

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

Felix Oberholzer-Gee

Harvard University

Koleman Strumpf

University of Kansas

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 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 that is statistically indistinguishable from zero. Our estimates are inconsistent

We would like to thank Bharat Anand, Gary Becker, Bob Frank, Shane Greenstein, Austan Goolsbee, Alan Krueger, Steven Levitt, Tom Mroz, Alan Sorensen, Joel Waldfoegel, Steven Wildman, Pai-Ling Yin, participants at numerous seminars, and two anonymous referees for helpful comments. 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 and Christina Hsiung Chen provided superb research assistance. The financial support of the George F. Baker Foundation (Oberholzer-Gee) and the Kenan Faculty Fund (Strumpf) is gratefully acknowledged. We appreciated the aural support from Massive Attack, Sigur Ros, and the Mountain Goats.

[*Journal of Political Economy*, 2007, vol. 115, no. 1]
© 2007 by The University of Chicago. All rights reserved. 0022-3808/2007/11501-0001\$10.00

with claims that file sharing is the primary reason for the decline in music sales during our study period.

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 percent in the last two years (Karagiannis et al. 2004; *Billboard* 2006). Sharing files is largely nonrivalrous because the original owner retains his or her 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 10 million simultaneous users on the major peer-to-peer (P2P) networks (Big Champagne 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 compact discs shipped in the United States fell by 25 percent 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 that 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 (Takeyama 1997; Bakos, Brynjolfsson, and Lichtman 1999; 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 because of the ever-present option of downloading or increase it through network ef-

fects and the greater ease of sharing (Takeyama 1994). Finally, it is possible that 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. In a recent study, Rob and Waldfogel (2006) find 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 data set includes 0.01 percent 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 10,000 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. Our instruments are relevant since they have a large impact on file transfer time, which in turn is a key determinant of the number of downloads.

We find that file sharing has had only a limited effect on record sales. After we instrument 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 one is considering the policy implications of these results, it is important to take into account the precision of our estimates. On the basis of 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 percent of sales and less than one-third of the observed decline in 2002). Models that consider the dynamics of file shar-

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 that 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 and the American public as the Boston strangler is to the woman home alone" (U.S. House, April 12, 1982). By 2004, 72 percent of domestic industry revenues came from rentals or sales of videotapes and digital video discs (Digital Entertainment Group 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.

ing 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 6.6 million CD sales or less than 10 percent of the 2002 decline. We arrive at similar conclusions if we allow the effect of international school holidays to vary by album. Our results continue to hold after we permit downloads to influence sales with a lag, omit data from the holiday shopping season, and restrict our sample to popular titles. 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 (2004) Consumer Expenditure Survey for the United States, households without a computer, which seem unlikely to engage in file sharing, report that they reduced their spending on CDs by 43 percent since 1999. Quasi-experimental evidence on the long-term effect of P2P on music sales also leads 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 by 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 microdata results. While the main estimates focus on high-frequency variation over several months, the experiments focus on long-term trends using data spanning several years.

Our results have broader implications beyond the specific case of file sharing. A long-standing 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 (2002) and Klein, Lerner, and Murphy (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 that studies the interactions

between the Internet and brick and mortar economies (Goolsbee 2000; Gentzkow, forthcoming).

The paper is organized as follows. Section II 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 in Section V. Section VI presents the results, and Section VII 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, because of shortcomings with the data. Most of what we know about the effect of file sharing on sales is based on surveys. There are numerous industry studies that arrive at a diverse range of conclusions. For instance, Forrester Research (2002) and Jupiter Media Metrix (2002) find neutral or positive effects, whereas 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 microdata 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 that sold more than 2 million copies since 1999, they find no relationship between downloading and sales. When the set of albums is expanded to include all music the students acquired in 2003, downloading five albums displaces the sale of one CD. These results could mean that piracy does not affect hit albums but hurts smaller artists, or it is also possible that file sharing had less of an effect on sales in earlier years. After we instrumented for downloads with the school the students attend—everyone at Penn has broadband access, whereas 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, nei-

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

ther the Rob and Waldfogel study nor Zentner's work allows 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 that allow the transfer of data. Each computer may agree to share some files and has the ability to search for and download files from other computers in the network. Our data come from the OpenNap network, an open-source descendant of Napster. OpenNap is an example of a centralized P2P network in which users log on to a central server that tracks all search requests and file downloads. During our study period in the fall of 2002, P2P networks were already quite large. FastTrack (which includes the popular KaZaA service; see Liang, Kumar, and Ross 2004) 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, the choice of about 1 percent of all P2P users, 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 (<http://home.soundscan.com/about.html>, 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 subsection.

A. *File-Sharing Data*

Our data were collected from two OpenNap servers, which operated continuously for 17 weeks from September 8 to December 31, 2002. The information on file transfers is collected as part of the log files that

