Supply Responses to Digital Distribution: Recorded Music and Live Performances*

Julie Holland Mortimer
Harvard University and NBER

Alan Sorensen
Stanford University and NBER

February 2007

Abstract

Technologies for reproducing and redistributing digital goods have made it more difficult to earn profits from their sale, leading to concerns that socially valuable digital products may not be brought to market. However, digital goods are often jointly supplied with non-digital products, so changes in their distribution technology affect not only the market for the digital product, but also the pricing and profitability of the non-digital good. In this paper we examine the impact of digital file-sharing (which by most accounts decreased sales of recorded music) on the market for live performances (a complementary, non-digital good). Using detailed data on weekly CD sales and concert performances for nearly 2,000 musical artists over a ten-year period, we show that while sales of recorded music declined after the introduction of file-sharing, concert revenues and the number of artists performing concerts increased dramatically. Overall, the patterns in the data suggest that while file-sharing may have eroded profits from CD sales, it also increased the profitability of live performances. Our numbers suggest that file-sharing actually *increased* the revenues from recorded and live music going to artists at typical royalty rates.

Preliminary and Incomplete: Please do not cite.

*The data for this study were provided by Pollstar and Nielsen SoundScan, and we thank Chris Muratore, Gary Bongiovanni, and Alan Krueger for their help in assembling the data. Elias Bruegmann, Natalie Chun, Yani Ioannou, Anna Levine, Maisy Samuelson, Matt Schefer, and Hassan Sultan provided outstanding research assistance. We are thankful to Tim Bresnahan and Austan Goolsbee for helpful comments. Any errors are ours.

1 Introduction

Distributing information goods (i.e., products that can be digitalized) has become an increasingly complex task in recent years. New information technologies have increased the variety of distribution channels available to consumers, but these same technologies have also raised the risk of illegitimate redistribution. Understanding how firms can create and distribute information goods while still protecting their intellectual property is the core issue of many current policy debates, including the recent passage of the Digital Millenium Copyright Act (DMCA), the proposed Uniform Computer Information Transactions Act (UCITA), and a recent case before the U.S. Supreme Court (*MGM v. Grokster*). The fundamental economic concern is that redistribution technologies may threaten markets for information goods by making it difficult for producers to capture the returns to their investments.

However, the debates about copyright protection for information goods have tended to overlook (or at least underemphasize) the simple fact that these goods typically have many different means of consumption. For example, recorded music can be downloaded easily from the internet (through legitimate means or not), but the experience of attending a concert cannot be downloaded. Similarly, while a movie's content has been easy to copy since the 1980's (through video tape), neither the theater experience nor movie-related merchandise can be easily duplicated for redistribution. Moreover, changes in the distribution technology for the digital use of a creative product will affect firms' pricing and supply decisions on non-digital uses because many non-digital uses are complementary in consumption (e.g., a concert may be more enjoyable if you first listen to the recorded song ahead of time, and children may be more interested in toys first featured in a movie).

In general, concerns about the viability of markets for digitally redistributable products may be tempered somewhat by the ability of sellers to recover their investments through the sale of complementary, non-digital goods. Losses due to illegal redistribution of a digital good may be offset by increases in demand for complementary non-digital goods. The implication, as argued by Teece (1986), is that public policy aimed at promoting innovation should not ignore the impact of an innovation on goods or assets that are complementary to it.

In this paper we study firms' responses to digital redistribution technologies in the specific context of the music industry, which has been at the forefront of recent debates about the impact of digital distribution, and has been the focus of several recent empirical studies.[1] Our goal is to examine artists' and record companies' responses to file-sharing, including responses in the market for live performances. If recorded music and live performances are complements, then increases in the consumption of recorded music due to file-sharing should lead to increased demand for live performances. To examine this hypothesis, we analyze a new and detailed dataset covering sales of both recorded music and live performances for 1,806 artists. The data span 10 years (from 1993 to 2002) and include all popular music concerts performed in North America during this period, as well as weekly CD sales from 100 cities, for each artist. The detail provided in the data is very rich: for each concert (ranging from small jazz clubs to stadium tours of international rock stars), we observe revenues, ticket quantity, high and low ticket prices, the identities of all performing bands, and the place and time of the concert. The data on CD sales provide the band and album name, and the quantity of each album sold, by week, in 100 geographic markets in the U.S.

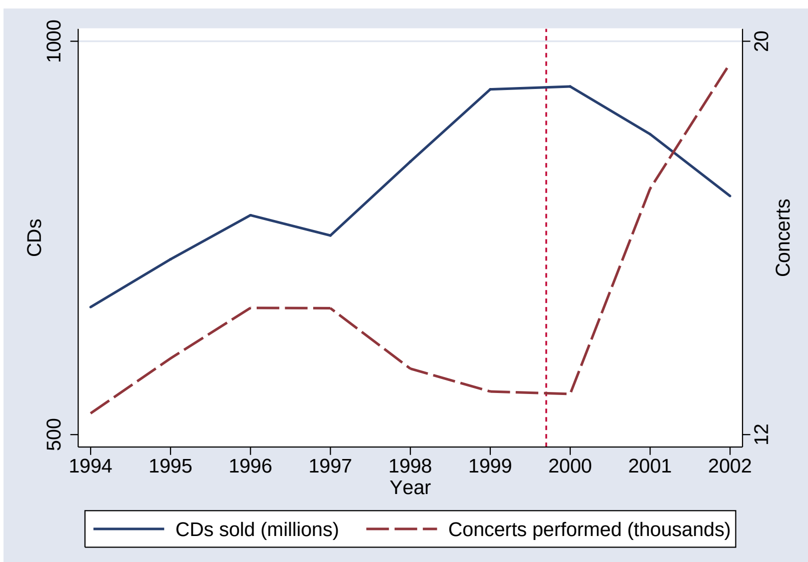

Figure 1: Album Sales and Concerts, 1994-2002

We use our data to describe several changes in music industry trends that occurred following the advent of file-sharing in 1999. The most basic fact, shown in Figure 1, is that the number of concert events increased dramatically in 2001 and 2002.[2] This surge in concert activity is a sharp contrast to the concomitant decline in album sales, which the record industry

[1]See, for example, Blackburn (2004), Hong (2005), Liebowitz (2004), Oberholzer and Strumpf (2004), Rob and Waldfogel (2004), and Zentner (2003). Using various approaches, these studies have all attempted to measure the extent to which downloads of recorded music displace CD sales.

[2]The vertical line marks the entry of Napster.

blamed on file-sharing. Examining the increase in concert activity more closely, we find that the size composition and geographic distribution of concerts also changed sharply in 2001 and 2002: there was a shift toward smaller concerts in smaller venues, and a slight shift toward performing in larger cities.

Our results allow us to say what the *net* impact of these technology changes has been for artists: i.e., overall, has file-sharing made artists better or worse off? Figure 2 shows time trends for total CD revenues, and the sum of CD and concert revenues for the bands in our sample. [3] Figure 2(a), shows total revenues from each of these sources for the industry as a whole, while figure 2(b) provides an estimate of revenues for a typical artist using the benchmark royalty rates discussed earlier. For artists, the decline in revenues from recorded music after 1998 is striking, but appears to have been more than offset by a concomitant increase in concert revenues. Total industry revenues, on the other hand, have not fully recovered, despite the increasing contribution of concert revenue to the total.

