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

Julie Holland Mortimer
Harvard University and NBER

Alan Sorensen
Stanford University and NBER

December 29, 2005
(Preliminary: please do not cite)

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 with non-convex production technologies may not be brought to market. However, digital goods are often jointly supplied with non-digital products, and changes in distribution technologies affect not only the market for the digital product, but also the pricing and profitability of the non-digital good. We outline a simple model illustrating these effects in the music industry, and test the model's implications using detailed data on weekly CD sales and individual 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. We examine whether these changes were most pronounced among artists or markets where file-sharing was likely to be most significant. 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.

*The data for this study were provided by SoundScan and Pollstar, and we thank Chris Muratore, Gary Bongiovanni and Alan Krueger for their help in collecting 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 Austan Goolsbee for providing data on broadband penetration, and to Tim Bresnahan 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. Although new information technologies have increased the variety of distribution channels available to consumers, 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 currently-debated 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 or impossible 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 uses and 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 theather experience nor the movie-related merchandise can be easily duplicated for redistribution. Moreover, because many of the non-digital uses of creative products are complementary in consumption with the digital uses (e.g., a concert may be more enjoyable if you listen to the recorded song ahead of time, and children may be more interested in toys and costumes after seeing them in a movie), changes in the distribution technology for the digital use of a product will affect firms' pricing and supply decisions on the non-digital uses. Importantly, losses due to the illegal redistribution of a digital good may be offset to some extent 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 in this paper is to examine artists' and record companies' responses to file-sharing, paying particular attention to the market for live performances (i.e., the complementary, non-digital good). We have collected a detailed dataset covering sales of both recorded music and live performances for 2,135 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 (the headline act as well as any supporting acts), 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 Designated Market Areas in the U.S. (similar to an MSA). The merged dataset contains all album sales and concert activity for every band in each of the 100 markets in the U.S. in every week over ten years.

We outline a simple model to illustrate the predicted effects of file-sharing on the market for recorded music and live performances, and then use our data to test the predictions of the model. Because we do not have detailed data on the prevalence of file-sharing itself, we rely on variation over time and across artist or market characteristics to draw inferences about its impact. For example, we document increases in concert ticket prices, concert revenues, and the number of artists performing concerts that are coincident with the diffusion of file-sharing technologies beginning in 1999. We then ask whether these changes were most dramatic for cohorts of artists where file-sharing activity is believed to have been most significant, or in markets with high levels of broadband penetration. To preview some of the results, we find that the impact of concerts on CD sales diminished after the introduction of file-sharing, and that the decline was largest in markets with high broadband penetration. Concert revenues and the number of bands on tour also increased after the advent of file-sharing, as

[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 answer the "sales displacement question": do music downloads displace music sales, and if so, at what rate?

did the time between album releases for established artists. In general, we find that all the time trends are consistent with the model's predictions; however, results from "differences in differences" tests (i.e., checks of whether the trend changes are most pronounced for artists or in markets where file-sharing is believed to have been most prevalent) are mixed.

While this study focuses specifically on the music industry, we wish to underscore that the economic phenomena we analyze here are relevant in many other markets. For example, 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. Digital copies of movies may cut into home video sales, but may also lead to higher demand for movie-related merchandise. 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.

The paper proceeds as follows. In the next section we describe the music industry and our data sources. Section 3 presents a stylized model of the market for music and outlines predictions regarding the impact of file-sharing. Evidence related to these predictions, along with further empirical analysis of supply responses to file-sharing, is presented in section 4. We consider alternative explanations of the observed data patterns in section 5, and conclude in section 6 with a discussion of the implications of our findings.

2 Industry and Data Description

2.1 Music Industry Background

Professional music artists earn revenues principally from recorded music sales and from live performances.[2] 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

[2]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.

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.[3]

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.[4] 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.

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 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.

[3]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 number for such contracts is $1.25 per CD.

[4]Tour support is typically a recoupable expense, but it is recouped from recorded music revenues.

Digital distribution technologies

Although reproduction technologies had threatened music sales for many years—even cassette tapes were relatively easy to copy—the arrival of digital file-sharing technologies represented a dramatic shift. In May of 1999, a software program (Napster) introduced an easy-to-use interface by which consumers could share and download digital copies of songs. 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.[5] 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.[6] 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.[7] Legal channels of digital music distribution are, by now, becoming well established.[8]

2.2 Data Description

Two organizations collect data on the concert and recorded music markets, respectively: Pollstar and Nielsen SoundScan. Our Pollstar data covers concert activity in the years 1990-2003, and includes information on revenues, ticket quantity, high and low ticket prices, the identities of all performing bands, and the place and time of each concert. In the current analysis, we examine headlining bands, not supporting bands. The SoundScan dataset contains weekly CD sales for several thousand artists, covering the years 1993-2002. SoundScan

[5]Original source: CNNMoney 2000. For an excellent review of the industry and the timing of filesharing events specifically, see Blackburn (2004), on which some of the present discussion is based.

[6]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.

[7]Park and Scotchmer (2004) analyze the impact of such technologies on the pricing of digital goods.

[8]Most notably, iTunes debuted in October, 2003.

provides the band and album name, and the quantity of each album sold, by week, in 100 Designated Market Areas in the US. An observation in the merged dataset contains concert activities and CD sales for every band in 100 U.S. cities in every week from 1993 through the end of 2002. In the analyses that follow, we summarize over the individual albums for a particular artist and report total album sales for each artist/city/week observation. The merged data contain full information on the concerts and CD sales for 2,135 artists. 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 problematic artist records and discard those artists. This process leaves us with a sample of 1,806 artists with verified data integrity. In addition to the data from SoundScan and Pollstar, we also augment each artist's CD sales with information from the RIAA awards database for all artists active prior to 1993. The reason for including these data is to account for historical sales of CDs that debuted prior to the SoundScan tracking service in 1993. While SoundScan provides data at the DMA level, RIAA awards are done on the basis of national sales. Thus, we apportion the pre-1993 sales across cities based on the relative proportion of sales across cities for an artist in 1993.

Clearly, the sample of "matched" observations (ie., artist with both touring and recording activity) may differ from the universe of artists doing one of the two activities. Comparisons of our matched sample to the universe of artists on tour are easy to do, because Pollstar essentially contains data on 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. 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—but our sample is representative in the sense that it covers a wide range of artist success (ranging from relative 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 Theory and Predictions

In this section, we characterize the objective function of artists and provide evidence on the predicted effects of digital distribution on artists' supply decisions. The model naturally represents a very simplified and stylized guide to the industry, and many alternative explanations can be offered to explain the observed trends in artists' supply functions. We attempt to address some of these alternative explanations as we go along. However, the goal of this section is to clarify the potential effects of digital distribution while capturing the most important features of the market.