the servers generate, and most users are unaware that their actions are being observed and recorded. An excerpt of a typical log file follows:

```
[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)": FILE NAME CONTAINS "kid rock devil" MAX_RESULTS 200 BIT RATE "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 entry in the log file shows user "evnormski" downloading the song "Devil without a Cause" by Kid Rock from user "bobo-joe." Information on downloads forms 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, or about 0.01 percent of all downloads in the world. We restrict the analysis to audio files by users in the United States. The server logs include the Internet protocol 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.³ 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 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 songs, we cannot reject a null that the two download samples are drawn from the same population (Pearson χ^2 statistic is 1,824.1). 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 that allows them to simultaneously participate in several networks. For example, roughly one-third of OpenNap participants use the WinMX software, which allows them to simultaneously access 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

³ A more comprehensive discussion of this point is in app. A of Oberholzer-Gee and Strumpf (2005).

little effect on the *distribution* of downloads. On the basis of these tests, we conclude that our sample is representative of the file transfers on the major P2P networks during our study period.

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.⁴ The sample is drawn from a *population* of albums on 11 charts produced by Nielsen SoundScan: alternative albums (a chart with 50 positions), hard music top overall (100), jazz current (100), Latin overall (50), rhythm and blues 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 that sold fewer than 100 copies during our study period, and the twenty-fifth percentile of sales in our data is only 12,493 copies.⁵ 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-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 2002.⁶ 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 percent 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

⁴The genre charts we sample from made up 81.8 percent 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 percent for the Big Five record companies, and 97 percent of the albums on the annual version of these charts were released on labels associated with the Recording Industry Association of America (RIAA).

⁵A typical measure of album success is gold certification, which occurs at sales of half a million copies.

⁶Albums 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.

TABLE 1
SAMPLE SALES BY CATEGORY

	Observations	Mean Sales	Standard Deviation	Minimum	Maximum
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 include sales only over our 17-week observation period. Most of the top-selling albums are classified as “current” for the purposes of this table.

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 songs that 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 song “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 percent.

There are two reasons why this rate is less than 100 percent. First, a downloaded song may not be 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 percent 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 a result, we believe that the 18 percent 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. A first stylized fact is that file sharing is truly global in nature. While over 90 percent of users are in developed countries, a total of 150 countries are represented in the data. U.S. users make up 31 percent 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 percent of their files from other U.S. users, with the remainder coming from a diverse range of countries including Germany (16.5 percent), Canada (6.9 percent), and Italy (6.1 percent).

While file-sharing activities are dispersed geographically, only a limited number of songs are transferred with any frequency. Table 3 shows that the average song is downloaded 4.6 times over the study period, but the median number of downloads is zero.⁷ 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 percent of the songs in our sample are never downloaded. When we aggregate 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 1,799 downloads, whereas the median number of downloads per album is 16, the seventy-fifth percentile is 63, the ninetieth percentile is 195, and the ninety-fifth 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. Table 3 shows that songs on the top current chart (Billboard 200) are most frequently downloaded. Downloads from this chart alone make up 48 percent of all file transfers. Another 25 percent come from the “alternative” category. The remaining nine charts are not particularly popular among file sharers. In view of the low cost of sharing and sampling music on P2P, one could expect 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. This suggests that

⁷ The seventy-fifth percentile of downloads per song is 2, the ninetieth percentile is 11, and the ninety-fifth percentile is 22.

TABLE 2
THE GEOGRAPHY OF FILE SHARING (%)

Country	Share of Users (1)	Share of Downloads (2)	Users in U.S. Download from (3)	Users in U.S. Upload to (4)	Share World Population (5)	Share World GDP (6)	Share World Internet Users (7)	Software Piracy Rate (8)
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	.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	.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	.6	1.7	1.3	47
Netherlands	2.1	2.1	1.9	1.6	.3	.9	1.6	36
Australia	1.6	1.9	.8	2.2	.3	1.1	1.8	32
Sweden	1.5	1.7	1.8	1.5	.1	.5	1.0	29
Switzerland	1.4	1.5	.9	1.0	.1	.5	.6	32
Brazil	1.3	1.4	1.2	1.3	2.9	2.7	2.3	55
Belgium	.9	1.2	.5	1.0	.2	.6	.6	31
Austria	.8	.6	.6	.4	.1	.5	.6	30
Poland	.5	.7	.7	.5	.6	.8	1.1	54

NOTE.—Shares of users and downloads are taken from the file-sharing data set described in the text. All other statistics are taken from the 2002 and 2003 editions of the *World Factbook* published by the Central Intelligence Agency, except the software piracy rates, which are taken from the Business Software Alliance (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 pertain to 2002 except population, which pertains to 2003.

TABLE 3
 DOWNLOADS BY GENRE

	Number in Sample	Mean Number of Downloads	Standard Deviation	Minimum	Maximum
A. Song Level					
All genres	10,271	4.645	21.462	0	1,258
Catalogue	714	4.361	10.370	0	152
Alternative	1,707	7.021	18.153	0	312
Hard	270	4.830	8.684	0	52
Jazz	261	.333	.920	0	7
Latin	309	.550	2.927	0	28
New artists	711	.609	7.039	0	184
R&B	2,249	1.635	7.680	0	159
Rap	1,227	.920	4.887	0	82
Current	1,342	17.182	51.286	0	1,258
Country	913	1.974	6.382	0	128
Soundtrack	568	1.673	5.301	0	61
B. Album Level					
All genres	680	70.162	158.628	0	1,799
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	1,799
Country	66	27.303	51.649	0	344
Soundtrack	33	28.788	36.611	0	185

common factors drive downloads and sales, which 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 at the album-week level for 17 weeks. These panel data allow us to estimate a model with album fixed effects,