Figure 2: Industry Revenue and Artist Revenue

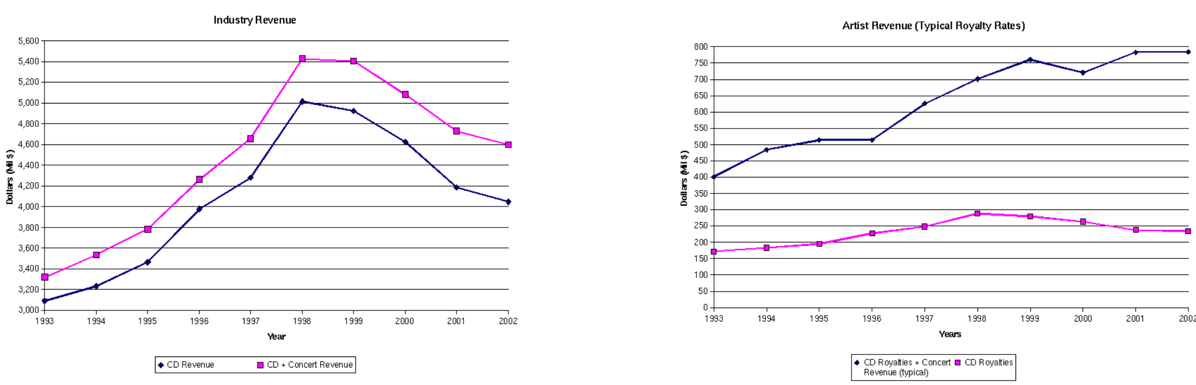

(a) Industry Revenue

(b) Typical Artist Revenue

Based on these facts, we build an empirical model that incorporates demand for recorded music, demand for live performances, and artists' supply decisions jointly in one system. We adopt a structural approach primarily so we can capture the interrelated effects of file-sharing

[3]CD revenues were obtained by multiplying total sales by average CD prices, with the average CD price series coming from Billboard magazine.

on album sales and concert demand. The structural approach allows us to leverage variation in both the live and recorded markets simultaneously in order to capture the overall effect of filesharing. Furthermore, estimates of the model enable a quantitative assessment of how file-sharing affected artists' and record labels' payoffs.

Because file-sharing was obviously not the only change that hit the music market during our sample period, we consider various alternative explanations for the observed increases in concert activity. One possibility is a shift in the costs of producing concerts. Another is the impact of consolidation in the concert promotion industry (i.e., Clear Channel). Another possibility is a shift in demand that resulted from demographic trends that increased the size and/or wealth of the concert-going population. We argue that none of these explanations by itself can explain all the patterns we see in the data. However, to check the robustness of our findings, we (plan to) incorporate assumptions such as changing costs or changing demand elasticities into our structural model to see if we would reach substantially different conclusions.

While our study focuses on the music industry, the economic phenomena we analyze are clearly relevant in many other markets. For example, digital copies of movies may cut into home video sales, but may also lead to higher demand for movie-related merchandise. An author's royalties from book sales may be reduced if the book is digitally shared, but the increased readership may lead to profits on the lecture circuit. Mass sharing of a pirated software program may displace paid licenses for that program, but may also generate increased sales of complementary physical products or technical support services.

The paper proceeds as follows. In the next section we provide a brief description of the music industry, focusing especially on the aspects most relevant to our analysis. Sections 2 and 3 describe the industry and several broad patterns in the data. In Section 4 we develop our empirical model and present our estimation method. Section 5 (will eventually) show results, and section 6 discusses alternative explanations. Section 7 concludes.

2 Music Industry Background

Professional music artists earn revenues principally from recorded music sales and from live performances.[4] Recorded music is produced under contract with a record label: the artist records an album as a work-for-hire, and the record label markets and distributes the album. Typical production costs are in the neighborhood of $100,000-$250,000, and industry executives report that marketing and distribution costs can easily eclipse the cost of production. The standard contract is a royalty contract: the artist is paid royalties on album sales, and receives an advance against those royalties in order to cover living expenses and studio costs during the production of the album. Royalty rates range between 10-18% of retail, with the typical rate being 12%; however, artists earn somewhat less than this due to various deductions that are usually built in to the contract. A reasonable estimate is that the artist earns around $1.00 for every CD she sells.[5] Somewhat surprisingly, record labels have traditionally held a negligible stake in the live performance business. Although labels usually offer some nominal tour support to new artists as part of the recording contract, and sometimes coordinate with concert promoters to advertise a show, they do not take a share of the touring revenues.[6]

Artists' live performances are coordinated and underwritten by concert promoters. The promoter finances almost every aspect of the concert production, including renting the venue, paying the artist and staff, and advertising. Artists are paid as a percentage of ticket revenues, subject to some minimum (called the "guarantee"). Artists also make money from merchandise sales; for some artists this can be a significant component of the net earnings. A typical deal gives 70-80% of merchandise revenues and 70-85% of the gross ticket revenues to the artist, although the actual percentages may be somewhat lower because

[4] Some very successful songwriters also earn significant revenues from music publishing fees, and some star artists have substantial income from endorsements, but the typical artist relies mostly on recorded music sales and concerts.

[5] Instead of using a standard royalty contract, some artists negotiate "penny contracts" specifying artist payments as a fixed dollar amount per CD sold. The typical artist share in these contracts is reportedly $1.25 per CD; however, artists who negotiate such contracts have more bargaining power than the average artist.

[6] Tour support is typically a recoupable expense, but it is recouped from recorded music revenues. This convention may be a holdover from an earlier era: historically, labels subsidized concert tours only as a way of promoting albums, and concerts were often not expected to be profitable on their own.

various deductions are made to the gross ticket revenues before the artist's cut is taken.

Although artists have virtually no say in the pricing of recorded music, most industry sources identify the artist as the primary agent with responsibility for setting concert ticket prices. The artist and/or artist's manager sets prices in consultation with the promoter and venue owner. The parties can have conflicting incentives; for example, aside from the rental fee for the venue, the venue owner's revenue comes primarily from concessions and parking, so they tend to push for low ticket prices in order to fill the house.

In May of 1999, the software program Napster introduced an easy-to-use interface by which consumers could share and download digital copies of songs. Napster and similar programs represented a dramatic shift in the distribution technology for recorded music.[7] Napster gained currency quickly, with a reported user base of over 20 million unique accounts at its peak and over a half million unique IP addresses connected at any given time on a routine basis.[8] The Recording Industry Association of America (RIAA) claimed that the presence of Napster eroded sales of CDs by facilitating copyright violations, and sued to have Napster dismantled in December of 1999. In 2003, the RIAA began suing individual participants of file sharing networks, and subsequent activity on these networks was reported to have declined.[9]

[7] Although earlier technologies also allowed for illegitimate reproduction (e.g. cassette tapes are easily copied), they were much more limited in scope, and typically had greater quality degradation.

[8] Original source: CNNMoney 2000. For an excellent review of the industry and the timing of filesharing events specifically, see Blackburn (2004).

[9] The Supreme Court ruling in *MGM v. Grokster* in June 2005 represented a significant legal victory for the RIAA, as the court held that distributors of file-sharing software could be held secondarily liable for copyright infringements facilitated by their software, essentially allowing the RIAA to go beyond merely suing individuals who shared files illegally to suing the companies whose software enables the sharing. In addition to the legal front, the music industry has also battled file-sharing on the technological front, using various encryption and digital rights management technologies to curb the flow of illegal music downloads. Park and Scotchmer (2004) analyze the impact of such technologies on the pricing of digital goods. Legal channels of digital music distribution are, by now, becoming well established. Most notably, iTunes debuted in October, 2003.

3 Data

Our dataset combines detailed information on album sales and concert events. The data on concerts come from Pollstar, a company that tracks virtually all concert activity in the United States. Our data cover 164,023 concert events performed by over 9,400 artists in the period 1991-2002.[10] For each concert event, we observe the number of tickets sold and the total ticket revenue; high, low, and average ticket prices; the identities of all performing bands; and the place (venue name and capacity) and time of each performance. In our analysis we focus on headlining bands, not supporting bands, and we ignore events where the number of tickets sold was fewer than 50.