Consider a simple model in which artists produce two goods: recorded music (good 1) and live music (good 2).[9] We write the artist's profit function as

$$\pi = \phi \left[(p_1 - c_1)Q_1 - K_1 \right] + (p_2 - c_2)Q_2 - K_2 - h(e_1, e_2) \,, \tag{1}$$

where p_j, c_j, and K_j are the price, marginal cost, and fixed cost (respectively) for good j, Q_j is the demand for product j, and ϕ is the artist's share of the profits from good 1. We include ϕ to capture the fact that artists' royalty payments and advances typically amount to a small share of the overall profits from recordings.[10] K_1 represents the production, marketing, and distribution costs associated with producing an album, and K_2 represents the fixed costs of performing live, which could include the opportunity costs of the artist's time. We assume that the costs of touring are heterogeneous in the population of artists: for a given artist, K_2 is a random draw from some continuous distribution G. The artist expends effort levels e_1 and e_2 on album production and live performances, respectively, at a cost $h(e_1, e_2)$. The effort variables could be interpreted as time spent by the artist on the respective activities; for example, e_2 could represent how much time the artist spends on tour, or how many

[9] The model is stylized to reflect features of the music industry, but the ideas and results could be generalized to other markets where firms sell an easily copied (i.e., digitalizable) product and another complementary product.

[10] For simplicity, the share of revenues from concerts that artists receive are assumed to be one in the model. In reality, artists usually receive 70 to 90 percent of concert revenues, but earn additional revenues from such peripheral activities as T-shirt sales, so that concert revenues are a reasonable approximation to the revenues captured by the artist.

concerts the artist performs.

We assume that the cost of effort function h is increasing and convex, and that $\partial^2 h/\partial e_1 \partial e_2 > 0$—i.e., going on tour increases the marginal cost of effort in album production. This formalizes the idea that artists face tradeoffs when allocating their effort between recording new albums and touring.

Consumers have unit demands for the two goods, and their willingness to pay depends on the artist's quality (θ) and on the effort ($e = (e_1, e_2)$) put into album production and live performances. Specifically, for an artist of quality θ, a consumer's valuations for the two goods are v_1 and v_2, drawn randomly from a joint distribution $F_\theta(v_1, v_2; e)$ on the support $[0, \infty) \times [0, \infty)$. A higher-quality artist generates more demand in the sense that if $\theta > \theta'$ then $F_\theta(v_1, v_2; e) < F_{\theta'}(v_1, v_2; e) \ \forall \, v_1, v_2$—i.e., holding effort fixed, the distribution of consumer valuations for an artist of type θ stochastically dominates the distribution for artists of types lower than θ. Similarly, holding artist quality fixed, demand for both products is assumed to be increasing in the artist's effort levels. In particular, the marginal distributions $F_i(v_i; e)$ are decreasing in *both* e_1 and e_2; so, for example, increased effort on concert performances may boost demand for albums. (In practice this demand increase could result from consumption complementarities or the effects of promotional activity and increased radio airplay that accompany a concert performance.) However, we assume that effort on album production affects album demand more than it affects concert demand (and vice versa) in the sense that $\mid \partial F_i/\partial e_i \mid > \mid \partial F/\partial e_i \mid > \mid \partial F_i/\partial e_j \mid$.

We represent the impact of file-sharing by assuming that with some probability γ a consumer can obtain good 1 for free (instead of paying p_1). We interpret γ as the the pervasiveness of file-sharing or the availability of albums on the network, compounded with consumers' propensity to download rather than purchase their music. The share of consumers purchasing good 1 is then $(1 - \gamma)[1 - F_1(p_1; e)]$.

The key assumption we make about concert demand is that only consumers who obtain the album (either via purchase or download) are in the market: i.e., consumers won't attend a concert without first listening to the album. This implies the share of consumers who purchase good 2 is

$$Q_2 = (1-\gamma)\int_{p_1}^{\infty}\int_{p_2}^{\infty} dF(v_1, v_2; e) + \gamma \int_{0}^{\infty}\int_{p_2}^{\infty} dF(v_1, v_2; e) \, .$$

$$= (1-\gamma) Prob(v_1 \geq p_1 \,\&\, v_2 \geq p_2) + \gamma Prob(v_2 \geq p_2)$$

Among downloaders (who occur in proportion γ), all those with $v_2 \geq p_2$ will purchase good 2. Among non-downloaders (in proportion $1-\gamma$), only those with $v_1 \geq p_1$ *and* $v_2 \geq p_2$ purchase good 2. Essentially, file-sharing increases demand for the concert by bringing into the market consumers who would not otherwise obtain the album, but whose valuations of the concert (after downloading the album) exceed the price.[11]

We note that a more realistic model would allow consumers who do not obtain the album to still attend the concert. The stark assumption that only album-owners are in the market is made primarily for convenience; the main ideas and results would hold if instead we simply assumed that album-owners have higher demands than non-owners. For example, a model with consumption complementarities implies that album-owners will value concerts more highly, and delivers the same predictions as we outline below.

Artists' contracts with record labels usually stipulate a nominal level of support for concert tours. However, record labels historically have not claimed a share of the revenues from live performances. We therefore assume that only albums with positive expected profits (from album sales alone) get produced. We also assume that artists can sell live performances only if they also sell recorded music—i.e., they cannot tour without an album. Together, these assumptions imply that only artists with θ's exceeding some threshold θ^* will be active in the market, since artists below that threshold cannot deliver positive expected profits from the sale of good 1 alone.

Finally, we treat p_1, the price of the recorded good, as fixed and exogenous. This means album prices are outside the control of the artist, so the artist's only pricing problem is to choose p_2. This is consistent with actual practice in the music industry: prices are set by the

[11] Note that we have implicitly assumed full displacement of a CD sale by a download. Alternatively, one could view γ as measuring the extent of displacement for CD sales. An attractive generalization of the model would allow for heterogeneity in γ across artists or genres.

record labels, not by the artists. The stronger restriction is that p_1 is fixed—i.e., artists act under the assumption that record labels will not change album prices even in response to file-sharing. This is obviously a strong assumption, but it is surprisingly descriptive of reality: CD prices are remarkably rigid, both across albums and over time. In particular, album prices were remarkably unresponsive to the diffusion of file-sharing technologies starting in 1999. Prices changed only slightly between 1999-2003 (and in fact they increased) and only in 2004 did prices begin to decline.

Given this framework, we can establish some basic predictions regarding the impact of an increase in the prevalence of file-sharing (γ) on artists' decisions. (We give mostly heuristic explanations here; detailed proofs are included in the appendix.)

Prediction 1: *All else equal, an increase in the prevalence of file-sharing causes artists to reduce effort spent on album production; i.e.,* $\partial e_1^* / \partial \gamma < 0$.

This result follows from the fact that increases in file-sharing diminish the marginal benefit of effort on album production (without doing anything to change the costs). As the prevalence of downloading increases, artists capture a shrinking fraction of any additional album demand their efforts generate, so they optimally allocate less effort to album production. Moreover, increases in file-sharing make artists more likely to perform concerts, so in that sense artists respond to file-sharing by shifting effort away from album production toward live performances:

Prediction 2: *If the artist's share in the recording profits is low, increases in file-sharing lead the artist to increase effort on live performances.*

If we interpret e_2 as the time the artist spends on tour or the number of concerts she performs, this prediction implies that file-sharing should have led to an increase in touring activity. To understand why, note that increases in file-sharing (γ) have two effects on artists' profits. First, file-sharing increases concert demand: consumers who download the artist's music (but who otherwise would not have purchased it) become an extra source of demand for the artist's live performances. Since artists keep the vast majority of concert revenues, the effect on artists' profits is potentially large. The second effect goes in the opposite direction: file-sharing diminishes artists' incentives to perform concerts as a way of boosting album sales, since a higher fraction of album sales is displaced by downloads. However, the second effect is likely to be dominated by the first, because artists' share of album revenues is typically very small.

