

**Peer-to-peer networks: Creative Destruction or just Plain
Destruction?**

**Stan J. Liebowitz
School of Management
University of Texas at Dallas**

July, 2004

The sales of music CDs have fallen dramatically in the last few years. A leading candidate among possible causes of this decline is the recent rise in file-sharing. The recording industry has tried to stem this decline by suing or threatening to sue individuals heavily engaged in file sharing. These law suits have attracted a good deal of publicity and discussion. A rather influential school of thought has formed in opposition to the corporate players in these industries, and these lawsuits have provided fodder for their critiques of traditional copyright.¹

The motion picture industry has expressed concerns that its sales are likely to suffer a similar fate to that of the sound recording industry if nothing is done to stem the unrestricted use of file-sharing software. Owners of advertising-based mediums, such as television, are also concerned about their future viability if the advertising is stripped out of the programming and made available through file sharing networks, as is sometimes contemplated.

The very high profile of the music and movie industries has given the file-sharing issue a great deal of visibility. Although the revenues from these industries are only a small portion of GDP, these industries provide products that consume a large portion of our leisure time and tend to be recognized throughout the world.

At the center of the file-sharing debate is the empirical issue of whether or not file-sharing decreases sales.

Empirical work should always be informed by theory. In this paper I first explore some of the theoretical issues involved with file-sharing. I examine several theoretical factors that have been proposed to explain the behavior of file-sharers—substitution, sampling, network effects and indirect

¹ These copyright critics, who are sometimes associated with the concept of the ‘creative commons,’ argue that copyright laws are being used by the sound recording, movie, and software industries so as to thwart competitive forces that would open up the market to new competition. This is the thesis of Laurence Lessig’s recent book *Free Culture* which views the current controversies as extensions of long-running debates regarding the power of cartels to monopolize access to creative works. In this view of the world, file-sharing is a wealth enhancing innovation, likely to democratize the entertainment industry by allowing artists to broadcast and distribute their works without intermediaries such as record companies. In this view, file-sharing systems should be promoted and if necessary, copyright law should be altered to allow file-sharing to proceed apace.

appropriability. One of the novelties of this analysis is the conclusion that sampling, the behavior most relied upon by those suggesting that file-sharing might have a benign impact on the industry, would actually have a negative impact. I conclude that of these factors, only network effects might have a positive impact on sales but suggest that this is unlikely to be an important factor. Thus economic theory would lead to an expectation that file-sharing should lead to a decline in sales relative to what they would have been in the absence of file-sharing.

I then turn to different methodologies that have been chosen in empirical examinations of this issue and I review these studies in some detail. I conclude that these studies tend to confirm what the theory predicts. Although some studies reach different conclusions, those studies appear to suffer from weak methodological choices or are based on questionable empirical specifications.

I. Background

Sound recordings, movies, and television occupy the large bulk of our time spent on leisure activities, with the average American watching four hours of television and listening to three hours of music each day.² The advertising industry is primarily based around these industries, as is the consumer electronics industry, which would include all forms of stereo equipment, televisions, DVD players, VCRs, and so forth. Thus even beyond its share of GDP, its share of the collective consciousness and behavior of consumers is very large. It is also possible that its value to consumers might be much larger than that indicated by market revenues.³

The products of these entertainment industries have proven amenable to digitization, allowing them to be transferred over peer-to-peer file-sharing networks. Napster was the first well-known peer-to-peer file-sharing system, but others have followed in the wake of the preliminary injunction that

² See Table No.909 "Media Usage and Consumer Spending: 1993 to 2003" in the 2000 US Statistical Abstract.

³ Paul Romer conflates revenues and value in an otherwise thoughtful paper "When Should We Use Intellectual Property Rights?" *American Economic Review*, May, 2002, 213-216. He states: "The stakes in the battle over the music business are small enough to get lost in rounding error for world GDP of about \$30 trillion. However, this battle creates a "teachable moment" that could help frame policy in more important areas." I would suggest that the interest in this topic reflects something more important than just a 'teachable moment.'

effectively shut Napster down. Current replacements have surpassed Napster in popularity. Some Napster replacements are amazingly popular. As an example, Kazaa has been downloaded 350,000,000 times, Morpheus has been downloaded 125,000,000 times and iMesh has been downloaded 75,000,000 times from a single website.⁴ It is estimated that file-sharing represents over one third of all material transferred over the Internet, and that music files are downloaded to the tune of .8 billion files per month.

In what follows I will focus on the sound recording industry because that is that market that has attracted the most attention. Since most computer users have enough bandwidth to download MP3 files and also have in place the requisite CD burners with which to allow the listening of this music in locations not tied to a computer, this is the arena where most file sharing is taking place. When bandwidth allows the downloading of high quality video files to take place in a matter of minutes, instead of the current length of hours, and when DVD burners become ubiquitous as are the CD writers, then movies, and to a lesser extent television, may become equally subject to file-sharing.