Our data on album sales comes from Nielsen SoundScan. These data contain weekly CD sales for several thousand artists, covering the years 1993-2002. SoundScan provides the band and album name, and the quantity of each album sold, by week, in 100 Designated Market Areas (DMAs) in the U.S.[11] In the analyses that follow, we summarize over the individual albums for a particular artist and report total album sales for each artist/DMA/week observation.

Merging the two data sources yields a matched sample of 1,806 artists for whom we have full information on both concerts and album sales.[12] Clearly, this matched sample may not be perfectly representative of the broader universe of artists. Comparisons of our matched sample to the universe of artists on tour are easy to do, because Pollstar tracks nearly the entire universe of touring artists. All the important patterns of concert activity look nearly identical in the entire Pollstar sample compared to the matched sample. Unfortunately, we cannot make this kind of explicit comparison to rule out biases in the selection of artists for whom we have SoundScan data. We only observe CD sales for the artists in our sample, not for the universe of relevant artists. However, our sample is broadly representative of album sales in the sense that it covers a wide range of artist success (ranging from relative

[10] We are also attempting to get concert data for the post-2002 years. CD data are not available after 2002.

[11] A DMA is similar to an MSA.

[12] There are 2,135 artists who appear in both the Pollstar and SoundScan datasets. However, due to the nature of our data collection process, a few artists are missing CD sales over a substantial period of time, or are missing sales of important albums. Comparing CD sales to the RIAA's data on awards, as well as to discographies collected from online databases and to time variation within the SoundScan data allows us to identify and discard artists with problematic data. This process leaves us with a sample of 1,806 artists with verified data integrity.

unknowns to major superstars), and we have no reason to believe CD sales patterns in our sample differ significantly from sales patterns in the broader population of artists.

3.1 Sample selection

One of our primary objectives is to describe changes in concert activity that may have resulted from the advent of file-sharing. In particular, we want to know whether artists were more or less likely to perform concerts after file-sharing became widespread, and whether these effects were more pronounced for some kinds of artists (or in some types of markets) than others. Answering this question is complicated by our inability to observe the entire universe of potential touring artists. Many artists sell albums but do not perform concerts, but we do not know exactly how many such artists exist in any given year, because our album sales data cover only a subset of artists.

More specifically, in trying to calculate the fraction of artists who perform concerts, we face two kinds of selection problems. First, we do not know the rate at which new artists are being "born." So, for example, we cannot immediately rule out the possibility that the increase in concerts shown in Figure 1 merely reflects a sharp influx of new artists in 2001 and 2002.[13] However, in our SoundScan sample, which was constructed specifically to capture artists whose careers began during the sample period, the birth rate of new artists is in fact slightly *lower* in 2000-2002 than in 1997-1999.

The second selection problem is that we do not know the rate at which existing artists "die." If a band is not observed releasing an album or performing a concert for four consecutive years, for example, we don't know if the band is "alive" but choosing not to record or tour, or "dead" (in the literal sense, or in the sense of having retired or broken up). In the analysis below, we address this problem by employing a simple rule: an artist is treated as "dead" if she doesn't release an album or perform a concert for three consecutive years. We then test the robustness of our conclusions against alternative rules (e.g., a 4-year rule instead of a 3-year rule).

[13] In a previous version of this paper we used Amazon.com's music catalog to get a rough count of new artists and new albums by year. There appeared to be a gradual increase in the number of artists over the sample period, but the numbers seemed unreliable, in particular because there was an implausibly large spike in the year 2000.

3.2 Basic patterns

The RIAA reported declining sales beginning in mid-2000, shortly after file-sharing became widespread, but there is some debate about how much of the decline can be attributed to file-sharing.[14] Regardless of its impact on sales, however, file-sharing undoubtedly increased the overall consumption of recorded music. The direct effect of file-sharing was to make millions of songs freely downloadable over the internet. Evidence from time-use surveys suggests that the consequent increase in music listening was dramatic. In one such survey, respondents in 2001 reported spending 3 times as much time listening to music as respondents from 1998. More tellingly, among respondents who reported having below-median internet usage rates, the increase in music listening was negligible (just over 10%), whereas the increase for those with above-median internet usage was more then tenfold.[15]

Naturally, we expect changes in the consumption of recorded music to affect demand for live music. Table 1 presents various patterns in concert activity over time. As shown in Figure 1 above, the number of concerts performed increased sharply in 2001 and 2002. Total tickets and total revenues also increased sharply. The trend in average prices was more variable, but there appears to have been an upward trend in prices from 1999 on. Interestingly, the number of unique artists performing concerts increased sharply in 2001, suggesting that the increase in concerts was partly due to performances by artists who previously would not have toured. The increases in concert activity do not appear to be a function of Pollstar's coverage, since the number of venues represented is relatively stable from 1995 on.

Table 2 examines trends in the average concert size, showing that the increases in the number of concerts (and the number of artists performing concerts) in 2001 and 2002 were associated with sharp decreases in average ticket totals and average venue sizes. In other words, most of the "additional" concerts performed in 2001 and 2002 were smaller than average, which is again suggestive of an increase in concert activity by artists who previously would not have toured.

The second panel of the table shows the geographic composition of concert activity. Using

[14] See the papers listed above related to the "displacement question," not to mention numerous reports in the popular press.

[15] Based on internet-accessible data from National Time Diary Studies conducted by the Survey Research Center at the University of Maryland: see *http://www.popcenter.umd.edu/sdaweb/diary9801/Doc/Diar.htm.*

Table 1: Concerts: changes over time

Year	Number of Concerts	Total Tickets (millions)	Total Revenues ($ millions)	Average Price	Number of Artists	Number of Venues
1991	11,092	39.27	962.82	24.52	1,721	1,502
1992	11,137	44.24	1,107.26	25.03	1,774	1,647
1993	11,064	42.16	1,069.26	25.36	1,904	1,666
1994	12,433	46.68	1,415.20	30.31	2,028	1,789
1995	13,551	44.96	1,179.86	26.24	2,126	1,947
1996	14,577	48.01	1,231.43	25.65	2,294	1,883
1997	14,571	51.78	1,472.45	28.44	2,337	1,955
1998	13,342	51.11	1,595.61	31.22	2,359	1,997
1999	12,877	51.33	1,789.36	34.86	2,306	1,837
2000	12,827	50.79	1,786.19	35.17	2,307	1,940
2001	17,003	56.80	1,988.85	35.02	2,929	2,040
2002	19,549	57.26	2,029.03	35.43	3,340	1,937

Based on Pollstar data. Revenues and prices are deflated to 1999 dollars using the CPI.

the Nielsen-assigned DMA ranks as a rough proxy for market size,[16] it appears that the increase in concert activity in 2001 and 2002 was accompanied by a slight trend toward performing in larger cities.

Table 2: Changes in the size and location of concerts

	Concert size			Market size (DMA ranks)			
Year	Average # of Tickets sold	Average Venue Capacity	Capacity Utilization	1-25	26-50	51-75	76-100$^+$
1991	3,540	4,060	.72	.52	.14	.10	.25
1992	3,973	4,192	.74	.49	.13	.11	.27
1993	3,810	3,994	.75	.48	.14	.11	.27
1994	3,755	3,735	.76	.48	.15	.10	.27
1995	3,318	3,720	.76	.49	.15	.09	.27
1996	3,294	3,757	.75	.48	.17	.10	.25
1997	3,553	3,716	.76	.49	.15	.11	.25
1998	3,831	3,903	.76	.49	.13	.11	.27
1999	3,986	4,307	.76	.53	.13	.10	.25
2000	3,959	4,283	.76	.53	.13	.10	.24
2001	3,340	3,635	.74	.54	.13	.11	.22
2002	2,929	3,379	.72	.57	.12	.10	.21

The "Concert size" panel reports averages across all concerts within each year. The capacity utilization column reports the average of (tickets/capacity), not (average tickets/average capacity). The "Market size" panel reports the fraction of concerts held in each group of markets. DMAs are ranked roughly by size, with DMA 1 (New York City) being the largest. The last group (76-100$^+$) includes concerts in cities that did not cleanly match to a SoundScan DMA.