Importantly, the model's prediction that artists reallocate effort away from recording toward live performances does not merely reflect the sales displacement effect of file-sharing. Even if the *only* effect of file-sharing were to decrease album sales, we would still predict a shift in effort toward concerts. However, the model implies that file-sharing has a non-neutral impact on concert demand:

Prediction 3: *All else equal, increases in file-sharing lead to higher demand and higher revenues for concerts.*

Using the argument above that concert profits weakly increase with file-sharing (since file-sharing's direct effect is to increase demand), it follows that revenues must go up too (since we're assuming marginal costs didn't change).

Note, however, that the impact of file-sharing on concert prices is ambiguous. The first-order condition that determines p_2^* is

$$(p_2 - c_2)\frac{\partial Q_2}{\partial p_2} + Q_2 = 0 \ .$$

File-sharing increases demand but also rotates it (making $\partial Q_2/\partial p_2$ more negative), so it is not clear whether p_2^* will go up or down when γ increases.

Our last prediction relates to record labels' decisions of whether to finance new albums:

Prediction 4: *All else equal, an increase in file-sharing results in fewer artists producing albums—i.e., increases in γ lead to increases in the threshold quality θ^*.*

This is a straightforward consequence of our assumption that albums only get produced if their expected standalone profits are positive. Profits are obviously lower when record labels and artists are forced to compete with free copies of their own product, so increased file-sharing means that fewer artists will be able to generate enough expected sales to justify the production of an album. We note, however, that more complex models of music demand could deliver an opposite prediction. For example, if widespread file-sharing facilitates the spread of information, it may make it easier for new artists to develop an audience. Alternatively, if profits from touring are substantial then some artists may self-finance the recording of an album.

However, even if record labels were to claim a share of the revenues from artists' concert tours, file-sharing should reduce the overall profitability of music production and push toward fewer active artists in equilibrium. If free distribution of recorded music could boost concert demand enough to increase overall profits, then presumably artists and record labels would have chosen to give away music rather than sell it. Of course, even if industry profits are decreased by file-sharing, it could still be the case that total social surplus is dramatically increased, especially if record labels' market power resulted in artificially high prices of the recorded good prior to the advent of file-sharing. It can also be the case that artists' profits increase under filesharing, because they benefit more from file-sharing than from the price chosen by record labels for CDs.

The objective of this paper is to document the nature of the complementarities between live and recorded music, and test whether the predictions outlined above are borne out in the data. In the following sections we show the basic trends in the recorded and live performance markets over time, and provide evidence on the presence of demand-side spillovers. The tests of the predictions indicate that after the advent of file-sharing in 1999: (1) established artists took longer to release new albums, (2) artists were more likely to tour, (3) concert revenues increased, and (4) fewer new artists debuted a recorded product. In addition to this time variation, we provide cross-sectional analyses of each of these trends based on artist-level observables (artists' ages and genres) and market-level observables (market size and the broadband penetration of each market). Unlike the simple time trends, which are uniformly consistent with the predictions of the model, the evidence from these cross-sectional analyses is more mixed: some of the tests appear to confirm our interpretation (that the changes over time reflect the impact of file-sharing), but others are noisy and/or difficult to interpret. One additional cross-sectional analysis that we intend to conduct is to compare trends across artists and markets based on the extent of the decline in CD sales immediately following the introduction of file sharing—e.g., to see if artists who were the hardest hit by file-sharing on the recorded music side were also the ones that had the greatest increases in concert demand. Additional "tests" of the predictions could also be proposed, but given the abilities (and limitations) of our dataset, these seem to be the most sensible correlations to examine with respect to the predictions above.

While these patterns are consistent with the model's predictions, we note that the model also identifies alternative explanations for these patterns. For example, changes in the distribution of K_2 (the opportunity costs of touring) or changes in the distribution of consumer valuations F (e.g., due to changes in the demographic characteristics of concert-goers) could also rationalize many of the patterns we observe. We consider the leading alternative explanations in section 5.

4 Empirical Analysis

Both live and recorded performances are important sources of revenue in the music industry. However, the technology for distributing the digital (i.e., recorded) good in this industry underwent dramatic changes after the advent of online file sharing, while the technology for producing/distributing live performances was not affected.[12] The stylized model outlined above generates several clear predictions about the impact of this technology shock, and the general purpose of our empirical analysis is to test whether the data are consistent with these predictions. We note, however, that we do not have rich information on the nature and extent of file-sharing to match our data on CD sales and concert activity. Instead, we focus on variation over time and across artist and market characteristics: we look for changes in the supply of recordings and concerts that coincided with the rapid diffusion of file-sharing beginning in 1999, and we ask whether those changes appear to be most pronounced for cohorts of artists or markets where file-sharing is thought to have been most prevalent.

[12]This is not to say that live performances were not affected by other changes. We provide a preliminary discussion of some of these changes in the section on alternative explanations.

Table 1: Summary Statistics, Matched Sample*

	Tickets	Concerts	Revenue	Ticket Price	Bands	Concerts /Band	Venues	CDs	Ratio= Rev/CDs
1993	10.0	2956	189	18.87	408	7.24	1024	172	1.10
1994	12.1	3890	258	21.40	477	8.60	1201	183	1.41
1995	13.8	4588	285	20.66	533	8.61	1315	195	1.46
1996	12.8	4803	272	21.22	562	8.55	1280	228	1.19
1997	14.3	5170	375	26.21	649	7.97	1312	248	1.51
1998	16.0	4610	422	26.36	730	6.32	1281	288	1.47
1999	16.0	4846	501	31.34	743	6.52	1224	280	1.79
2000	14.6	4788	490	33.66	764	6.27	1214	263	1.86
2001	16.0	6537	601	37.56	838	7.80	1300	238	2.53
2002	17.4	7609	622	35.77	1000	7.61	1245	234	2.66

*Tickets, Revenue and CDs divided by 1,000,000. Concert revenues and ticket prices in December, 1997 dollars using the Entertainment CPI. Ratio calculated using the Entertainment CPI to deflate estimated revenues from both concerts and CD sales.

4.1 Basic Patterns Over Time

Table 1 provides summary statistics from our matched sample of artists' concerts and CD sales and illustrates the important time trends. The first column contains total concert tickets sold. The number of concert tickets sold increases over the sample period, as do the number of concerts, as shown in column two. Columns three and four show concert revenues and average ticket prices in each year, which also increase sharply from an average ticket price of $18.87 in 1993 to an average ticket price of $35.77 in 2002 (reported in December 1997 dollars). Note that the number of concerts and concert revenues both increased quite sharply in 2001 and 2002, at which time file-sharing had become widespread. The increases in price were also most dramatic in 1999-2002.