$$S_{it} = \mathbf{X}_{it}\beta + \gamma D_{it} + \omega_s t^s + \nu_i + \mu_{it} \quad (1)$$

where i indicates the album, t denotes time in weeks, S_{it} is observed sales, \mathbf{X}_{it} is a vector of time-varying album characteristics that includes a measure of the title's popularity in the United States, D_{it} is the number of downloads for all songs on an album, and ω_s controls for time trends (a flexible polynomial or week fixed effects). The key concern in our

TABLE 4
 DOWNLOADS BY SALES: ALBUM LEVEL

	Mean Sales Quartile	Maximum Sales Quartile	Observations	Mean Number of Downloads	Standard Deviation	Minimum	Maximum	Mann- Whitney
1st quartile	7,235	12,493	170	11.358	38.472	0	402	-14.067**
2nd quartile	21,022	31,115	170	20.929	52.082	0	433	-12.431**
3rd quartile	57,940	100,961	170	48.088	55.223	0	264	-8.187**
4th quartile	486,184	3,430,264	170	200.270	265.369	0	1,799	

NOTE.—Mann-Whitney test statistics pertain to the null that the fourth quartile with the highest sales comes from the same population as the other sales quartiles.

** Significant at the 1 percent level.

empirical work is that the number of downloads is likely to be correlated with unobserved album-level heterogeneity. As the descriptive statistics suggest, the popularity of an album is likely to drive both file sharing and sales, implying that 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 two-stage least-squares 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. 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).⁸

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.⁹ German school vacations produce an increase in the supply of files and make it easier for U.S. users to download music.¹⁰ During holidays German teens can spend more time trading music online, since they do most of their file sharing at home (Niesyto 2002). School vacations also allow the German kids to stay up later, which means that they can engage in file sharing during the peak U.S. trading hours (early evening, eastern standard time). 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.

⁸ Appendix B of Oberholzer-Gee and Strumpf (2005) presents a formal model of purchase and download behavior that is the foundation for our econometric approach. In particular, it shows why we can use linear demand equations rather than the more complicated transformations that are typical in this literature (Berry 1994; Bresnahan, Stern, and Trajtenberg 1997).

⁹ The important role of German file-sharing users is documented in the authoritative BigChampagne database (OECD 2004). Oberholzer-Gee and Strumpf (2005) provide intuition on why this connection is so strong.

¹⁰ Appendix C of Oberholzer-Gee and Strumpf (2005) shows that German users are always net suppliers to file-sharing networks, and this effect is accentuated during weeks in which many kids are on vacation.

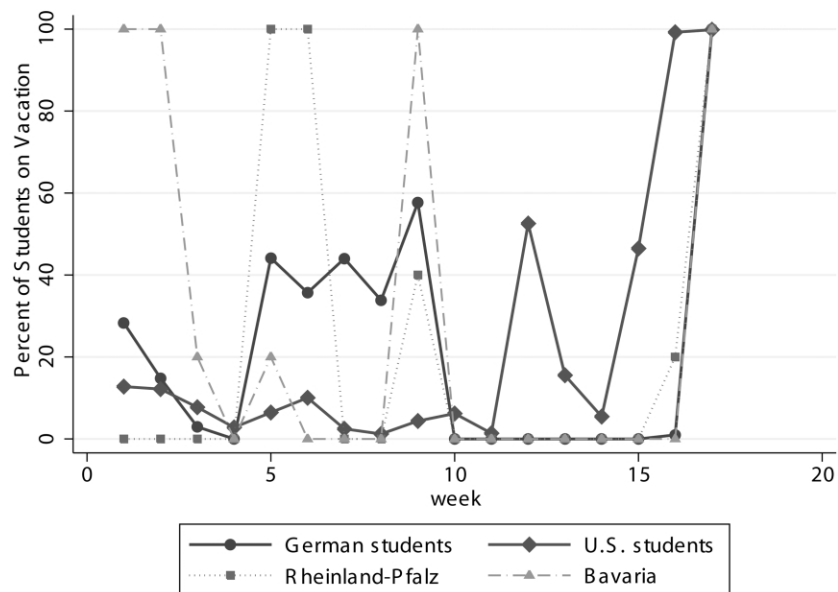


FIG. 1.—Timing of German and U.S. school vacations

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 16 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 whereas Rheinland-Pfälzer kids were already back in school (see fig. 1). A second difference from 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, <http://www.agentur-lindner.de/special/schulferien/index.html>, 2004). As we explain in greater detail below, there are additional reasons to believe that 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 \times *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 such as 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 \times *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.¹¹ 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 Microsoft Word's spell checker to determine whether an album has any song titles with an unconventional spelling. We expect misspellings to reduce the size of the positive supply shock coming from German vacations.

German kids on vacation \times *rank of album on German charts*.—Songs from popular albums 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 differences in the dynamics of song popularity in the two countries due to taste differences and differences in album 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*

Our analysis presumes that each instrument influences download costs and that these costs affect the number of file transfers. We test this idea by analyzing more detailed server log files that allow us to calculate the download time and success rate of download attempts. We construct

¹¹ For example, "lose yourself," the name of a popular song, would typically return over 1,000 results, but mistyping even one character (such as "lose yourse;f") or omitting part of a word ("lose yours") returned zero results.

TABLE 5
SUMMARY STATISTICS ($N = 10,093$)

	Mean	Minimum	Maximum
Sales (1,000s)	9.580 (34.361)	0	874.137
Downloads	4.360 (13.644)	0	368
German kids on vacation (millions)	9.855 (3.576)	0	12.491
Band on tour in Germany	.003 (.053)	0	1
Misspelling indicator	.062 (.187)	0	1
Rank of single on German charts (calculated as $101 - \text{rank}$)	1.576 (10.268)	0	100
Rank of single on MTV charts (cal- culated as $101 - \text{rank}$)	2.158 (13.568)	0	100
Billboard rank previous album (cal- culated as $201 - \text{rank}$)	61.136 (82.314)	0	200
Best Billboard rank ever (calculated as $201 - \text{rank}$)	83.548 (89.994)	3	200
Number of previous releases	6.718 (15.574)	0	194
Herfindahl-Hirschman index downloads	2.460 (3.672)	0	10,000

NOTE.—Standard deviations are in parentheses.

five measures of download costs: 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 that 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. Columns 1–3 show that the typical file takes 20 minutes to download, starting from the initial search until the transfer is complete.¹² There are also long delays for top-selling albums, suggesting that there is a ubiquitous scarcity of supply. While slow download speeds are the norm in our data, the estimates in table 6 show that searching and downloading audio files in the United States are considerably easier when a larger number of German school children are on vacation. This reduction is even larger

¹² Gummedi et al. (2003) independently document these long download times. This likely reflects the fact that only a third of the U.S. users in our data had a broadband connection.

TABLE 6
 DOWNLOAD TIMES: RELATION TO INSTRUMENTS AND IMPACT ON NUMBER OF TRANSFERS

	TIME: DOWNLOAD REQUEST TO INITIATION (Sec.) C_1 (1)	TIME: SEARCH REQUEST TO DOWNLOAD REQUEST (Sec.) C_2 (2)	TIME: INITIATION DOWNLOAD TO COMPLETION (Sec.) C_3 (3)	RATIO: NUMBER OF SEARCH REQUESTS TO NUMBER OF DOWNLOADS C_4 (4)	PERCENTAGE: DOWNLOAD REQUESTS NOT COMPLETED C_5 (5)	IMPACT OF DOWNLOAD TIME ON DOWNLOAD QUANTITY (6)	
						Download Time (1st Stage) $C_1 + C_2 + C_3$	Downloads (2nd Stage) D_{it}
German kids on vacation (millions)	-32.005 (5.51)**	-4.336 (.29)**	-26.031 (2.69)**	-.453 (.05)**	-2.351 (.10)**	-62.420 (5.24)**	
German kids × band on tour	-49.914 (20.31)*	-3.966 (1.73)*	-35.015 (13.35)**	-.480 (.22)*	-2.927 (.51)**	-89.010 (17.83)**	
German kids × misspellings	22.494 (33.66)	6.157 (2.182)**	8.609 (17.76)	.672 (.25)**	1.963 (.58)**	7.302 (40.59)	