In general, trends in concert activity by artists in our matched subsample (i.e., the artists for whom we have both concert data and album sales data) seem to be similar to those de-

[16] This ranking isn't perfect, but it is reasonably accurate. DMAs 1-3 are New York, Los Angeles, and Chicago, respectively; DMAs 20, 40, 60, and 80 are Phoenix, Oklahoma City, Richmond (VA), and Huntsville (AL).

scribed above. However, sample selection issues make it difficult to make clean comparisons of the matched subsample to the broader set of artists appearing in the Pollstar data. In particular, because the SoundScan sample is slightly biased toward artists who had moderately successful careers (at least during the 1993-2002 period), the artists in this sample will tend to survive longer than the average artist. Combined with the fact that we pick up new artists in each year of the sample period, this means the net birth rate of new artists is likely to be higher in our matched subsample.

Table 3 describes basic trends within the matched subsample, where we correct for artist deaths using a simple three-year rule: an artist is considered "active" if within the past three years she either (a) performed a concert, or (b) had a quarter in which album sales increased more than 50% over the previous quarter. The reason for including criterion (b) is to account for album releases, publicity campaigns, television appearances, etc. that revive (or keep alive) an artist's career. As the table shows, even with this rule the annual net change in artists is always positive. This is the sense in which we think our matched sample might differ from the broader set of artists in the Pollstar data: we suspect the net birth rate is lower overall than within our subsample.

The table shows that the fraction of artists on tour (i.e., the fraction performing at least one concert within the year) was high in 1996, then dropped from 1997-2000, and then increased sharply in 2001 and 2002. The number of concerts per artist (conditional on having at least one) shows a similar pattern. The last three columns of the table give an indication of concerts' relative importance to the artist's bottom line. As mentioned above, typical artist royalty rates translate to roughly \$1.00 of artist income per CD sold. Using this as a benchmark, the last column provides the ratio of total concert revenues to album sales (using the CPI to deflate concert revenues). This ratio is increasing over the sample period from 1.40 in 1993 to 2.31 in 2002, with the most dramatic change coming in 2001. In other words, in 1993, artists' revenues from concerts are estimated to be roughly 40% more than album revenues, whereas after 2000, concert revenues are well over twice as large as album revenues.

We generally expect artists with the strongest album sales to also be the ones earning the most revenues from concerts. Table 4 describes the joint distribution of album sales and concert revenues in 1998 and 2002. Columns represent approximate quintiles of nonnegative concert revenues (plus a column for zero-revenue artists), rows represent deciles of album

Table 3: Matched sample: basic trends

Year	Artists in sample	Active Artists	Fraction on tour	Average # of concerts	Total concert revenue	Total album sales	(Concert revenue)/ (album sales)
1993	704	704		7.14	239.65	171.77	1.40
1994	819	796		8.03	312.17	183.12	1.70
1995	929	895		8.56	329.27	195.04	1.69
1996	998	973	0.57	8.48	295.51	226.12	1.31
1997	1293	1199	0.53	7.93	390.23	241.97	1.61
1998	1418	1363	0.52	6.30	419.40	284.56	1.47
1999	1519	1422	0.51	6.41	470.20	272.34	1.73
2000	1607	1494	0.50	6.21	446.40	254.20	1.76
2001	1677	1542	0.53	7.74	520.22	230.75	2.25
2002	1806	1623	0.59	7.55	519.50	224.91	2.31

Artists are "active" if they performed a concert or had a 50% increase in album sales (from one quarter to the next) sometime in the previous three years. We cannot calculate the fraction of active artists on tour in 1993-1995 because we don't observe sales data prior to 1993, so we cannot calculate the denominator using our three-year "death" rule. The average number of concerts is conditional on having at least one concert. Concert revenues and album sales are in millions; revenues are in constant 1999 dollars. The ratio of concert revenue to album sales is a rough estimate of the ratio of artists' aggregate concert income to aggregate income from album sales, based on the assumption that artists get $1 for every album sold.

sales, and cells report frequencies—i.e., the fraction of active artists with the corresponding combination of concert revenues and album sales. It is immediately clear from the table that levels of album sales in 2002 were lower at each decile than the levels in 1998. For example, in 1998 album sales ranged from 0.41 million to 10.9 million in the highest decile. The same part of the distribution in 2002 ranged from 0.27 million to 8.3 million. In contrast, concert revenues were higher at each decile in 2002 than they were in 1998.

It is also clear from the table that while album sales and concert revenues are positively correlated, the correlation is far from perfect. Obviously, artists with very high album sales may choose not to perform any concerts in a given year. But even when such artists do perform concerts, the resulting revenues are sometimes lower than the revenues of bands with much weaker album sales.

In any given year, the fraction of artists earning no concert revenues is typically large. Table 4 shows that this fraction was smaller in 2002 (41 percent) than in 1998 (48 percent), and the shift came disproportionately from artists with low album sales. In other words, artists were more likely to tour in 2002 than in 1998, and this was especially true of artists with relatively low album sales. It is also interesting to note that the distribution is somewhat more "spread out" in 2002. For example, conditional on having low album sales, an artist is more likely to be in the upper deciles of concert revenues. Conversely, conditional on being a top-performing concert artist, album sales are more likely to be in the lower deciles.

4 Empirical Model

Motivated by the basic evidence outlined above, in this section we outline a structural empirical model in which file-sharing affects the markets for both recorded and live music. The direct effect of file-sharing is to reduce album sales, since some consumers are able to download the music for free. However, file-sharing increases the number of people listening to a given artist's music, thus expanding the market for that artist's concert performances.

In our basic model, we make some fairly restrictive assumptions so that changes in album sales and concert activity are almost forced to be explained by file-sharing. We do this to determine whether file-sharing alone can serve as a sufficient explanation for the broad patterns in the data. We then relax some of the assumptions to accommodate alternative

Table 4: Joint distribution of Album Sales and Concert Revenues

1998

Album Sales	Concert Revenues (thousands)						
(thousands)	0	1-7	7-18	18-53	53-222	222-34300	Total
0-3	6.24	2.27	0.88	0.37	0.22	0.00	9.98
3-5	6.60	1.83	0.95	0.51	0.07	0.00	9.98
5-10	5.94	1.39	1.32	0.88	0.37	0.07	9.98
10-17	5.50	1.54	1.39	1.10	0.29	0.22	10.05
17-25	4.55	1.32	1.47	1.69	0.95	0.00	9.98
25-43	4.18	0.73	1.54	1.39	1.83	0.29	9.98
43-79	4.11	0.66	1.03	1.69	1.76	0.81	10.05
79-153	4.18	0.15	0.88	1.03	1.91	1.83	9.98
153-412	3.74	0.22	0.81	0.88	1.47	2.86	9.98
412-10900	2.57	0.29	0.22	0.95	1.61	4.40	10.05
Total	47.62	10.42	10.49	10.49	10.49	10.49	100.00

2002

Album Sales	Concert Revenues (thousands)						
(thousands)	0	1-8	8-23	23-66	66-234	234-39700	Total
0-2	6.47	1.66	0.92	0.55	0.12	0.25	9.98
2-4	5.42	1.91	1.66	0.49	0.43	0.06	9.98
4-8	5.11	1.85	1.66	0.80	0.49	0.06	9.98
8-14	4.50	1.79	1.79	1.23	0.62	0.12	10.04
14-23	3.88	1.42	1.73	1.91	0.92	0.12	9.98
23-34	3.76	0.86	1.42	1.85	1.54	0.55	9.98
34-58	3.27	1.23	0.86	1.85	1.85	0.99	10.04
58-111	3.64	0.62	0.68	1.66	1.79	1.60	9.98
111-270	2.59	0.31	0.55	0.99	2.34	3.20	9.98
270-8347	2.59	0.06	0.49	0.43	1.66	4.81	10.04
Total	41.22	11.71	11.77	11.77	11.77	11.77	100.00

explanations, to see if they significantly weaken our conclusions about the impact of file-sharing.