The number of artists on tour more than doubles from 1993 to 2002, with the total number of bands in those years increasing from 408 to 1000. The venues in which concerts are performed can range from small auditoriums and clubs to large stadiums. The number of venues covered by the data is roughly stable over time, as shown in column seven. The next to last column of table 1 provides information on the sales of CDs. CD sales rise quickly until 1998. In 1999, sales of CDs drop modestly, and in the years after 1999, CD sales drop quickly, back to 1996 levels. In the last column, we provide a rough estimate of the relative importance of concert and CD revenues for a typical band in the industry. 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 CD sales using the Entertainment CPI to deflate concert revenues. This ratio is increasing over the sample period from 1.10 in 1993 to 2.66 in 2002, with the most dramatic change coming in 2001. In other words, in 1993, total concert revenues for bands are estimated to be roughly equal to total CD revenues, while in 2003, total concert revenues for bands are estimated to be over 2.5 times larger than CD revenues. Note that this ratio provides information on total concert revenues to bands versus total CD revenues to bands. Clearly, there is significant heterogeneity across artists, and one might also be interested in calculating this ratio separately for each artist and then reporting its average. We discuss artist heterogeneity in the coming sections and provide related tabulations from the data. In table 1, however, we wish to convey the following three significant trends in this industry. First, concert prices and quantities are increasing over time. Second, CD revenues and quantities are decreasing in the second half of the sample. Finally, the aggregate importance of concert revenues is increasing over time.

The sample described in table 1 is used in all later analyses. As discussed in the data description, this sample consists of bands that both tour and sell CDs on at least one occasion between the years of 1993 and 2002. This sample is a subset of the set of all bands that tour; throughout the sample period, the proportion of all concert tickets sold that are included in our sample is around one-third, and the same is true for the total proportion of bands. However, the selected sample appears to be representative of the population: all trends over time look qualitatively the same between the two datasets.

4.2 Evidence of Demand-Side Spillovers

In the model of section 3, the demands for albums and for concerts are explicitly interrelated: sales (or downloads) of albums boost demand for concerts, and concerts stimulate demand for CDs. The presence of the latter effect is usually taken as given in the music industry—indeed some people describe the purpose of a concert tour as simply to promote an album—and could result from various demand-side mechanisms. For example, concert-goers may purchase additional recorded music by the artist in anticipation of the concert event to enhance their enjoyment of the concert, or because concert-goers receive more enjoyment from listening to a CD if they have also attended a performance of the music. Promotional activity surrounding a concert event, including increased radio airplay, could also stimulate CD sales.

In this section we measure the impact of concert performances on CD sales, and test whether the magnitude of the spillover changed after file-sharing became widespread. We estimate the following model of demand for CD sales for artist i in year j in market m at week t:

$$sales_{ijmt} = \alpha_{ijm} RON\_sales_{ijmt} + \beta_{ij} D_{ijmt} + \epsilon_{ijmt}, \tag{2}$$

where RON_sales is "rest-of-nation sales" (i.e., national sales excluding the sales in this DMA), and all sales numbers are in levels.[13] By conditioning on the rest-of-nation sales and allowing its coefficient to be market-specific, we essentially allow each market to represent a constant share of national weekly sales for the artist. Rest-of-nation sales varies over weeks, and captures observable nationwide shocks to demand for an artist's CDs in each week. For example, if winning a Grammy award or playing on a national television show increases an artist's national CD sales in subsequent weeks, we capture that effect. The vector D_{ijmt} is a series of dummy variables that turn on in the weeks surrounding a concert event. Here, we use a vector of dummy variables of length 17: thus, the first element of D_{ijmt} turns on

[13] We get the same qualitative effect in logs, with the effect of a concert on the log of sales being very similar across artists. Here we show the level effects, because the effect of a concert in artists' profitability utimately depends on the levels of sales, not the logs, and we wanted to preserve this source of heterogeneity for artists' profits.

8 weeks before a concert, the second element turns on 7 weeks before a concert, and so on, until 8 weeks after the concert event. Concerts in a given year by a given artist tend to be of similar size, so we do not interact D_{ijmt} with the size of a concert. All coefficients have i and j subscripts, allowing them to vary across artists and years. In other words, the specification leads to a series of regressions that are run separately for each artist in each year. This takes advantage of the extensive variation in the dataset across the 100 geographic markets.[14] These "year-by-year" results indicate a large change between the years 1993-1998 and the years 1999-2002, with the total effect of a concert on CD sales dropping sharply between those two periods. We therefore estimate a slightly more restrictive model, which pools the 1993-1998 and 1999-2002 years, and includes year fixed-effects. For each artist in these two sets of years, we can graph the β_{ij} coefficients to see the effect of a concert on CD sales. Figure 1 shows the median effects (across artists) in each of the two periods.

Figure 1: CD Sales Surrounding a Concert Event

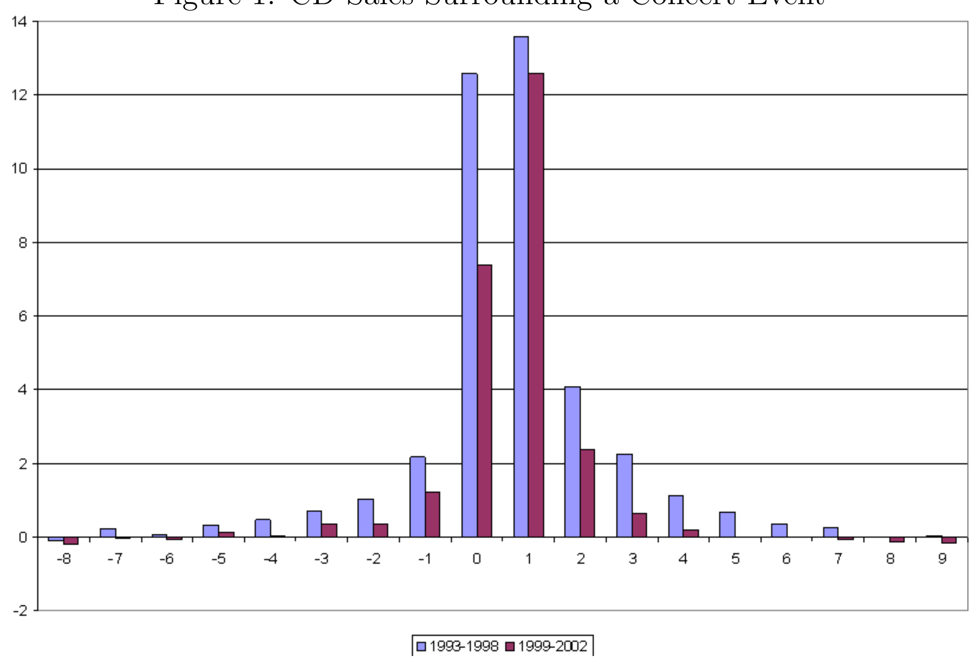

[14]A pooled regression would yield smaller standard errors, but our sample sizes are sufficiently large that statistical precision is attainable even in artist-specific regressions, with the exception of artists that do not have extensive concert tours across markets.

There are two important facts to take away from figure 1. The first is that concert events are strongly correlated with a concomitant increase in local CD sales. The "bump" in sales is strongest in the week immediately following the concert, and in the week of the concert itself. There is a relatively smooth run-up in CD sales on either side of the concert, with a larger effect after the concert than before. Although not shown here, the effects in logs indicate that the total percentage increase in sales over the 17 weeks surrounding a concert is roughly 150 percent. The second important fact is that the increase in CD sales surrounding a concert is lower in 1999-2002 than in 1993-1998. These periods correspond with the periods of overall CD sales growth and decline, respectively. Figure 1 shows the median effects across artists in each of the two periods, indicating this decline. Because the pool of artists changes in each of the two periods—e.g., there are more bands touring in 2001 and 2002—the effects might reflect compositional changes. However, the figure looks essentially the same if it is drawn using only artists who performed concerts in both periods.

