared is the sum of the explained within-variance and the fraction of the variance that is due to the fixed effects. Album-weeks prior to the release date are excluded from the sample.
** 1% level of significance * 5% level of significance

Table 10 – Robustness Checks

(IV) Coefficient downloads (std. error)	(VI) Coefficient downloads (std. error)	(VII) Coefficient downloads (std. error)	*N*	Specification
-0.010 (0.158)	0.005 (0.062)	0.037 (0.065)	10093	Benchmark specifications, models (IV), (VI) and (VII) in Table 7
	Changes in Sample			
0.064 (0.376)	-0.001 (0.108)	-0.013 (0.112)	7399	Without holiday sales
0.018 (0.166)	0.034 (0.071)	0.079 (0.075)	7890	Without albums that are not downloaded
0.051 (0.184)	0.083 (0.090)	0.161 (0.097)	5033	Albums that sell more than 151,284 copies (50th percentile) during the sample period
0.037 (0.135)	0.062 (0.055)	0.092 (0.058)	8567	Without Latin and Country albums
	Changes in Model Specification			
-0.006 (0.007)	0.001 (0.003)	0.004 (0.003)	10093	Dependent variable is log of sales
0.083 (0.029)**	0.019 (0.026)	0.005 (0.022)	3232	Sales and downloads are expressed as percentage changes
	Does the estimated effect vary by popularity?			
Main effect downloads	Interaction	H_0 sum = 0 (Prob > F)		Downloads (instrumented) are interacted with…
-0.095 (0.185)	0.001 (0.001)	0.6119	10093	Billboard rank of artist's prior album
-0.130 (0.192)	0.001 (0.001)	0.5015	10093	Best Billboard rank for artist during career
0.002 (0.181)	0.002 (0.007)	0.9822	10093	Number of previous albums
-0.128 (0.175)	0.039 (0.026)	0.5917	10093	Herfindahl index measuring concentration of downloads

Dependent variables are album sales (1,000s) and # downloads at the 1st stage. Robust standard errors are in parentheses. For the popularity results in the lower panel, the specification is model (V) in Table 7. Album-weeks prior to the release date are excluded from the sample.
** 1% level of significance * 5% level of significance

Table 11 – Hypotheses Tests

Class of Models	Can reject hypothesis that impact of file sharing is larger than (in million CDs)
All models (Tables 7 through 9)	-20.1
GMM models	-14.3
Models with German vacation × Album FE interactions	-10.6
Models with scaled downloads (Table 9)	-10.3
5 models with smallest standard errors	-7.4

These values represent the overall, industry-wide impact of file sharing for 2002 as implied by the various specifications. The lower bound is the minimum of the 95% confidence interval around the mean. This is calculated as $\sum_t\sum_i (D_{it}\times5.04\times1000)\times(\gamma-2\times se(\gamma)) = 240m\times(\gamma-2\times se(\gamma))$, where γ is the point estimate from equation (1). The factor 5.04 scales the results from our sample to all releases and the entire year 2002. It is calculated as: Aggregate impact = (Effect of file sharing on sample sales over observation period) × (population sales/sample sales) × (file sharing activity over year/file sharing activity in observation period). From our sales data, the ratio (population sales/sample sales) is 2.27. The second ratio is (File sharing activity over year/file sharing activity in observation period) = 2.22, which is calculated from weekly file sharing traffic rates over the 2002 calendar year on the Internet2 backbone (Internet2 Netflow Statistics, 2004) and the monthly average number of U.S. file sharing users (BigChampagne, 2006). Note that the second conversion factor is close to a naïve correction based simply on time, (52 weeks in year/17 weeks in observation period) = 3.06.