The empirical approach is to generate moment conditions corresponding to the four basic outcomes we observe in the data: album sales for a given artist/market/time period, ticket sales for each concert, ticket prices for each concert, and a binary indicator for whether an artist performs a concert in a given market/time period. To clarify the exposition, we use r superscripts to denote recorded music, d superscripts for downloads, and c superscripts for concerts.

4.1 Demand for recorded music

If consumer i in market m obtains an album by artist j in period t, we assume he gets utility

$$u^r_{ijmt} = \delta_{jt} - \alpha^r(1 - \Gamma_{ijmt})p^r_{jt} + \phi C_{jmt} + \xi^r_{jmt} + \epsilon^r_{ijmt} \ ,$$

where δ_{jt} is an artist-quarter fixed effect (intended to capture fluctuations in the artist's national popularity, due to album releases, TV appearances, etc.), and ξ^r_{jmt} is a mean-zero error term intended to represent unobserved market-specific shocks to the artist's popularity. Γ_{ijmt} is a binary random variable equal to one if consumer i was able to download the recorded music instead of buy it, and p^r_{jt} is the purchase price, which we assume is constant across markets but potentially variable across artists and over time. C_{jmt} is an indicator equal to one if artist j performed a concert in market m in period t. The parameter ϕ measures the extent to which concerts increase demand for the artist's albums; importantly, we are assuming the strength of this complementarity effect is constant over time. Also, we are implicitly assuming there is no quality degradation from downloading instead of purchasing—i.e., downloaded music is not inferior to purchased music.

We parameterize the probability of download as a function of covariates:

$$P(\Gamma_{ijmt} = 1) = 1 - e^{-x_{jmt}{}'\beta} \ ,$$

where x_{jmt} includes observable factors like the artist's age and genre, the broadband penetration in market m, and of course controls for time t to capture the impact of file-sharing beginning in 1999. For notational convenience, call this probability γ_{jmt}.

We assume that the ϵ^r_{ijmt}, the consumer's idiosyncratic preference shock, is distributed exponentially with unit mean, so that expected sales (Y^r) of artist j in market m and quarter t are

$$\begin{aligned} Y^r_{jmt} &= N_{mt} \cdot (1-\gamma_{jmt}) \cdot [1 - F(-\delta_{jt} + \alpha^r p^r_{jt} - \phi C_{jmt} - \xi^r_{jmt})] \\ &= N_{mt} \cdot (1-\gamma_{jmt}) \cdot e^{\delta_{jt} - \alpha^r p^r_{jt} + \phi C_{jmt} + \xi^r_{jmt}} \end{aligned}$$

where N_{mt} is the number of consumers in market m in period t.

Our assumption about the distribution of ϵ^r_{ijmt} is motivated largely by computational convenience. Although the inclusion of the δ_{jt} terms is necessary to account for the dramatic variation in sales we observe in the data (both across artists, and over time for a given artist), estimating the δ_{jt}'s is impractical given the size of our sample. Assuming the distribution of ϵ^r_{ijmt} is exponential allows us to derive an expression in which the fixed effects enter linearly, and from which they can be differenced out. Log sales (y^r) are given by

$$y^r_{jmt} = log(N_{mt}) + log(1-\gamma_{jmt}) + \delta_{jt} - \alpha^r p^r_{jt} + \phi C_{jmt} + \xi^r_{jmt} \ ,$$

so the difference between y^r_{jmt} and $\bar{y}^r_{j.t}$ (the average of log sales for artist j across markets in period t) is

$$y^r_{jmt} - \bar{y}^r_{j.t} = \left[log(N_{mt}) - \overline{log(N_{.t})}\right] + \left[log(1-\gamma_{jmt}) - \overline{log(1-\gamma_{j.t})}\right] + \phi(C_{jmt} - \bar{C}_{j.t}) + \xi^r_{jmt} - \bar{\xi}^r_{j.t} \ . \tag{1}$$

We assume the number of consumers N_{mt} is simply $a_{mt}N_t$, where N_t is the total number of consumers across all markets. In other words, we think of market m's size as some fraction of the national market. Then the expression

$$\left[log(N_{mt}) - \overline{log(N_{.t})}\right]$$

becomes

$$log(a_{mt}) + log(N_t) - \left(\frac{1}{M}\sum_m (log(a_{mt}) + log(N_t))\right) = log(a_{mt}) - \frac{1}{M}\sum_m log(a_{mt})$$

In the estimation we compute a_{mt} as market m's share of overall sales across artists in period t.

Substituting for γ_{jmt}, we can rewrite equation (1) as

$$y^r_{jmt} - \bar{y}^r_{j.t} - (log(a_{mt}) - \bar{a}_{.t}) = -(x_{jmt} - \bar{x}_{j.t})'\beta + \phi(C_{jmt} - \bar{C}_{j.t}) + \xi^r_{jmt} - \bar{\xi}^r_{j.t}\ . \qquad (2)$$

Because we are averaging over nearly 100 markets, we can safely ignore the $\bar{\xi}^r_{j.t}$: the law of large numbers says it will be effectively zero.

The main virtue of our differencing strategy is that it sweeps out the fixed effects δ_{jt}. Of course, it also sweeps out any other variables that are constant across markets for a given artist. This includes prices (p^r_{jt}), which is convenient because we do not have useful data on prices anyway. But we also lose any covariates in the parameterization of γ_{jmt} that are constant across markets. In other words, using the demand for recorded music alone we would not be able to identify variation in γ that is purely a function of time (e.g., pre- vs. post-Napster) or purely a function of artist characteristics (e.g., genre or artist age). Intuitively, the reason is that we are allowing each artist's popularity to change so flexibly over time (through the δ_{jt}'s). So, for example, if sales fall in general over time, our model would attribute this to a decline in album quality or artist popularity (δ_{jt}), not a change in γ. To identify anything about how the general level of file-sharing changed over time or how it varied across artists, we need to incorporate moment conditions from the concert side of the market.

We do not have data on the extent of downloading activity. However, our model implies that downloads are proportional to sales (which we do observe). A consumer will download if $\Gamma_{ijmt} = 1$ and $(\delta_{jt} + \phi C_{jmt} + \xi^r_{jmt} + \epsilon^r_{ijmt}) > 0$. Expected downloads ($Y^d$) are therefore given by

$$
\begin{aligned}
Y_{jmt}^{d} &= N_{mt} \cdot P(\Gamma_{ijmt} = 1) \cdot P(\epsilon_{ijmt}^{r} > -\delta_{jt} - \phi C_{jmt} - \xi_{jmt}^{r}) \\
&= N_{mt} \cdot \gamma_{jmt} \cdot e^{\delta_{jt} + \phi C_{jmt} + \xi_{jmt}^{r}} \\
&= Y_{jmt}^{r} \cdot \left(\frac{\gamma_{jmt}}{1 - \gamma_{jmt}} \right) \cdot e^{\alpha^{r} p_{jt}^{r}} \ .
\end{aligned}
$$

With the parameterization $\gamma_{jmt} = 1 - e^{-x_{jmt}{}'\beta}$, this can be rewritten as

$$
Y_{jmt}^{d} = Y_{jmt}^{r} \cdot \left(e^{x_{jmt}{}'\beta} - 1 \right) \cdot e^{\alpha^{r} p_{jt}^{r}}
$$

This means we can write the combined number of purchases and downloads as a function of observed sales and parameters:

$$
\begin{aligned}
Y_{jmt}^{r} + Y_{jmt}^{d} &= Y_{jmt}^{r} \left[1 + \left(\frac{\gamma_{jmt}}{1 - \gamma_{jmt}} \right) \cdot e^{\alpha^{r} p_{jt}^{r}} \right] \\
&= Y_{jmt}^{r} \cdot e^{x_{jmt}{}'\beta} \cdot e^{\alpha^{r} p_{jt}^{r}}
\end{aligned}
$$

This expression captures the idea that if downloading is prevalent (i.e., γ is large), then consumption of the recorded good is some potentially large multiple of sales. Even if only a small number of consumers actually purchase an artist's album, there may be a large number of consumers listening to that album.