Figure 1 shows median effects primarily for ease of presentation; the underlying estimation results contain the same information at the individual artist level. Not surprisingly, heterogeneity in the artist-specific estimates is substantial, as for some artists the impact of a concert performance on CD sales is much larger than for others. Some of this heterogeneity reflects differences in artist popularity: in absolute terms, the magnitude of the spillover effect will naturally be larger for superstars than for relative unknowns. Our ability to estimate artist-specific effects means we can perform within-artist tests of the hypothesis that spillover effects from concerts diminished after the advent of file-sharing. In particular, we can test whether the sum of the coefficients in 1993-1998 is larger than the sum of the coefficients in 1999-2002 for each individual artist. For nearly two-thirds of all artists active in both sets of years, the difference is positive (meaning that the sales "bump" is smaller for those artists in 1999-2002 compared to 1993-1998), and in 60 percent of those cases the difference is statistically significant.

The shrinking spillovers from concerts to CD sales are consistent with the file-sharing hypothesis, but of course other factors could have been changing around that time as well. Table 2 shows changes in the concert spillovers broken down by artist and market characteristics; if file-sharing was the main driver of spillover declines, then the changes should be most pronounced for artists and markets where file-sharing was most important. To construct the table, we first estimate a simplified version of regression 2, and include interaction terms with various artist and market characteristics. Specifically, we estimate the regression

$$sales_{ijmt} = \alpha_{ijm} RON\_sales_{ijmt} + \beta_{ij} D_{imt} + \delta_{ij} D_{imt} X_{im} + \epsilon_{ijmt}, \qquad (3)$$

where j now denotes period (i.e., we group the years as 1993-1998 vs. 1999-2002, instead of estimating separate coefficients for each year), D_{imt} is a dummy equal to one if artist i performed a concert in market m within 8 weeks on either side of week t (i.e., we estimate the average increase over the 17-week period rather than estimating the week-by-week effects shown in figure 1), and X_{im} are the artist or market characteristics we interact with the concert dummy.

Table 2: Average increases in CD sales (17-week period surrounding a concert)

	1993-1998	1999-2002	% Change
Overall	162.93	83.28	-48.9
Country/Folk	165.66	99.43	-40.0
Jazz/Latin	127.37	24.37	-80.9
Rock	161.33	76.91	-52.3
Urban/Rap	178.54	114.70	-35.8
Age $<$ 5	150.23	62.25	-58.6
Age 6-10	239.58	161.37	-32.6
Age 11-20	123.69	57.96	-53.1
Age 21+	129.21	84.91	-34.3
High broadband	203.30	71.29	-64.9
Low broadband	119.54	90.90	-24.0
Large market	183.96	86.16	-53.2
Small market	128.60	87.44	-32.0

Based on coefficients from regression 3. Ages are as of 1999. High broadband markets and large markets are the top quartile of the sample in terms of 2003 broadband penetration and market size, respectively.

Averaged across all artists and markets, the increase in local CD sales around the time of a concert was around 163 units in the years 1993-1998. After the advent of file-sharing, the spillovers from concerts were roughly half as large. Although little concrete evidence has been brought to bear on the question of which artists or genres have been most affected by file-sharing, the conventional wisdom seems to be that file-sharing was more active for popular music genres (like rock and rap) than it was for jazz or classical music. If in fact this was the case, we should expect the spillovers to have declined more sharply for rock and rap artists than for jazz artists. The table suggests the opposite: the change in the spillover effects of concerts appears to have been most pronounced in the Jazz/Latin category.[15] However, we are reluctant to draw strong inferences from genre comparisons, because our grouping of artists into genres is very coarse, and (more importantly) we suspect the accepted wisdom about file-sharing—i.e., that it was most prevalent for pop music—is at best an oversimplification.

The table also reports breakdowns based on artist age. For the purposes of this table (and all remaining analyses in the paper), we calculate an artists's age as 1999 minus the artist's debut year. Some authors have suggested that file-sharing should be most important for young bands, since downloading is a way of costlessly sampling new music. Of course, similar reasoning suggests that concert spillovers will be largest for young bands: the impact of increased airplay and promotional activity around the time of the concert is greatest when many consumers have not yet heard of or purchased the artist's music. The numbers in table 2 appear to be consistent with both of these ideas: spillovers were in fact largest for young artists, but young artists also experienced the largest decline in the spillover effect.

A somewhat more direct check of the file-sharing hypothesis is to test whether the concert spillovers declined more dramatically in cities with high broadband penetration. Downloading music, especially during the period of our sample, was considerably faster and more convenient over a broadband (i.e., high-speed) internet connection. Naturally, we should expect file-sharing to have the most significant impact in cities with the highest rates of broadband usage. This appears to be true for concert spillovers: as the table shows, declines in the spillover effect were much more dramatic in cities that were in the highest quartile

[15] The Jazz and Latin categories were grouped together because there were no significant differences in any of the basic time trends from table 1 for these groups. Furthermore, the sample sizes in both groups are relatively small, so drawing conclusions from the individual genres would be difficult. The same logic was applied when grouping Country with Folk, and Urban with Rap.

of broadband penetration.[16] The effects were also more pronounced in large cities than in small cities. Note that although broadband penetration is positively correlated with market size, the correlation is far from perfect. The differences based on broadband penetration rates are robust to the inclusion of controls for market size.

One important caveat about the estimates underlying figure 1 and table 2 is that artists' decisions about where and when to perform concerts are endogenous. Given the flexibility of our specification, however, this endogeneity only poses a problem if artists are able to time their concerts to coincide with unobserved local demand shocks. We have adequate controls for the fact that artists choose markets where their music is popular, and for the fact that they tour most heavily when their music is selling well nationwide; our estimates will only be biased if artists can somehow schedule their events to correspond with specific weeks of high local demand in a specific city. A related issue is the question of what exactly a "concert event" is. There may be promotional activities or extra radio airplay surrounding a concert event, and the effects of these activities are also included in the sales bump in figure 1. A second caveat is that the bump in CD sales surrounding a concert could reflect intertemporal substitution, rather than additional sales. In other words, a concert-goer could just make a planned purchase of a CD earlier, rather than buying a CD that he otherwise would not have purchased.[17]

4.3 Changes in the Demand for Concerts

Assuming concert demand is an increasing function of listenership, digital sharing of recorded performances should increase the overall demand for live performances. With sharing, each purchased CD potentially results in many listeners. This could result from direct sharing, as when an individual purchases a CD and then burns copies for several friends using a

[16] We use the fraction of households with broadband connections at home, as reported in a 2003 survey by Forrester Research. We are thankful to Austan Goolsbee for providing these data.

[17] The results in figure 1 include sales of all CDs produced by the artist. If intertemporal substitution is important, then we might expect the measured spillovers to be largest for the artist's most recently released album; for older albums, potential purchasers have had plenty of time to make a purchase decision. When estimating the model using only the most recent album, we see similar patterns as those useing all albums, which makes us less concerned about the importance of intertemporal substitution here.

CD-writer, or from broad distribution via peer-to-peer networks. In principle, even one purchased copy of a CD could translate into tens of thousands of actual listeners. In the absence of such digital redistribution technologies, by contrast, the number of consumers listening to an artist's album would roughly equal the number of albums purchased.