German kids × rank German charts	-.347 (.18)*	-.034 (.02)	-.471 (.16)*	-.005 (.00)*	-.024 (.01)*	-.849 (.22)**	
Download time							-.006 (.00)**
Album fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,662	1,952	1,332	2,164	1,952	1,332	1,332
Mean for dependent variable	609.08	91.02	796.20	12.21	62.96	1,491.18	7.25

NOTE.—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 (the data come from more detailed log files that are available only during these weeks).

* Significant at the 5 percent level.
 ** Significant at the 1 percent level.

when the artist is on tour and when the album is highly ranked on the German charts.¹³ 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 , and C_5), but it has little effect on the time it takes to transfer a file (C_1 and 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 German 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 that they are insensitive to the time cost of file sharing. To check whether the variation in download time that is due to our instruments has a significant impact on the number of transfers, we estimate the system

$$\begin{aligned} C_{it} &= \mathbf{Z}_{it}\delta + \nu_i + \mu_{it} \\ D_{it} &= C_{it} + \nu_i + \epsilon_{it} \end{aligned} \quad (2)$$

where \mathbf{Z}_{it} is the full list of instruments and C_{it} denotes total download time ($C_1 + C_2 + C_3$). Columns 6 of table 6 show 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 equation (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 United States.¹⁴

D. *Specific Concerns with Individual Instruments*¹⁵

German kids on vacation.—A potential difficulty with the vacation variable is that it might be correlated with time-varying album popularity

¹³ 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 United States is driving the parameter estimates.

¹⁴ 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. Appendix C of Oberholzer-Gee and Strumpf (2005) presents this point in detail.

¹⁵ A general concern is that the instruments are based on high-frequency variation in download costs. Unfavorable conditions might lead users to simply defer downloads to a later time, in which case our second-stage estimates will be attenuated to zero. Oberholzer-Gee and Strumpf (2005) show that this concern is not warranted, since users are impatient and quickly lose interest in an album.

in the United States. We perform a number of tests to see whether this is the case. First, we check whether German vacations happen to coincide with official U.S. holidays. We find that there is little overlap.¹⁶ 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 are no centralized data on holidays for all 14,000 U.S. school districts, we collect information on the number of 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 United States. 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 United States, 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 whether 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.¹⁷

In a final test, we check more directly whether 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.¹⁸ 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. This interaction is included in both stages of equation (1):

$$D_{it} = \mathbf{X}_{it}\beta + \mathbf{Z}_{it}\delta + \varphi_1 \text{Gkids}_t \times \text{MTV}_{it} + \omega_{1s}t^s + \nu_i + \epsilon_{it}$$

$$S_{it} = \mathbf{X}_{it}\beta + \gamma \hat{D}_{it} + \varphi_2 \text{Gkids}_t \times \text{MTV}_{it} + \omega_{2s}t^s + \nu_i + \mu_{it} \quad (3)$$

where \mathbf{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 United States. In the second stage, however,

¹⁶ Estimates over our 17-week observation period yield

$$\text{U.S. Holidays}_t = 1.148 (1.61) - 0.182 (0.16) \times \text{German Kids}_t$$

where U.S. 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.

¹⁷ When we control 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.

¹⁸ We thank one of our referees for this suggestion. We also used the Billboard airplay ranking to explore these effects, with similar results.

φ_2 is economically small and statistically insignificant. When an album becomes more popular in the United States, 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 percent are also on the top 100 German charts. More generally, we study Amazon rankings to compare sales ranks in the two countries (Chevalier and Goolsbee 2003). With the exception of Latin and country music, Wilcoxon matched-pairs signed-ranks tests cannot reject the null of equal distributions for the 11 genres in our sample. In the robustness section of the paper, we test whether the undersupply of Latin and country music affects our estimates. We show that this is not the case, suggesting that the measured effect of downloads on sales is likely to be a good estimate of the average population effect.

German kids on vacation \times *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.¹⁹ Second, all our specifications presented in the results section include album fixed effects that 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 important, the second-stage estimates would be attenuated toward zero since the variation in fitted downloads would be largely due to noise. Several pieces of evidence suggest that this is not true. First, the estimates in the last subsection show that misspellings do in fact have real effects on transfer times and user behavior. Second, we can check for misspellings in unmatched downloads. If the criticism is correct, there

¹⁹ 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 (less than 10 percent of the mean) during our entire study period.

should be more misspellings in the unmatched than in the matched sample. This is not the case.²⁰

German kids on vacation \times *rank of album on German charts*.—The idea underlying this instrument is that vacation periods in Germany will boost downloads in the United States 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 identifies 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 Section VI 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.²¹ 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

$$D_{it} = Z_{it}\delta + \varphi_1 \text{Ukids}_t + \varphi_2 \text{Ukids}_t \times \text{Gcharts}_{it} + \varphi_3 \text{Ukids}_t \times \text{MTV}_{it} + \varphi_4 \text{Gkids}_t \times \text{MTV}_{it} + \omega_s t^s + \nu_i + \epsilon_{it}, \quad (4)$$

where 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 United States 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 \times German chart interaction in our later specifications. The data show that relative popularity in Germany interacts with German but not with U.S. vacations.

VI. Results

Before we turn to the estimates, it is instructive to graph some of the data. Figure 2 shows the weekly time series of sales and purchases for one of the most popular albums in our sample. This “superstar” album was largely ignored in file-sharing networks until it became available for

²⁰ The rates are 0.041 ($N = 35,614$) and 0.038 ($N = 7,163$) in the unmatched and matched samples, respectively. The Pearson χ^2 statistic is 1.402 .

²¹ We thank one of our referees for this suggestion.




FIG. 2.—Dynamics of downloads and album purchases for a popular album (by week, sales in thousands).