II. Economic Theory of File-Sharing's Impact

One topic that has received too little attention in the recent literature is the theory underlying predicted impacts of file-sharing on the marketplace.⁵

One set of theories largely mirror the claims made in the Napster case by the opposing sides but there are additional theories as well. On the one hand, a downloaded file can substitute for the purchase of an original CD or single song. Substitution of a free alternative is easily understood to have a negative impact on sales. On the other hand, the claim has been made that users might merely use downloaded songs to become more familiar with potential music. Although this was originally referred

⁴ On the download.com website, as of mid May 2004.

⁵ For discussions of the economics of copying and copyright see Watt (2004) or Varian (forthcoming).

to as the *exposure* effect,⁶ it is currently called the *sampling* effect. Under this scenario users sample from available music and then purchase those songs and albums that are found to be most suitable to matching the tastes of the users. This sampling hypothesis is usually associated with a claim that sales will increase if consumers are allowed to become more familiar with the product before they purchase it, although, there has not been much analysis of this claim.

There is also a claim of potential network effects. As more downloaders listen to music, this theory goes, consumers derive greater value from their legitimate purchases. It is suggested that this might lead to an increase in the sales of CDs.

Finally, there is a possibility that sellers of original files can capture the value from later copiers indirectly in the price of originals, a concept known as indirect appropriability..

I examine each of these claims in turn.

A. Sampling

Although it is often asserted that sampling would obviously be expected to have a positive impact on sales, the impacts of sampling are far more subtle. Indeed, a more complete analysis tends to lead one to expect that sampling would lead to a decrease in sales in this market.

The sampling story basically argues that file sharing allows consumers to experience music in a more complete manor prior to purchase than they would have been able to do were they to use the more traditional methods of learning about music: hearing it on the radio or at a friend's house. With file-sharing, listeners can become as familiar with the music as they wish, listening to it over and over again until they are sure they like it.

At that point, according to the sampling theory, the listeners go out and purchase the music through traditional outlets. A natural question is to ask why they would make a purchase when they

⁶ See Liebowitz 1981, 1985.

already have the item for free. There are several possible answers. First, they might be uncomfortable listening to music which they have not purchased. This discomfort might arise from a sense of honesty or a sense of wishing to support their favorite musicians. Alternatively, listeners might get to know three or four songs on an album which then allows them to feel comfortable buying the entire album and avoid the efforts involved with downloading another 6 songs.

Assuming that sampling occurs in the manner described above, what would be the likely economic impacts of sampling?

Assume that those engaged in sampling have no intention of listening to MP3 files after the sample period. Instead, they either purchase the music or throw it away. This is a pure analysis of sampling independent of any pirating motive.

I am not aware that any proponents of the sampling story have attempted to fill in the details of the pure impact of sampling.⁷ Presumably, after sampling, consumers have more information about which CDs to purchase allowing them to purchase CDs that provide greater utility than they would have purchased without sampling. Although it is natural to think that consumers would be led to purchase more CDs if CDs can provide greater utility than they did without sampling this is not necessarily the case.

To see this it helps to think of a CD as a candy bar, following a line of reasoning developed by Jack Hirshleifer (1971).⁸ Each individual consumer has particular tastes in music and some CDs are

⁷ The closest might be a paper by Gopal, Bhattacharjee, and Sanders (forthcoming). They attempt to analyze the theoretical impacts of file-sharing using a fairly typical model. Sampling plays an important role in their model but they do not analyze the impacts of sampling by itself. If the full cost of sampling were zero, consumers in their model would sample all music, in order to find the most highly valued music. Whether they would then purchase the preferred music depends on other costs, such as costs of punishment if caught pirating, the sound quality differentials between sampled and purchased music, and the revealed value of music. If sound quality was identical between original and copies and if there were no punishment if copying, consumers would completely pirate and purchase no music.

⁸ By working with the underlying characteristic of the good we can avoid the heterogeneity problem brought about by the fact that CDs are not perfectly homogenous. This is done by assuming that CDs are differentiated by containing different quantities of the underlying music-listening service characteristic. This keeps the analysis simple. It also allows the analysis to treat the indirect valuation of copiers as altering the quantity of music services created by the original as an analog to increasing durability of a product. I provide this type of analysis in Liebowitz (1982a).

better than others as far as satisfying these tastes. Consumers, after all, do not derive utility from the CDs *per se* but derive enjoyment from listening to the music contained within the CD. The underlying demand can be thought of as the demand for music-listening services, which is met to differing degrees of success by various CDs. Those CDs that better satisfy the consumer can be thought of as providing more music-listening services within the fifty or so minutes of music contained within the CD. Since those CDs contain more of what the consumer wants, they can be analogized to providing consumers larger candy bars containing more of the candy which the consumer ultimately desires.