The numbers in Table 1 suggest the demand for concerts expanded after 1999, when digital file-sharing technologies became widely available. While CD sales declined, both the total number of concert tickets sold and the total number of concerts performed increased, in spite of sharply increasing ticket prices. The number of different artists performing concerts also increased dramatically. However, the patterns in the table reflect aggregate movements, and do not necessarily imply that concert demand increased for the typical artist. Also, the composition of artists touring could have changed over the sample period, or the effects could be driven by increases affecting only a small subset of artists.

In order to investigate these issues more carefully, we analyze a simple regression model. Letting $\tilde{r}_{ijmt}$ represent the log of ticket revenues for artist i's concert in year j in market m in week t, we estimate

$$\tilde{r}_{ijmt} = \theta_{0j} PAST_{ijmt} + \theta_{1j} CURR_{ijmt} + x'_{ijmt}\beta + \eta_i + \delta_m + \epsilon_{ijmt} \ , \tag{4}$$

where $PAST_{ijmt}$ is the logs of cumulative past CD sales for artist i in year j market m as of week $(t-26)$ (i.e., a measure of the *stock* of purchased recordings up to six months before the concert), $CURR_{ijmt}$ is the log of cumulative sales over the six months prior to the concert (a measure of the current *flow* of purchases), η_i and δ_m are artist and market fixed effects (respectively), and x_{ijmt} is a vector of additional controls (e.g., week-of-year dummies to control for seasonality (not currently included), and the log of CD sales for artist i in week t in all other markets besides m to control for nationwide shocks to the artist's popularity).

For artists whose careers began after January 1993, our data allow us to calculate the key explanatory variables directly. However, many successful touring artists have albums dating back to the 1970's, and for these artists we do not directly observe cumulative past album sales in each market. For such artists, we impute cumulative past sales from Recording

Industry Association of America (RIAA) album awards. The RIAA issues Gold and Platinum status to albums when they reach sales of 500,000 or 1,000,000, respectively.[18] The dates of these sales awards are documented in a publicly available database that extends back to 1958, which allows us to construct estimates of cumulative past sales (as of 1993) for older artists in the sample.[19]

A central objective is to test whether θ_0 and θ_1 changed after file-sharing became widespread in 1999. The results from this regression are reported in Table 3. We provide two specifications: the first specification estimates equation 4 above. The second specification pools the years from 1993 - 1998 together, and the years of 1999 - 2002 together, as the coefficients within these two groups of years are not significantly different from each other.

Table 3: Regression Analysis of Concert Revenues

Specification:	(1)										(2)	
Year	93	94	95	96	97	98	99	00	01	02	93-98	99-02
θ_{0j}'s	0.29	0.30	0.33	0.34	0.33	0.35	0.32	0.31	0.30	0.29	0.33	0.32
θ_{1j}'s	0.22	0.18	0.16	0.14	0.15	0.13	0.19	0.22	0.23	0.26	0.16	0.21
$E[\theta_{0j} PAST +$ $\theta_{1j} CURR]^*$	3.85	3.82	3.95	3.96	4.01	4.02	4.20	4.24	4.23	4.19	4.00	4.24
R^2	0.45										0.45	
N	47,675										47,675	

*Evaluated at average value of $PAST$ and $CURR$ in each year.

The relationship between the number of CD owners in a geographic market and subsequent concert sales is clearly changing over the course of the sample, with an apparent break in 1999. Although the number of fans in the market for a concert ticket is a relatively stable function of past CD sales, the effect of recent CD sales (which is presumably where the

[18] For each multiple of one million units that an album sells, the RIAA upgrades the award to "Multi-Platinum," and albums that sell 10 million copies are given "Diamond" status.

[19] The imputation need only be done for sales prior to 1993; regardless of the age of the artist, we observe actual sales from January 1993 onward.

biggest effect of file sharing is likely to be felt) clearly changes after 1998. Pooling the years into 1993-1998 vs. 1999-2002 (specification 2), the estimates indicate that before 1999, a 100 percent increase in the number of CDs sold within six months prior to a concert event is associated with a 16 percent increase in concert revenue. After 1999, this number increases to 21 percent, and the difference is statistically significant.

Another way to interpret these numbers is to ask how many additional CDs must be sold in order to generate one additional sale of a $20 concert ticket. As Table 4 reports, this number declined from 8.5 in the 1993-1998 period to 6.4 in 1999-2002. The decrease is consistent with the prediction that file-sharing increases the "conversion rate" of CD sales into listeners: with file-sharing, artists can acquire fans without ever selling them any recorded music. The table also reports breakdowns across artist and market characteristics. Unlike the differences shown in table 2 for concert-to-CD spillovers, the differences here provide little additional support for the file-sharing hypothesis. The changes are roughly the same for all age groups, and for high vs. low broadband cities.

Table 4: How many additional CDs must be sold to generate $20 of concert revenue? *

	1993-1998	1999-2002	% Change
Overall	8.47	6.36	-24.85
Country/Folk	6.07	4.61	-24.08
Jazz/Latin	13.04	9.04	-30.69
Rock	7.13	7.13	-25.22
Urban/Rap	40.23	19.39	-51.80
Age<5	18.36	13.62	-25.80
Age 6-10	9.21	7.88	-14.40
Age 11-20	5.37	4.42	-17.84
Age 21+	4.79	4.00	-16.44
High broadband	9.38	7.06	-24.78
Low broadband	7.81	5.87	-24.84
Large market	10.07	7.72	-23.35
Small market	5.10	3.41	-33.09

* Based on coefficients from regression (4), with the *CURR* variable interacted with dummies for the respective groups.

4.4 Changes in the supply of live performances

An implication of the theoretical discussion outlined earlier is that increases in file-sharing lead an artist to increase effort on live performances. An increase in effort could be manifested in several ways: artists might produce higher quality shows (which may be reflected in higher revenues and/or higher prices and is consistent with the results from the previous section), or artists could increase touring activity (a direct shift out in the supply curve for a given quality level). Table 5 explores the second of these two potential effects. The first two columns of table 5 gives the baseline probabilities of touring in the 1993-98 period, and the 1999-2002 period, broken out for the same genre and age categories as the earlier results. We also provide the relative percentage change in these probabilities between the two periods, as well as estimates from a linear probability model, which corrects for DMA and artist fixed effects. The linear probability model specifies the probability of touring at the artist-year-DMA level as:

$$y_{ijm} = \alpha_i + \delta_m + \beta \text{ Period} + \epsilon_{ijm}$$

We count an artist as having toured in a DMA in a year if they performed at least one concert there in the given year, regardless of the size or revenues of the concert. Older (11-20 year old) bands and, perhaps surprisingly, rock bands, were actually less likely to tour in the 1999-02 period relative to the 1993-98 period. Overall, however, bands were more likely to tour in the later period. The largest increases in touring activity were seen by young bands (28.4 percent relative increase in touring activity), and among Jazz/Latin and Urban/Rap bands. We do not find large differences in the probability of touring across cities with high vs. low broadband penetration. We do, however, see larger increases in touring probabilities in larger cities.