4.2 Demand for live music

We write the consumer's utility from attending a concert by artist j in market m and period t as

$$
u_{ijmt}^{c} = \omega_{j} - \alpha_{j}^{c} p_{jmt}^{c} + \xi_{jmt}^{c} + \epsilon_{ijmt}^{c} \ ,
$$

where ω_j is an artist-specific quality term, p^c_{jmt} is the concert ticket price, ξ^c_{jmt} is an unobserved demand shock that is common across consumers, and ϵ^c_{ijmt} is a consumer-specific preference shock. We again assume that ϵ^c_{ijmt} is exponentially distributed. For a consumer in the market for the concert, the probability of attending the concert is therefore

$$Pr(\epsilon^r_{ijmt} > -\omega_j + \alpha^c_j p^c_{jmt}) = e^{\omega_j - \alpha^c_j p_{jmt} + \xi^c_{jmt}} \ .$$

The key assumption we make is that the size of the market for an artist's concert is a function of that artist's album sales. Specifically, we assume the market size is proportional to cumulative sales and downloads over the past s periods, which we write as

$$L_{jmt} \equiv \sum_{t-s}^{t} (Y^r_{jmt} + Y^d_{jmt}) = \sum_{t-s}^{t} Y^r_{jmt} \cdot e^{x_{jmt}'\beta} \cdot e^{\alpha^r p^r_{jt}} \ . \tag{3}$$

Expected concert ticket sales are then

$$Y^c_{jmt} = L_{jmt}(x_{jmt}, p^r_{jt}; \beta, \alpha_r) e^{\omega_j - \alpha^c_j p_{jmt} + \xi^c_{jmt}} \ ,$$

and taking logs we get

$$y^c_{jmt} = log(L_{jmt}(x_{jmt}, p^r_{jt}; \beta, \alpha_r)) + \omega_j - \alpha^c_j p_{jmt} + \xi^c_{jmt} \ . \tag{4}$$

The price coefficient α^c_j varies across artists (i.e., it is equivalent to interacting price with artist fixed effects).

4.3 Supply of live music

On the supply side of the market for live music, we model the two primary choices artists make: whether to perform a concert in a given market and time period, and what price to set (conditional on performing a concert).

If an artist chooses to perform a concert in market m, period t, we assume she sets the revenue-maximizing monopoly price for that concert. Given the functional form for concert demand, this leads to the simple pricing rule

$$p^*_{jmt} = \frac{1}{\alpha^c_j} \ . \tag{5}$$

Notice that we are making three important assumptions: (1) marginal costs are zero; (2) competition from other concerts is irrelevant; and (3) dynamic concerns can be ignored in the pricing decision. These are strong assumptions, but they are not obviously bad ones. The zero marginal cost assumption is approximately true given that most of the concerts in our sample were not sold out.[17] Competition from other concerts is probably of second-order importance, because concert offerings are differentiated musically and chronologically. (It is rare to see two concerts by musically similar artists on the same weekend in the same city, for example.) Ignoring dynamic considerations may be less defensible, since new artists may have strong incentives to underprice their concerts in order to build a larger fan base for the future. We set aside dynamics in the present analysis primarily for the sake of tractability, but we intend to explore the issue further.

Note that the monopoly pricing rule implies that if the price coefficient is constant across markets for an artist, price won't vary across markets for that artist (i.e., $p^*_{jmt} = p^*_j = 1/\alpha^c_j$). In reality, this is often the case, although it need not be. There are a number of modeling choices that one can make when actually estimating the α^c_j parameter, which we discuss in the estimation section.

The decision to perform a concert depends on whether the expected revenues that an artist receives from that concert exceed the artist's unobserved fixed cost of performing. Assuming that an artist receives an amount equal to gross ticket revenues, and that prices are set as above, we expect to observe a concert by artist j in market m in period t if

$$p^*_{jmt} Y^c_{jmt} > \kappa_{jmt} \ ,$$

or

[17]Our estimation ignores capacity constraints entirely. To the extent that artists can choose from among many different venue sizes in a given market, this assumption is not too restrictive. Obviously, however, the choice of venues will be limited in some markets, so that capacity constraints are sometimes binding. I most cases, artists can relax this constraint by playing for multiple nights.

$$\frac{1}{\alpha_j^c}\left(L_{jmt}e^{\omega_j-1+\xi_{jmt}^c}\right) > \kappa_{jmt} \ ,$$

where κ_{jmt} is the artist's unobserved fixed cost.[18] This cost could represent the literal fixed costs of producing a concert, but also represents the artist's opportunity cost of time: presumably in some time periods an artist will choose to work on recording a new album instead of going on tour. We model the fixed costs as

$$\kappa_{jmt} = e^{\bar{\kappa}+\Delta\kappa_{jmt}} \ ,$$

with $\Delta\kappa_{jmt}$ a mean-zero error term. In other words, we estimate only the mean level of fixed costs, and assume actual fixed costs are distributed around that mean identically across artists, markets, and time periods. (We consider relaxations of this assumption in the robustness checks below.)

Taking logs and rearranging terms, the above condition becomes

$$log(L_{jmt}) + \omega_j - 1 - log(\alpha_j^c) - \bar{\kappa} > \Delta\kappa_{jmt} - \xi_{jmt}^c \ .$$

Assuming the composite error term $\Delta\kappa_{jmt} - \xi_{jmt}^c$ has some distribution function G, the probability that artist j plays a concert in market m at time t is:

$$E(C_{jmt}) = G\left(log(L_{jmt}) + \omega_j - 1 - log(\alpha_j^c) - \bar{\kappa}\right) \ .$$

In specifying this condition, we have ignored some potentially important factors that could influence the artist's decision of whether to perform a concert. First, in principle the artist should consider the spillover benefits of concerts for album sales: by performing a concert, the artist not only gets the revenues from that concert, but also a share of the revenues from the extra album sales generated by that concert. Given our specification of album demand,

[18] We don't observe the precise contractual terms governing individual artists' concert royalties, nor do we observe ancillary revenue sources such as T-shirt sales. As a result, we follow the "rule of thumb" given by industry participants, which is that an artist's share of ancillary revenues is generally around 20 percent of gross ticket revenues, and an artist's share of gross ticket revenues is 80 percent. Thus, we model artist's decisions as if an artist takes home 100 percent of gross ticket revenues.

the incremental sales are just $(e^{\phi} - 1)$ times the observed sales. However, because most artists don't earn any royalties on the marginal album sale (and even for those who do, the royalty is at most 15%), as a practial matter we think it is reasonable to ignore this aspect of the decision.