sale in week 10 of our sample. This suggests that it is the publicity associated with an official release that drives downloads as well as sales. Notice also the rapid but nonmonotone decay in sales and downloads, which highlights the importance of using high-frequency data.

A. Panel Analysis

In table 7 we report results for equation (1). The unit of observation is the album-week. The models include a control in both stages 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 ordinary least squares specification yields a large positive effect of 1.093 with a standard error of 0.023. A model that adds album fixed effects is given in specification 1. 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 (model 2). The first-stage estimates imply that a one-standard-deviation increase in the number of children on vacation boosts weekly album downloads by slightly more than one-half of their mean, an effect that is statistically significant and economically meaningful. Once we instrument for down-

TABLE 7
 PANEL ANALYSIS: DOWNLOADS AND ALBUM SALES ($N = 10,093$)

	MODEL 1	MODEL 2		MODEL 3		MODEL 4		MODEL 5		MODEL 6		MODEL 7	
	Sales	1st Stage Down- loads	2nd Stage Sales	1st Stage Down- loads	2nd Stage Sales	1st Stage Down- loads	2nd Stage Sales	1st Stage Down- loads	2nd Stage Sales	1st Stage Down- loads	2nd Stage Sales	1st Stage Down- loads	2nd Stage Sales
Number of downloads	.277 (.025)**		.003 (.194)		.024 (.189)		-.010 (.158)		.005 (.062)		-.027 (.270)		.037 (.065)
German kids on vacation		.671 (.054)**		.670 (.054)**		.667 (.054)**		1.818 (.125)**					
German kids \times band on tour				.469 (.168)**		.474 (.167)**		.470 (.161)**		.464 (.167)**		.451 (.161)**	
German kids \times misspellings						-.288 (.124)*				-.290 (.124)*			
German kids \times German charts						.012 (.001)**		.007 (.002)**		.012 (.001)**		.007 (.002)**	
U.S. MTV rank	.079 (.020)**	.036 (.008)**	.089 (.021)**	.037 (.008)**	.088 (.021)**	.035 (.008)**	.089 (.021)**	.058 (.103)	-.194 (.256)	.036 (.008)**	.092 (.022)**	-.042 (.102)	-.183 (.255)
German kids \times album fixed effects?	No	No	No	No	No	No	No	Yes	No	No	No	Yes	No
MTV \times album fixed effects?	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 fixed effects?	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Album fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prob $\chi^2 > 0$ on excluded instruments		.0000		.0000		.0000		.0000		.0000		.0000	
Sargan test (p -value)					.73		.70		.98		.50		.97
R^2	.75	.74	.76	.74	.76	.73	.76	.74	.79	.82	.77	.85	.79

NOTE.—The unit of analysis is the album-week. Dependent variables are the number of downloads at the first stage (summing all songs on an album) and album sales (1,000s). Robust standard errors are in parentheses. Since all models include album fixed effects, the reported R^2 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.

* Significant at the 5 percent level.

** Significant at the 1 percent level.

loads, 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 (models 3 and 4). 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 model 5 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 United States, our estimates would be biased. Specification 5 addresses this issue in four ways. As before, we include album fixed effects to make sure that it is the timing of the supply shocks that identifies 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 \times album fixed effects 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 equation (3), now with German kids \times album fixed effects \times U.S. MTV interactions. In the sales equation, these interactions are individually and collectively not different from zero.

Model 5 of table 7 reports results with the album interactions. Our instruments retain their statistical significance.²² The mean of the co-

²² The vacations \times misspellings interaction is collinear with vacations \times album fixed effects and cannot be included in this specification.

efficients on the vacation \times album fixed effects interactions is -1.143 , leaving the average effect of vacations on downloads almost unchanged from the earlier specifications. Grouping the album interactions by genre, we find that vacations increase downloads the most for music types that are popular in Germany: the mean of the vacation \times album fixed effects coefficients is -0.71 for international albums and -0.91 for rock. In contrast, the effect of vacations is much smaller, but still positive, for genres that are less popular in Germany (the mean interactions are -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 whether our results are driven by our modeling choice for the time trend in downloads and sales, we replace the polynomial time trend with week fixed effects in models 6 and 7 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 (model 7).

Table 7 suggests that 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 affect the music industry as a whole. If file sharing were to be eliminated, the most negative estimate (model 6) implies that industry sales for all of 2002 would increase by 6.5 million albums. With the most positive estimate (model 7), industry sales would fall by 8.9 million copies.²³ 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 allow only 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 that they

²³ The impact is the difference between predicted sales and the fitted value when downloads are set at zero. From eq. (1), the summed impact for our album sample and for our 17-week observation period is

$$\sum_t \sum_i S_u(D_{it}) - S_u(0) = \gamma \times \sum_t \sum_i D_{it}$$

We multiply this number by a scaling factor to get the annual impact for the entire music industry, $\gamma \times 240$ million (this calculation is described in more detail below in table 11).

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 method of moments (GMM) models.

A difficulty with the first approach is that 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 instrumented using the core set of instruments (specification 4 in table 7) or the extended set (specification 5). Our formal measure is the weighted stock of current and previous weekly downloads, $D_t^{\text{Stock}} = \sum_{s \geq 0} \delta_s \times D_{t-s}$.²⁴ In these models, we continue to find small and statistically insignificant effects for the weighted sum of three weeks of downloads, in specifications with both a polynomial time trend (models 1 and 2 of table 8) and week fixed effects (models 3 and 4). As in the panel results, standard errors drop significantly with the extended set of instruments (models 2 and 4). 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 5 and 6 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.²⁵ 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 that are used as instruments are endogenous, but the tests reveal no such problem. The results of these models, with a polynomial time trend as in model 5 or with week fixed effects as in model 6, are similar to our previous findings. The estimates are