4.5 Changes in the production of recorded music

Among the predictions of the model outlined in Section 3 is that file-sharing erodes the profits from recorded music, increasing the quality threshold required for an album to be

Table 5: Probability of Touring

	Prob(tour) 1993-98	Prob(tour) 1999-02	Relative % Change	Relative % Change with DMA, Band FEs*
Overall	3.20	3.33	4.1	4.1*
Country/Folk	3.04	3.02	-0.7	5.3*
Jazz/Latin	1.04	1.42	36.5	38.5*
Rock/Metal	4.13	3.88	-6.1	-2.4*
Urban/Rap	1.31	2.11	61.1	59.5*
Age < 5	2.64	3.28	24.2	28.4*
Age 6-10	3.51	3.57	1.7	1.1
Age 11-20	3.57	3.12	-12.6	-1.1*
Age 21+	2.96	2.98	-24.7	0.3
High broadband	5.01	4.88	-2.6	0.2
Low broadband	2.59	2.71	4.6	6.6*
Large market	7.74	8.35	7.9	10.3*
Small market	1.67	1.54	-7.8	-5.4*

* Indicates significance at 5%.

commercially viable. Absent any changes in the distribution of talent, file-sharing should therefore reduce the supply of new artists and new albums.

This straightforward prediction turns out to be difficult to test, because it is surprisingly difficult to count new artists and albums. Defining the universe of relevant artists is itself problematic—there are over 200,000 independent musicians registered with *garageband.com*, for example—and although various databases containing detailed album information exist, most are incomplete and/or have unreliable information about album release dates. In order to estimate artist and album counts by year and genre, we started by gathering the list of all artists with albums for sale on Amazon.com in early 2005. We use this as our universe of artists for two reasons: first, it imposes a reasonable size threshold (you're only an artist when you have an album to sell); second, that size threshold is still relatively low (Amazon's inventory includes over 80,000 artists, the vast majority of whom are unknown artists with small, local fanbases). We then gathered discography information for these artists from an online database that we found to have reliable release dates for major studio recordings.[20] It is then straightforward to count the number of new artists[21] and new albums by year.

This procedure for counting new artists and albums is obviously imperfect. One concern is that Amazon's inventory may be a selected sample. If Amazon drops artists from the online inventory when the artist released an album but then faded to obscurity, then we may undercount new artists in earlier years. Although we have no way of verifying this, we suspect it is not a major problem, because we observe a large number of one-album-and-gone artists who debuted in the early 1990's. To be conservative, we only report numbers as of 1995. Another question is whether it makes sense to count artists on small "boutique" labels. Our universe obviously includes some very small artists who may be essentially self-publishing their music; one could argue that the predictions of our model apply more cleanly to major labels. The patterns we report below are qualitatively similar if we exclude artists whose albums were produced by obscure labels.

Table 6 reports the counts for five major genres. For every genre, the number of new artists and new albums peaks sharply in 2000 and is followed by a steep decline from 2001-2004. Given that file-sharing would likely affect album production with a time lag of at least one year—e.g., if decisions to cut back artist rosters began in early 2000, artists signed in 1999 would still end up releasing albums in 2000—the general pattern is loosely consistent with the predictions of the model in Section 3. However, we are reluctant to draw any strong conclusions from these broad patterns, because (a) we have no clear explanation for the dramatic spike in 2000, and (b) there are obviously many other factors we are not controlling for. There may have been important time trends in the underlying demand for certain genres, and of course record labels' investments in new artists and albums are likely to track the overall economy. Indeed, the time series of new artists is positively correlated with real GDP for the years 1995-2004 (the simple correlation is 0.52). However, the decline in the number of new artists after 2000 appears to be much sharper than would have been predicted by the business cycle alone.

[20] Compiling discographies from Amazon's inventory is possible in principle; however, it is extremely difficult to weed out singles, compilations, live recordings, imports, interview CDs, etc. The online database we used lists major studio recordings separately from these other kinds of minor releases.

[21] We count an artist as new in the year of her debut album.

Table 6: New Artist and New Album Counts, by Year

New Artists

	Country / Folk	Jazz / Latin	Rock	Urban / Rap	**Total**
1995	642	1,982	1,359	902	4,885
1996	621	1,537	1,422	1,006	4,586
1997	565	1,315	1,301	1,056	4,237
1998	596	1,217	1,395	1,177	4,385
1999	644	1,453	1,984	1,493	5,574
2000	1,563	1,894	3,257	1,983	8,697
2001	924	1,331	2,667	1,758	6,680
2002	785	1,252	2,519	1,477	6,033
2003	676	1,337	2,316	1,174	5,503
2004	904	1,627	1,919	1,297	5,747

New Albums

	Country / Folk	Jazz / Latin	Rock	Urban / Rap	**Total**
1995	1,616	6,510	3,682	1,726	13,534
1996	1,622	4,726	3,711	1,855	11,914
1997	1,374	3,927	3,412	2,047	10,760
1998	1,426	3,646	3,688	2,387	11,147
1999	1,464	3,983	4,428	2,922	12,797
2000	2,712	4,846	6,352	3,401	17,311
2001	1,867	4,072	5,675	3,332	14,946
2002	1,799	4,203	5,619	2,879	14,500
2003	1,672	4,430	5,770	2,565	14,437
2004	1,871	4,680	4,903	2,789	14,243

Did artists reduce effort on album production?

Because artists' effort levels cannot be easily observed or measured, we of course cannot directly test whether they shifted effort away from recorded music toward live performances. However, the increase in touring activity documented in section 4.4 plausibly came at the expense of album production. Anecdotally, record company executives often complain that a busy touring schedule can delay the production of an artist's next album. This suggests an observable outcome of the predicted shift in effort: elapsed time between album releases should have increased after 1999.

Measuring time between releases is complicated by attrition—many artists never release a next album because they are dropped by their record labels—and by the right truncation of our sample. Even if we wanted to measure time-to-next-release *conditional* on having a next release, toward the end of our sample period we only observe releases if the production time was very short. As a crude way of circumventing these problems, we instead ask the question, "For albums released in the current year, what is the average length of time since the previous release?" Table 7 shows these averages for the years 1995-2004, based on the Amazon data described above. Album production times do appear to have increased after 1999, but there appears to have been an upward trend even before the diffusion of file-sharing. The last column of the table suggests increases in production times after 1999 were largely attributable to a higher fraction of artists who were taking four or more years to record a new album.

5 Alternative Explanations

While many of the patterns described above are broadly consistent with our theory, there are at least two plausible alternative explanations. One is changes in the cost structure, which could be driven in part by the consolidation of concert promotion by Clear Channel. 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.