Second, the costs of performing a concert are almost certainly not independent across markets in a given time period. For example, conditional on performing in Seattle, it's not as costly to perform in Portland. More importantly, however, the costs are probably increasing in the total number of concerts performed in the time period. Put another way, nothing in our model says that an artist can't perform a concert in *every* market in period t. But if the time period is a quarter (which is what we assume below for purposes of estimation), that is basically impossible. One crude way to control for this is to include the total number of concerts performed by artist j in period t as a covariate in the specification of fixed costs, κ_{jmt}.

Finally, nothing in the model says an artist won't perform concerts in every period for a given market—i.e., if the model says it was profitable for the artist to perform in Boston last period, it will almost certainly be profitable to do so again in this period. This is unrealistic and obviously inconsistent with observed touring patterns. We adjust for this by including controls for past concerts as covariates in κ_{jmt}.

4.4 Estimation details

We use four of the conditions described above to generate a set of moment conditions for estimation. The conditions consist of demand equations for both the CD market and the concert market, as well as two supply decisions in the concert market (whether to perform a concert in a given market and time period, and what price to set). We don't use a supply equation for the CD market for two reasons. First, artists don't set prices in this market. Second, artists are often under contractual schedules by a record label to produce albums, and they receive relatively low royalty rates on CD sales. Furthermore, we view artists' decisions in the concert market as being relative to their baseline effort in the market for recorded music. (If they choose to spend more time touring, this is equivalent to spending less time in the recording studio.)

The parameters of primary interest are the coefficients on x_{jmt}, the covariates affecting the probability of downloading. In the estimation x includes artist characteristics such as age and genre, as well as market characteristics such as broadband penetration. We interact these variables with year dummies (or pre- vs. post-Napster) dummies, since the idea is to see whether the impact of file-sharing was larger for some kinds of artists or in some kinds of markets.

We construct the sample so that a time period (t) is a quarter—i.e., we aggregate the weekly data up to a quarterly frequency. We use equation (2), describing the demand for CDs, to construct the first moment condition. For a set of instruments Z^r, we assume that $E(Z^{r\prime}\xi^r_{jmt}) = 0$, with

$$\xi^r_{jmt} = (y^r_{jmt} - \bar{y}^r_{j.t}) - (log(\lambda_{mt}) - \bar{\lambda}_{.t}) + (x_{jmt} - \bar{x}_{j.t})'\beta - \phi(C_{jmt} - \bar{C}_{j.t}) \ .$$

A concert event is assumed to affect sales for 6 weeks before and after a concert in addition to the week of the concert. Thus, a concert event that takes place during the middle week of a quarter is assigned as a full concert in that quarter ($C_{jmt} = 1$), while a concert event that takes place in the last week of a quarter is assigned fractionally across two quarters ($C_{jmt} = 7/13$ and $C_{jm(t+1)} = 6/13$).

However, the concert variable is not an appropriate instrument, because it is likely to be correlated with ξ^r_{jmt}. The potential endogeneity problem is seen most clearly if we decompose the unobserved shock into two components: $\xi^r_{jm} + \Delta\xi^r_{jmt}$. In other words, any given artist may be disproportionately popular in a given market, and this popularity will persist over time (hence the ξ^r_{jm}). While it may be reasonable to assume that C_{jmt} is orthogonal to $\Delta\xi^r_{jmt}$ (artists do not strategically schedule their concerts to coincide with transitory positive shocks to market-specific popularity), C_{jmt} is almost certainly correlated with ξ^r_{jm}, because artists will obviously be more likely to perform concerts in cities where they have a consistently strong fan base.

Recall that the differencing strategy in equation (2) causes us to lose any covariates in the parameterization of γ_{jmt} that are constant across markets. Thus, the only parameters we identify from the CD sales equation are the β's associated with market-varying x's. If we had a good instrument for C_{jmt}, we could also identify ϕ from this equation.

Turning to the concert market, the demand equation 4 allows us to form moment conditions based on a set of instruments Z^c, where

$$E(Z^{c\prime}\xi^c_{jmt}) = E(Z^{c\prime}(y^c_{jmt} - log(L_{jmt}) - \omega_j - \alpha^c_j p_{jmt})) = 0 \ . \tag{6}$$

We assume that the market for a concert (L_{jmt}) consists of everyone who consumed a download or purchased of the recorded good in the last three years (so that the summation in equation 3 is over 12 quarters).[19]

Note the large number of parameters involved in this moment condition. We need to identify the β's for all X_{jmt}'s that do not vary across markets, α^r, and the set of ω_j's and α^c_j's. We assume that the x's in L_{jmt} are exogenous, and these are used as instruments to identify the β's. The α^r parameter, which appears in the expression for L_{jmt}, is normalized to zero. The fact that we have no data on CD prices essentially forces us to make this normalization. However, we expect it to have very little impact: prices are in fact quite stable over time and across artists, so the $e^{\alpha^r p_{jt}}$ term in L_{jmt} is effectively just a scalar multiple, and will be absorbed into the ω_j's. The ω_j's are estimated using a two-step procedure described below.

Obviously, ticket prices p^c_{jmt} are endogenous, and cannot be included in the instrument set Z^c. Instead, we estimate the α^c_j's directly from the pricing equation, and treat them as data when estimating the other moment conditions. By estimating the α^c_j separately, we are essentially assigning a large weight to the pricing condition in the GMM weighting matrix. One could also include the pricing decision directly in the GMM objective function and estimate using the usual Hansen (1982) weight matrix.

There are a few ways one could estimate the α^c_j's. The most direct application of the pricing rule would imply setting an $\alpha^c_{jmt} = 1/p_{jmt}$ for each concert. However, this is only feasible for markets and time periods in which artist j chose to play a concert. Otherwise, we don't observe a p_{jmt}. For this reason, we consider a price coefficient that varies across artist, but not across markets.[20] One option is to assign α^c_j to be equal to one over the mean of the observed p_{jmt}'s for each artist. The interpretation would presumably be that p^*_j is measured with error across markets and time periods (i.e., $p_{jmt} = 1/\alpha^c_j + \zeta_{jmt}$). The presence of this additional noise would lead to an upward bias on the estimate of σ^2_κ. Alternatively, one could

[19] We will check the robustness of this assumption across different window lengths.

[20] One could also allow α^c_j to vary across artist and calendar year if the artist tours sufficiently often.

estimate the α_j^c's from the artist's smallest concert (based on revenues). This would give an interpretation of α_j^c that follows from a model in which an artist plays concerts each year until the revenues of the smallest concert just equal the fixed cost.

Finally, if one assumes the composite error term $\Delta\kappa_{jmt} - \xi_{jmt}^c$ is normally distributed with variance σ_κ^2, we can get moment conditions of the form

$$E\left(W^{c\,\prime}\left(C_{jmt} - \Phi\left(\frac{1}{\sigma_\kappa^2}(log(L_{jmt}) + \hat{\omega}_j - \hat{\alpha}_j^c p_{jmt} - \log(\hat{\alpha}_j^c) - \bar{\kappa})\right)\right)\right) = 0\ , \tag{7}$$

where W^c is a set of instruments that include the concert demand instruments, Z^c, as in equation (6) above, as well as any additional instruments that could affect the decision to play a concert (such as whether or not a concert was played in the same city very recently).[21] The idea behind the identification of the fixed-cost parameters separately from the β parameters is that changes in the fixed cost of touring can affect the propensity to tour, but do not disproportionately change the predicted revenues of touring as a funciton of previous CD sales (i.e., it won't be able to explain why small artists can do larger concerts in high broadband cities).