²⁴ 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 $(\delta_0, \dots, \delta_7)$ are $(1, 0.1, 0.1)$. It is interesting that the weights that best fit our data give much importance to downloads in the current week, whereas downloads further back in the past do not heavily influence sales. Oberholzer-Gee and Strumpf (2005) present additional results showing that file sharers are impatient. These findings are consistent with those of Einav (forthcoming) for movie consumption.

²⁵ The formal model is

$$S_{it} = \alpha S_{i,t-1} + \mathbf{X}_{it}\beta + \gamma D_{it} + \omega_s t_s + \nu_i + \mu_{it}$$

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 album fixed effects. We instrument for the first-differenced $S_{i,t-1}$, which are now endogenous.

TABLE 8
DYNAMIC PANEL ANALYSIS: DOWNLOADS AND LAGGED ALBUM SALES ($N = 8,739$)

	MODEL 1: 2nd Stage Sales	MODEL 2: 2nd Stage Sales	MODEL 3: 2nd Stage Sales	MODEL 4: 2nd Stage Sales	MODEL 5: GMM Δ Sales	MODEL 6: GMM Δ Sales
Weighted sum of three weeks of downloads (instrumented)	.097 (.115)	.048 (.039)	.022 (.170)	.045 (.041)		
Δ downloads					.029 (.074)	.047 (.078)
U.S. MTV rank	.092 (.015)**	-.016 (.169)	.097 (.016)**	-.022 (.168)	.085 (.091)	.041 (.080)
Lagged sales					.166 (.100)	.261 (.117)*
German kids \times album fixed effects in 1st stage?	No	Yes	No	Yes	No	No
MTV \times album fixed effects?	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	4	5	6	7		
Arellano-Bond test in first differences: $Pr > z$:						
For AR(1)					.302	.204
For AR(2)					.638	.522
R^2	.92	.96	.92	.97		

NOTE.—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 that minimizes the unexplained fraction of the variance in our models. Models 5 and 6 use the GMM estimator developed by Arellano and Bond (1991). In this model, the typical standard error estimator tends to be downward 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 that are used as instruments are endogenous. The tests reveal no such problem. Album-weeks prior to the release date are excluded from the sample.

* Significant at the 5 percent level.

** Significant at the 1 percent level.

fairly precise, making these GMM models an alternative to using our extended set of instruments.

C. “Dropout” 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 “dropout” 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 2005). Ninety percent reported that they recently purchased a CD, a value reaching 100 percent among the most active downloaders.

Second, we can test the dropout 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 dropouts 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.²⁶ Because the number of users increased by over a third over our observation period, we should be able to detect a dropout effect if it exists. Table 9 reports these estimates for three panel models, three models using a stock of previous downloads, and two GMM models. In all these specifications, downloads still do not have a significant effect on sales. A third approach to testing the dropout hypothesis is to compare the long-run sales growth of individual genres of music. We return to this point in Section VII.

²⁶ We use 22 data points on the number of KaZaA users in the period from September 9, 2002, to February 4, 2003, to fit a fractional polynomial trend in the number of users. The model explains 85 percent of the variation.

TABLE 9
 ROBUSTNESS CHECK WITH SCALED DOWNLOADS: TESTING THE DROPOUT HYPOTHESIS

	MODEL 1		MODEL 2		MODEL 3		MODEL 4	MODEL 5	MODEL 6	MODEL 7	MODEL 8
	1st Stage Downloads	2nd Stage Sales	1st Stage Downloads	2nd Stage Sales	1st Stage Downloads	2nd Stage Sales	2nd Stage Sales	2nd Stage Sales	2nd Stage Sales	GMM Δ Sales	GMM Δ Sales
Scaled downloads		-.009 (.126)		.022 (.046)		.029 (.049)					
Weighted sum of three weeks' downloads							.078 (.093)	.038 (.030)	.037 (.031)		
Δ downloads										.072 (.053)	.123 (.072)
German kids on vacation (millions)	.856 (.073)**		2.608 (.171)**								
German kids × band on tour	.602 (.225)**		.600 (.216)**		.585 (.216)**						
German kids × misspellings	-.377 (.167)*										
German kids × rank German charts	.014 (.002)**		.008 (.002)**		.008 (.002)**						

U.S. MTV rank	.036 (.011)**	.089 (.020)**	-.084 (.137)	-.198 (.255)	-.059 (.137)	-.182 (.255)	.093 (.015)**	.139 (.158)	-.023 (.168)	.085 (.097)	.044 (.077)
Lagged sales										.166 (.101)	.261 (.118)*
German kids × album fixed effects in 1st stage?	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes	No	No
MTV × album fixed effects?	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes	No	No
Polynomial time trend?	Yes	Yes	Yes	Yes	No	No	Yes	Yes	No	Yes	No
Week fixed effects?	No	No	No	No	Yes	Yes	No	No	Yes	No	Yes
Album fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Specification as in table (model)	7 (4)	7 (4)	7 (5)	7 (5)	7 (7)	7 (7)	8 (1)	8 (2)	8 (4)	8 (5)	8 (6)
Observations	10,093	10,093	10,093	10,093	10,093	10,093	8,739	8,739	8,739	8,739	8,739
R^2	.74	.76	.85	.79	.87	.79	.82	.86	.87		
Arellano-Bond test:											
For AR(1)										.305	.201
For AR(2)										.643	.531

NOTE.—Dependent variables are album sales (1,000s) and are scaled downloads at the first stage. Downloads are scaled to reflect the growth of KaZaA users over the sample period. For the fixed-effects models, the reported R^2 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.

* Significant at the 5 percent level.

** Significant at the 1 percent level.

D. Robustness Tests

To further corroborate our results, we perform a large number of robustness checks, some of which we report in table 10.²⁷ 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 because of 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 4, 6, and 7 of 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 minuscule 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 whether the undersupply of Latin and country music influences our estimates. Recall from Section V.D that this would cause a problem only if the substitutability of downloads and album purchases varies across music genres. The last specification in panel A of table 10 reestimates 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 4 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.

In panel B of table 10, we explore two alternative specifications. To reduce the importance of outlier albums with a large number of sales, we use $\log(\text{sales})$ 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-differ-

²⁷ 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 that file sharing has little impact on sales.

encing and the weeks with zero downloads or sales. Using our core set of instruments, we now find a positive and statistically significant but economically small effect of downloads on sales. However, the estimated coefficient drops considerably and is insignificant when we introduce week fixed effects.

The previous models constrained the effect of downloads on sales to be identical for all releases. In panel C 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 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 minus actual rank so that larger numbers indicate greater popularity.²⁸ We estimate these models using specification 4 of 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 is 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 an album. As documented earlier, most file sharers obtain just a few songs from an album. One might suspect that P2P is a fairly good substitute for albums 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 stages. There is no evidence that albums with more concentrated downloads suffer disproportionately from file sharing.