Table 7: Time between album releases

Year of Album Release	Average # years since previous album	Fraction of albums with lag: 0-1 years	2-3 years	4-7 years
1995	1.163	0.707	0.208	0.085
1996	1.318	0.674	0.246	0.080
1997	1.460	0.594	0.332	0.074
1998	1.623	0.536	0.373	0.091
1999	1.647	0.540	0.350	0.110
2000	1.637	0.563	0.304	0.133
2001	1.702	0.550	0.319	0.131
2002	1.757	0.532	0.328	0.140
2003	1.755	0.528	0.338	0.134
2004	1.832	0.507	0.333	0.160

Unfortunately, our concert data do not include information on costs, although we can get a rough estimate of average costs in the industry from trade journals. Trade journals indicate a slight increase in production costs. Venue costs may have changed more significantly during the period of analysis, due to changes in the structure of the market for concert promotion. Most notably, changes in the regulation of radio outlets allowed national radio station owners to enter the promotion business. Prior to this change, artists contracted with local promoters, who typically held exclusive dealing contracts with local venues. Thus, "Bill Graham Presents" held exclusive rights to promote concerts at the Shoreline Amphitheatre in San Francisco from 1966 until 2002, when it was acquired by Clear Channel Entertainment, which also owns many radio stations nationally. Similar mergers and acquisitions across the country have consolidated Clear Channel's national market share in the promotion industry. If demand across local markets is independent, theoretical predictions for concert pricing give ambiguous results for markets in which a series of local monopolists operate (such as Bill Graham Presents), versus markets in which a single national monopolist operates (such as Clear Channel). Of course, venue prices paid by the band may differ, and so the share of rents appropriated by the different agents may change.[22] Our data contain relatively rich information on venues, venue owners, and concert promoters associated with each event,

[22] If double-marginalization occurs because both the venue and the band apply a markup to ticket prices, increased venue costs could exacerbate this problem and lead to higher ticket prices.

so to some extent we will be able to assess the importance of market structure changes as alternative explanations for rising concert ticket prices.

Another alternative explanation is that there was a shift in the demand for live music, which would most likely have been driven by demographic changes. Evidence from tables 4 and 5 shows that the "conversion rate" of CD sales to concert revenues has the smallest increase for older bands, which are presumably the bands with older, wealthier fan-bases. In addition, older bands are not more likely to tour according to the results in table 5, which one would expect if demographic changes were swamping any other effects of sharing recorded music content. Nevertheless, demographic changes could certainly co-exist with other changes to demand in this industry–although a story based entirely on demographic changes seems unlikely given the patterns we see among younger bands over this period.

6 Discussion and Conclusions

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. 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. [23] 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.

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

Figure 2: Industry Revenue and Artist Revenue

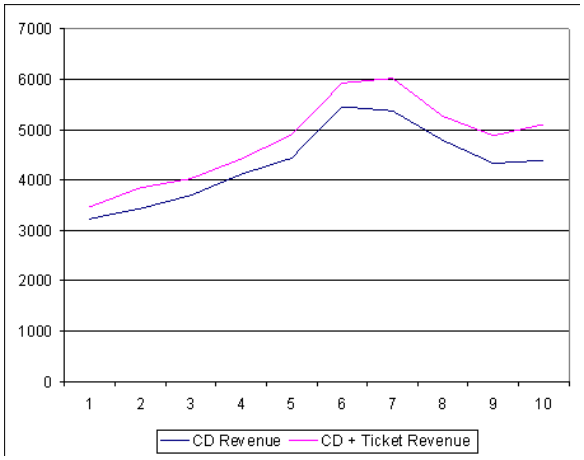

(a) Industry Revenue

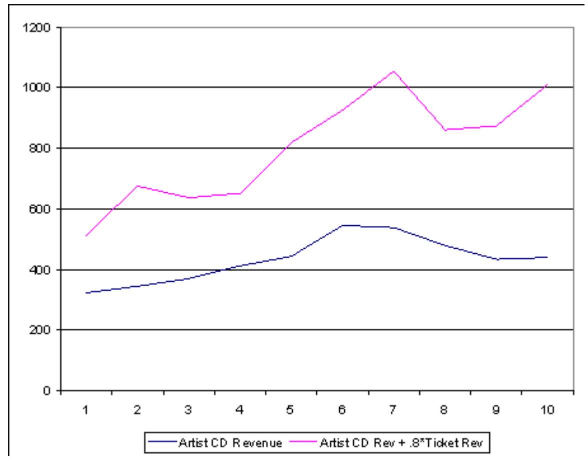

(b) Typical Artist Revenue

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 extension 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 after 1999 (see Table 1) suggests this may be true.

These kinds of questions certainly merit further exploration, since their answers will speak to the eventual impact of digital distribution on product variety and industry structure.

Bibliography

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*.

Appendix A

Here we provide detailed derivations of the first two predictions in Section 3. (The third and fourth predictions are straightforward.)

Prediction 1. The optimal effort level e_1^* is determined by the first-order condition

$$\phi(p_1 - c_1)\frac{\partial Q_1}{\partial e_1} + (p_2 - c_2)\frac{\partial Q_2}{\partial e_1} = \frac{\partial h}{\partial e_1} \ . \tag{5}$$

First, note that changes in γ have no direct impact on the cost-of-effort function, h. Assuming π is concave in e_1, $\partial e_1^*/\partial\gamma < 0$ as long as

$$\phi(p_1 - c_1)\frac{\partial^2 Q_1}{\partial e_1 \partial\gamma} + (p_2 - c_2)\frac{\partial^2 Q_2}{\partial e_1 \partial\gamma} < 0 \ .$$

Since $Q_1 = (1-\gamma)[1 - F_1(p_1; e)]$, $\partial^2 Q_1/\partial e_1 \partial\gamma = \partial F_1/\partial e_1$, which is negative by assumption. What about $\partial^2 Q_2/\partial e_1 \partial\gamma$? Note that

$$\begin{aligned}
\frac{\partial Q_2}{\partial\gamma} &= Prob(v_1 < p_1 \,\&\, v_2 \geq p_2) \\
&= \int_0^{p_1}\int_{p_2}^{\infty} dF^1(v_1, v_2; e_1) \\
&= \int_0^{p_1}\int_0^{\infty} dF^1(v_1, v_2; e_1) - \int_0^{p_1}\int_0^{p_2} dF^1(v_1, v_2; e_1) \ .
\end{aligned}$$

So

$$\frac{\partial^2 Q_2}{\partial\gamma \partial e_1} = \frac{\partial F_1(p_1; e)}{\partial e_1} - \frac{\partial F(p_1, p_2; e)}{\partial e_1} \ ,$$

which is negative by assumption.

Prediction 2. The optimal effort allocated to live performances (e_2^*) is determined by the first-order condition

$$\phi(p_1 - c_1)\frac{\partial Q_1}{\partial e_2} + (p_2 - c_2)\frac{\partial Q_2}{\partial e_2} = \frac{\partial h}{\partial e_2} \ . \tag{6}$$

Assuming π is concave in e_2, effort on live performances will increase with file-sharing (i.e., $\partial e_2^*/\partial\gamma > 0$) if

$$\phi(p_1 - c_1)\frac{\partial^2 Q_1}{\partial e_2 \partial \gamma} + (p_2 - c_2)\frac{\partial^2 Q_2}{\partial e_2 \partial \gamma} > 0\ .$$

The first piece of the above equation is in fact negative: $\partial^2 Q_1/\partial e_1 \partial \gamma = \partial F_1/\partial e_2$, which is negative by assumption. This reflects the idea that effort on concerts can boost album demand, but this effect becomes less important as album sales are displaced by downloads.

The second piece of the equation is:

$$\frac{\partial^2 Q_2}{\partial e_2 \partial \gamma} = \frac{\partial F_1(p_1; e)}{\partial e_2} - \frac{\partial F(p_1, p_2; e)}{\partial e_2}\ ,$$

which is assumed to be positive. Since the first (negative) effect is premultiplied by ϕ, it will be outweighed by the second (positive) effect as long as ϕ is sufficiently small—i.e., if the artist's share of album revenues is low enough, then the direct effect of increased concert demand will dominate the opposing effect of diminished incentives to promote album sales through concerts, and the net impact will be an increase in effort allocated to live performances.