The estimation combines these conditions and runs GMM. Define the set of parameters to be estimated as $\theta = (\beta, \phi, \omega_j, \sigma_\kappa^2, \kappa)$. The objective function is:

$$\psi(\theta, Z, \hat{\alpha}_j^c) \equiv \begin{pmatrix} \xi_{jmt}^r x_{jmt} \\ \xi_{jmt}^c Z^c \\ W^{c\,\prime}\left(C_{jmt} - \Phi\left(\frac{1}{\sigma_\kappa^2}(log(L_{jmt}) + \omega_j - 1 - log(\alpha_j^c) - \bar{\kappa})\right)\right) \end{pmatrix} \tag{8}$$

Generalized Method of Moments solves:

$$\hat{\theta} = \text{argmin}\left(\sum_i \psi(\theta, Z_i)\right)' \ A \ \left(\sum_i \psi(\theta, Z_i)\right).$$

where $\psi(\theta, Z_i)$ is the set of moment conditions and A is a weight matrix, chosen as in Hansen (1982).

[21] There are almost certainly more elegant ways to model this decision–we are still working on this condition. There is also the issue that the ω_j's enter non-linearly here, which may affect our procedure for estimating them.

Direct estimation of $\hat{\theta}$ is difficult because of the large number of parameters in ω_j. However, ω_j enters ξ^c linearly, which allows us to simplify the estimation by using a two-step procedure. We first guess values for the ω vector, then given those values, we minimize the objective function (which is non-linear) over values of the other parameters. Given the new values of the non-linear parameters, we construct fitted values of ω_{jmt}'s, which we then regress on artist dummies to recover new estimates of the ω_j vector. We iterate this procedure until convergence.

5 Results and Discussion

5.1 Results

5.2 Discussion

While changes in distribution technology appear to have eroded the profitability of selling recorded albums, our preliminary findings suggest that these changes may have simultaneously boosted demand for live performances. Given an understanding of the changes in the total surplus, we can also explore changes in the sharing of that surplus. For example, record labels have historically claimed the lion's share of revenues from recorded music, while leaving the concert business to the artists. Indeed, artists usually contract with independent promoters to produce their concerts, with little (if any) of the concert revenues reverting to the artist's record label. Not surprisingly, the loudest complaints about the effects of internet file-sharing have come from record labels and their parent distributors. Since concerts capture returns to investments at least partially made by record labels, it seems likely a new equilibrium will emerge in which those labels play a larger role in concert promotion and claim a larger share of concert profits.

An additional outcome of our research will be to uncover and explain any heterogeneity underlying aggregate patterns like those in Figure 2. For the music industry, some of the most interesting unanswered questions concern the differential impact of internet file-sharing across artists. It is quite likely that file-sharing is a boon to some artists and a bane to others, but to date there is little empirical evidence indicating which types of artists gain vs. lose. For instance, digital distribution of recorded music may have made it easier for new or unusual artists to establish a large enough fan base to profitably tour: the dramatic growth in the number of artists performing concerts in 2001 and 2002 suggests this may be true.

6 Alternative Explanations and Robustness Checks

The advent of digital file-sharing was not the only major change that occurred during the sample period, and although the patterns in the data are broadly consistent with the predicted impact of file-sharing, there are at least three alternative explanations worth considering. One possibility is that the supply of live performances was affected by changes in costs. There are two important types of costs in the concert industry: production costs (including salaries and hotel expenses for the crew, transportation of crew and equipment, etc.) and venue costs, such as the rental fee and any revenue-sharing arrangements with the concert hall. Our concert data do not include information on costs, but trade journals indicate a slight increase in production costs over our sample period. This pattern was apparently true more generally: the concert promoters with whom we spoke described a gradual upward trend in their overall costs, but could not identify any specific or dramatic shocks to their cost structure. Note that higher costs may partly explain the price increases during our sample period, but obviously cannot rationalize the expanded supply of live performances.

Another possibility is that the patterns in our data reflect the impact of consolidation in the concert promotion industry. The Telecommunications Act of 1996 relaxed previous restrictions on the number of radio stations a single company could own, and also effectively allowed national radio station owners to enter the promotion business. By the late 1990's, many venues and radio stations had been purchased by Clear Channel Entertainment (CCE). In principle, this consolidation could help explain the observed increase in the supply of concerts, because cross-ownership of radio stations and concert performances may have made marketing investments more efficient, and more importantly because it may have been easier to book national tours through a promoter that owned hundreds of venues nationwide. In reality, however, we doubt that this was the principal cause of the increase in concert activity. Prior to consolidation, artists arranged concert tours through booking agencies, and these agencies obviously already had relationships with numerous venues across the country. Moreover, while the increase in the supply of concerts was most dramatic among new artists, there is no evidence that these artists were disproportionately more likely to use

CCE as their promoter.[22] Similarly, there is no evidence that increases in concert activity were larger in cities where CCE became the dominant promoter.

Finally, it is possible that the shift in demand for live music was not due to file-sharing, but rather due to demographic trends that increased the size and/or wealth of the concert-going population. In particular, one hypothesis is that the increases in concert prices and in the number of concerts overall were driven by "dinosaur bands" like the Rolling Stones cashing in on their older, wealthier, baby-boomer fans. However, our data show no indication that older bands were disproportionately more likely to tour during the period in question. Indeed, the biggest increase in touring activity was among young (new) bands, whose fans were presumably younger—and based on the Census Bureau's national intercensal estimates, the population of 15-34 year-olds actually *shrank* relative to the overall population from 1995-2000. Although we do suspect that demographic trends affected the demand for concerts during our sample period, they cannot alone explain the patterns in our data.

Although these alternative explanations have difficulty rationalizing some features of the data, we use them to motivate several alternative specifications of our structural model. By relaxing some of the model's assumptions, we hope to determine how much of the data variation can potentially be explained by factors other than file-sharing. [...]

7 Conclusion

[22] Our data indicate the primary promoter for each concert, so we estimated probit models predicting promotion by CCE as a function of city dummies, dummies for age groups, and dummies for genres. Concerts by new bands (age 0-5) were no more likely to be promoted by CCE.

References

Blackburn, D., (2004), "On-line Piracy and Recorded Music Sales," *working paper.*

BusinessWeek, February 13, 2003, "Big Music's Broken Record."

Galenson, D., (2004), "A Portrait of the Artist as a Very Young or Very Old Innovator: Creativity at the Extremes of the Life Cycle," *NBER working paper W10515.*

Gentzkow, M., (2004), "Valuing New Goods in a Model with Complementarities: Online Newspapers," *working paper.*

Hong, Seung-Hyun, (2005), "The Effect of Digital Technology on the Sales of Copyrighted Goods," *Stanford University Dissertation.*

Krasilovsky, M. and S. Shemel, (2000), *This Business of Music*, New 8th Edition. Billboard Books, Watson-Guptill Publications.

Krueger, A., (2002), "Music Sales Slump and Rock Fans Pay the Price," *The New York Times: The Economic Scene*, October 17th.

Krueger, A., (2004), "The Economics of Real Superstars: The Market for Rock Concerts in the Material World," *working paper.*

Oberholzer, F. and K. Strumpf, (2004), "The Effect of File Sharing on Record Sale: An Empirical Analysis," *working paper.*

Park, Y. and S. Scotchmer, (2004), "Digital Rights Management and the Pricing of Digital Products," *NET Institute Working Paper #04-09.*

Passman, D., (2000), *All You Need to Know About the Music Business*, Simon & Schuster.

Rob, R. and J. Waldfogel, (2004), "Piracy on the High C's: Music Downloading, Sales Displacement and Social Welfare," *working paper.*

Takeyama, L., (1994), "The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities," *The Journal of Industrial Economics* V. 42, pp 155-66.

Teece, D. J., (1986), "Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy," *Research Policy* v. 15, pp. 285-305.

USA Today, August 6, 2004, "Summer Tours Help Bands Pay Bills."

Varian, H., (1997), "Versioning Information Goods," *working paper*.

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

Zentner, A., (2003), "Measuring the Effect of Online Music Piracy on Music Sales," *working paper*.