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 due to seasonality, geography, music genre, or 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.

²⁸ More precisely, the term is a three-way interaction: downloads \times indicator that the artist had a Billboard ranking \times (201 - Billboard rank).

TABLE 10
ROBUSTNESS CHECKS

COEFFICIENT DOWNLOADS			OBSERVATIONS	SPECIFICATION
Table 7 (4)	Table 7 (6)	Table 7 (7)		
-.010 (.158)	.005 (.062)	.037 (.065)	10,093	Benchmark specifications, models 4, 6, and 7 in table 7
A. Changes in Sample				
.064 (.376)	-.001 (.108)	-.013 (.112)	7,399	Without holiday sales
.018 (.166)	.034 (.071)	.079 (.075)	7,890	Without albums that are not downloaded
.051 (.184)	.083 (.090)	.161 (.097)	5,033	Albums that sell more than 151,284 copies (50th percentile) during the sample period
.037 (.135)	.062 (.055)	.092 (.058)	8,567	Without Latin and country albums

B. Changes in Model Specification

-.006 (.007)	.001 (.003)	.004 (.003)	10,093	Dependent variable is log of sales
.083 (.029)**	.019 (.026)	.005 (.022)	3,232	Sales and downloads are expressed as percentage changes

C. Does the Estimated Effect Vary by Popularity?

Main Effect Downloads	Interaction	H_0 Sum = 0 (Prob > F)	Observations	Downloads (Instrumented) Are Interacted with:
-.095 (.185)	.001 (.001)	.6119	10,093	Billboard rank of artist's prior album
-.130 (.192)	.001 (.001)	.5015	10,093	Best Billboard rank for artist during career
.002 (.181)	.002 (.007)	.9822	10,093	Number of previous albums
-.128 (.175)	.039 (.026)	.5917	10,093	Herfindahl index measuring concentration of downloads

NOTE.—Dependent variables are album sales (1,000s) and number of downloads at the first stage. Robust standard errors are in parentheses. For the popularity results in panel C, the specification is model 5 in table 7. Album-weeks prior to the release date are excluded from the sample.

* Significant at the 5 percent level.

** significant at the 1 percent level.

The first experiment involves variation over time. The number of file-sharing users in the United States drops 12 percent over the summer (estimated from BigChampagne [2006]) because college students are away from their high-speed campus Internet connections. If downloads crowd out sales, we should observe that the share of albums sold in the summer increases following the advent of file sharing. 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 (1995–98) preceding the introduction of Napster, the average share of summer sales was 37.0 percent with a range of 36.4–37.8 percent. During the more recent period of extensive file sharing (1999–2005), the average share of summer sales was 37.2 percent with a range of 35.9–37.8 percent.

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. Because of time zone differences, such transfers are easier for East rather than West Coast users. The reason is that the peak file-sharing period (7:00 p.m. to 3:00 a.m.) 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.²⁹ 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 101 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 percent. This share has hardly moved since then. In 1999–2002, the mean was 43.5 percent and the range was 42.7–44.0 percent. This is consistent with some common national factors, rather than file sharing, driving sales trends.

A third experiment, which also provides a test of the dropout hypothesis, is to see whether download intensity influences long-run sales

²⁹ Unfortunately, Internet protocol addresses can be matched only imperfectly to locations, so this finding is merely suggestive.

growth after explicitly controlling for trends in music format popularity. The model for the period 1999–2005 is

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

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 dropout hypothesis predicts a greater sales reduction for genres that are popular on file-sharing networks. The estimated γ is not statistically significant using either download levels or downloads relative to purchases. For example, when we use mean downloads per album and control 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 slumps are common in the music business, even prior to file sharing. While real revenues have fallen 28 percent over 1999–2005, real revenue fell 35 percent during the collapse of disco music in 1978–83. Real sales also dropped 6 percent over 1994–97.³⁰ More direct evidence comes from regressing total album sales, including paid digital downloads, on the average number of simultaneous file-sharing users in the United States (BigChampagne 2006),

$$\text{Sales}_t = \gamma \times \text{Users}_t + \nu_m + \mu_p, \quad (6)$$

where t indicates a month, and ν_m are monthly fixed effects that account for seasonality. When we use monthly data from August 2002 to May 2006 ($N = 46$) and define Sales and Users in millions (with respective sample means of 56.0 million and 5.0 million), the estimated $\gamma = -0.427$ 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.³¹ The estimates remain insignificant if equation (6) is estimated in first differences.

³⁰ These are calculated from nominal RIAA revenues listed in RIAA (1998, 2006) and Lesk (2003).

³¹ If file sharing were eliminated, the point estimates imply that monthly sales would increase by only 2.1 million.

TABLE 11
HYPOTHESES TESTS

Class of Models	Lower Bound of 95% Confidence Interval ^a
All models (tables 7–9)	–24.1
Models with German vacation × album fixed effects interactions	–12.7
Models with scaled downloads (table 9)	–12.4
GMM models with scaled downloads (table 9)	–6.6
Five models with smallest standard errors	–6.0

NOTE.—These values represent the overall, industrywide impact of file sharing for 2002 as implied by the various specifications. The lower bound is the minimum of the 95 percent confidence interval around the mean impact. Details of this calculation are listed below. The entries are the median lower bound for that class of models. The lower bound is calculated as

$$\sum_i \sum_j (D_i \times 5.04 \times 1,000) \times [\gamma - 2 \times se(\gamma)] = 240m \times [\gamma - 2 \times se(\gamma)],$$

where γ is the point estimate from eq. (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 naive correction based simply on time: (52 weeks in year/17 weeks in observation period) = 3.06.

^a Can reject the hypothesis that the impact of file sharing is larger than the entry in the column (in millions of albums).