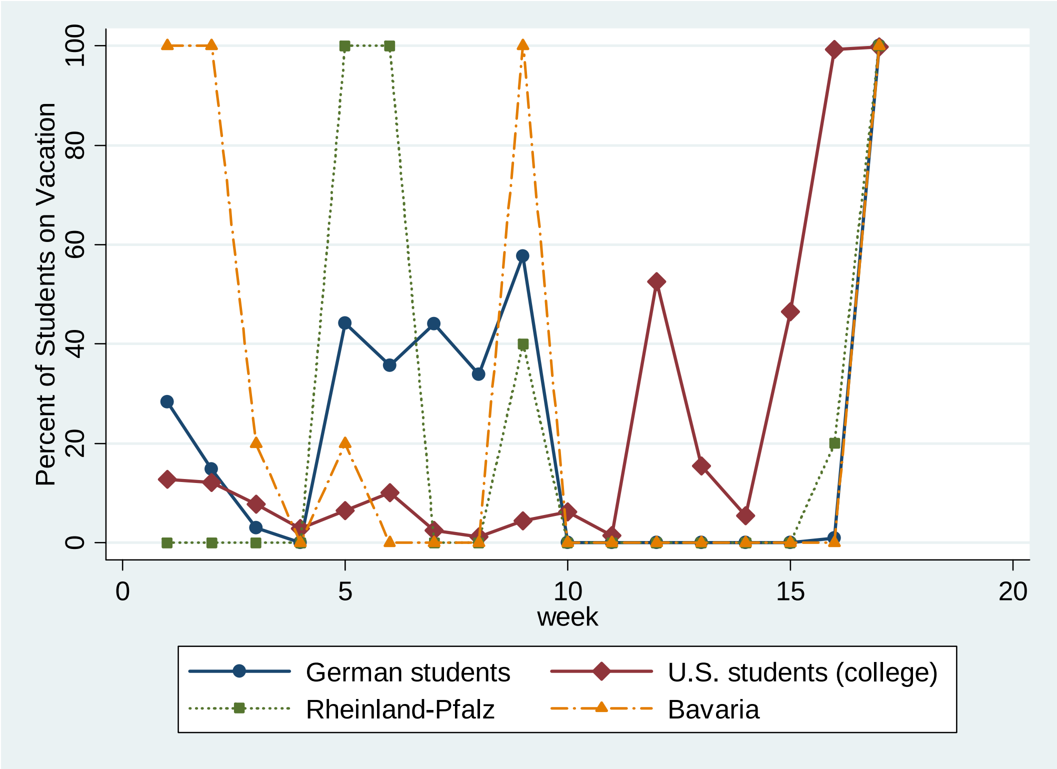

Figure 1: Timing of German and U.S. School Vacations

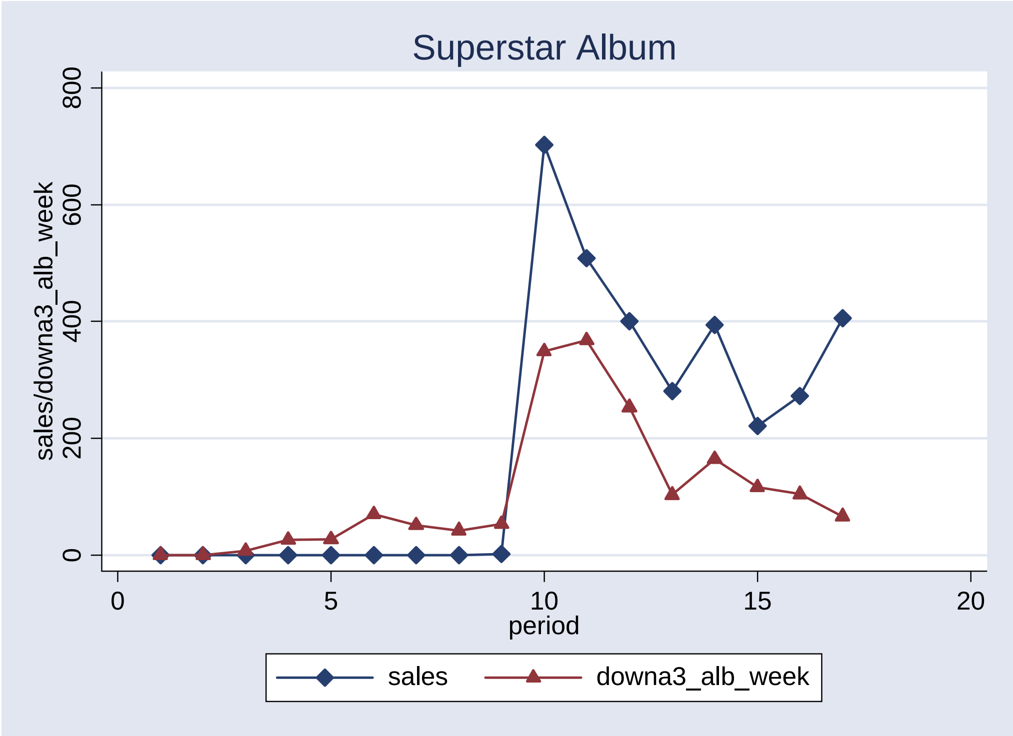

Figure 2: Dynamics of Downloads and Albums Purchases
(by week, sales in thousands)