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 technology 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 (specification 6 of table 7) 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 95 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 P2P caused a sales decline greater than 24.1 million albums. For reference, the music industry sold 803 million CDs in 2002, which was a loss of 80 million from the previous year (RIAA 2004). Our estimates become more precise if we relax the assumption that file sharing affects only contemporaneous sales and if we allow for growth in the number of file sharers. For example, the scaled GMM models in table 9 reject a null of losses greater than 6.6 million. Relying on our five most precise estimates, we conclude that the impact could not have been larger than 6.0 million albums. While file sharers downloaded billions of files in 2002, the consequences for the industry amounted to no more than 0.7 percent 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, more than 14 percent of music sales shifted from record stores to more efficient discount retailers such as Wal-Mart, possibly reducing inventories. As a result, album shipments, which are often cited to document the decline in the legal demand for music, fell much more than actual sales.³² 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 toward 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 percent, 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, in 2005 retail music sales rose in four of the five largest national markets. Third,

³² In the 1999–2003 period, the number of shipped albums fell by 301 million, but the number of albums that were sold declined by only 99 million.

in the United States the entire drop in 2005 album sales is due to losses at a single firm, the recently merged Sony-BMG, which has experienced severe postmerger integration difficulties. If file sharing were responsible for the observed sales decline in the United States, we would not expect this activity to affect the products of only a single firm.

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. On the basis of our results, we do not believe that 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 songs downloaded each year, suggests that the added social welfare from file sharing is likely to be high.

References

- 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." *Rev. Econ. Studies* 58 (April): 277–97.
- Bakos, Yannis, Erik Brynjolfsson, and Douglas Lichtman. 1999. "Shared Information Goods." *J. Law and Econ.* 42, no. 1, pt. 1 (April): 117–56.
- Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *Rand J. Econ.* 25 (Summer): 242–62.
- BigChampagne. 2006. "Average Simultaneous U.S. Users: August 2002–May 2006." Manuscript, BigChampagne, Beverly Hills, CA.
- Billboard*. 2006. "Fuzzy Math." July 1, pp. 24–26.
- Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *J. Econometrics* 87 (November): 115–43.
- Boldrin, Michele, and David K. Levine. 2002. "The Case against Intellectual Property." *A.E.R. Papers and Proc.* 92 (May): 209–12.
- Bresnahan, Timothy F., Scott Stern, and Manuel Trajtenberg. 1997. "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the late 1980s." *Rand J. Econ.* 28 (special issue): S17–S44.

- Business Software Alliance. 2003. *Eighth Annual BSA Global Software Piracy Study*. <http://www.bsa.org/globalstudy2003/index.cfm>.
- Chevalier, Judy, and Austan Goolsbee. 2003. "Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com." *Quantitative Marketing and Econ.* 1 (June): 203–22.
- Digital Entertainment Group. 2005. "Industry Boosted by \$21.2 Billion in Annual DVD Sales and Rentals." <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. "Gross Seasonality and Underlying Seasonality: Evidence from the U.S. Motion Picture Industry." *Rand J. Econ.*
- Forrester Research. 2002. "Downloads Save the Music Business." <http://www.forrester.com>.
- . 2004. "US Antipiracy Bill Won't Stop File Sharing." <http://www.forrester.com>.
- Gentzkow, Matthew. Forthcoming. "Valuing New Goods in a Model with Complementarity: Online Newspapers." *A.E.R.*
- Goolsbee, Austan. 2000. "In a World without Borders: The Impact of Taxes on Internet Commerce." *Q.J.E.* 115 (May): 561–76.
- Gummadi, Krishna P., et al. 2003. "Measurement, Modeling, and Analysis of a Peer-to-Peer File-Sharing Workload." In *Proceedings of the 19th ACM Symposium on Operating Systems Principles* (Bolton Landing, NY, October 19–22). New York: Assoc. Computing Machinery.
- Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (March): 467–75.
- International Federation of the Phonographic Industry. 2002. *Recording Industry in Numbers 2001*. London: Internat. Federation Phonographic Industry.
- Internet2 Netflow Statistics. 2004. *Internet2 NetFlow: Weekly Reports*. <http://netflow.Internet2.edu/weekly/>.
- 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?" Paper presented at Globecom, Dallas, November–December. <http://www.caida.org/outreach/papers/2004/p2p-dying/>.
- Klein, Benjamin, Andres V. Lerner, and Kevin M. Murphy. 2002. "The Economics of Copyright 'Fair Use' in a Networked World." *A.E.R. Papers and Proc.* 92 (May): 205–8.
- Kultusministerkonferenz. 2002. *Statistische Veröffentlichungen*, no. 162 (August).
- Lesk, Michael. 2003. "Chicken Little and the Recorded Music Crisis." *IEEE Security & Privacy* 1 (September–October): 73–75.
- Liang, Jian, Rakesh Kumar, and Keith Ross. 2004. "Understanding KaZaA." Manuscript, Polytechnic Univ.
- MPAA (Motion Picture Association of America). 2005. "U.S. Entertainment Industry: 2004 MPA Market Statistics." Motion Picture Assoc. America, Worldwide Market Res., Washington, DC.
- 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 (March): 307–42.
- Niesyto, Horst. 2002. "Digitale Spaltung—digitale Chancen: Medienbildung mit

- Jugendlichen aus benachteiligten Verhältnissen." Manuscript, Pädagogische Hochschule Ludwigsburg.
- Oberholzer-Gee, Felix, and Koleman Strumpf. 2005. "The Effect of File Sharing on Sales: An Empirical Analysis." Manuscript, Harvard Bus. School and Univ. North Carolina at Chapel Hill.
- OECD. 2004. *OECD Information Technology Outlook 2004*. Paris: Organisation Econ. Cooperation and Development.
- Plant, Arnold. 1934. "The Economic Aspects of Copyright in Books." *Economica* 1 (May): 167–95.
- Posner, Richard A. 2005. "Intellectual Property: The Law and Economics Approach." *J. Econ. Perspectives* 19 (Spring): 57–73.
- RIAA (Recording Industry Association of America). 1998. *RIAA 1996 Statistical Overview*. Archived copy from the Internet archive, <http://web.archive.org/web/19980124173720/>, <http://www.riaa.com/market/releases/statover.htm>.
- . 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>.
- . 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." *J. Law and Econ.* 49 (April): 29–62.
- Shapiro, Carl, and Hal R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Bus. School Press.
- Takeyama, Lisa N. 1994. "The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities." *J. Indus. Econ.* 42 (June): 155–66.
- . 1997. "The Intertemporal Consequences of Unauthorized Reproduction of Intellectual Property." *J. Law and Econ.* 40 (October): 511–22.
- U.S. House. 1982. *Home Recording of Copyrighted Works: Hearings before the Subcommittee on Courts, Civil Liberties, and the Administration of Justice of the Committee on the Judiciary*. 97th Cong., 2nd sess. <http://cryptome.org/hrcw-hear-htm>.
- Varian, Hal R. 2000. "Buying, Sharing and Renting Information Goods." *J. Indus. Econ.* 48 (December): 473–88.
- Windmeijer, Frank. 2000. "A Finite Sample Correction for the Variance of Linear Two-Step GMM Estimators." Working Paper no. W00/19, Inst. Fiscal Studies, London.
- Zentner, Alejandro. 2006. "Measuring the Effect of File Sharing on Music Purchases." *J. Law and Econ.* 49 (April): 63–90.