An equivalent analysis is to think of sampling as reducing the risk associated with the purchase of a CD. If consumers are not sure whether they will like the music on a CD they will consider it a somewhat risky purchase. Sampling has the possibility of reducing or removing this risk. What this means is that sampling allows consumers to achieve greater expected utility with a CD purchase than they would have achieved without sampling. At a given price of CDs, the cost of the expected satisfaction gained by consumers falls. Again, it is equivalent to putting more expected yummy music services in each CD.

It is natural to think that if candy bars remain constant in price while increasing in size, that the quantity of candy bars sold will increase. After all, each large candy bar provides more utility than a small candy bar. This is apparently the thinking of those claiming that sampling MP3s increases the sales of CDs.

The story, however, is not so simple.

First note that the price of candy is effectively lowered when the bar becomes bigger holding the price of the candy bar constant. If the demand *for candy* is elastic, then revenue in the market will increase when candy goes down in price, as it does when constant-price bars become larger. If revenues increase in this way, and the price of bar is unchanged, then more candy-bars are purchased. The inverse of this story will hold when the demand for candy is inelastic. Making candy bars larger

rotates the demand curve so that it is higher at small quantities but lower at large quantities—in other words, the demand curve becomes steeper. Satiation occurs at a smaller number of bars since each bar is bigger.

A simple analogy can be made to the introduction of cable television. Cable allowed viewers accustomed to having a choice of only a handful of broadcast signals to have a choice of dozens of channels. This should have increased the probability that viewers would find, in any half hour period, a program more to their liking than they were likely to have found with only the limited original choices. The analogy to the sampling hypothesis is very strong since allowing greater choice in a thirty minute television time slot is similar to providing information allowing greater choice in the selection of a CD.

Previous research examining the impact of increased viewer choice on total time viewing television tended to find that that providing more choice to consumers did *not* increase the time they spent viewing television.⁹ Thus the claim that providing consumers the ability to fine-tune their product selections need not increase their consumption is seen to have real-world explanatory power.

Understood in this way, it becomes clear that better information about the product might lead to either an increase or a decrease in the number of units sold. Equivalently, it might lead to either an increase or a decrease in revenues. This was a conclusion I had reached in previous work (Liebowitz 2004). There is, however, some additional information in this market that helps to resolve this imprecision.

CDs are thought to have low variable costs of production and high fixed costs. It is common in theoretical models of markets like software or music to assume a zero marginal cost of production.

⁹ See Liebowitz (1982) who compares the link between viewing hours and cable penetration across different Canadian metropolitan areas and finds an insignificant but sometimes negative relationship. Also see Weimann (1996) who examines viewers in Israel after the introduction of a multi-channel cable system where previously there had been but a single public channel (a more extreme increase in choice than would normally be found). After a year, there was virtually no difference in changes in viewing between a group with cable and a control group that did not receive cable (the cable group increased its viewing by 16 minutes over the control group).

Although this is merely a theoretical convenience, since the variable costs are clearly not zero, variable costs do appear to be quite low in the case of sound recordings.

What are some of these variable costs? The cost of a blank CD is only a few cents and the pressing of CDs appears to be less than a dollar. Although the artists normally receive a royalty that is expressed as a function of sales, those payments are usually paid up-front as a non-refundable advance against future royalties, so for most units sold, marginal royalties to the performers are effectively zero. Promotional costs for CDs are also usually also taken out of up-front advances, removing another potential variable cost from the variable cost column. There is a variable payment made to the composers of songs that are included in the CD, however, with a statutory maximum payment of approximately seventy cents per CD.¹⁰

It seems reasonable to conclude, therefore, that variable costs are quite low relative to the wholesale price of CDs, which is in the vicinity of twelve dollars.¹¹

This datum of low variable costs provides some important additional information about elasticity of demand facing each CD. Profit maximization, when marginal costs are zero, is equivalent to revenue maximization. The elasticity of demand must be one at the profit maximizing price chosen by sellers of CDs.

The elasticity of demand for CDs maps directly into the elasticity of demand for music-listening services. If the price elasticity of demand for CDs equals one, so too must the price elasticity of demand for music-listening services. After all, if the marginal revenue of another unit of music service is negative, so too must be the marginal revenue of the CD containing that unit.

¹⁰ This is a compulsory license which amounts to ninety cents for each CD. In the common case where the performer is the composer, however, it is typical for there to exist a “controlled composition” clause which pays less than (75% appears to be a typical rate) the statutory payment.

¹¹ If the composer is not the performer of the song, and if the song has not been recorded, this payment will be negotiated. After a song has been recorded by one performer, however, anyone can record the song if they purchase a ‘compulsory license’ which is the statutory payment described in the text

¹¹ According to statistics on the RIAA web page, the average price in 2003 was \$11.91.

Because there is competition between record titles, we should expect that the elasticity of demand for music-listening services (whether defined by musical genre or the entire industry) will be less than the elasticity of demand individual firms or individual record titles, for the same reasons that industries have lower elasticities than their constituent firms. This implies that the elasticity for the industry will be less than one.

It matters little, however, whether the price elasticity is less than one or equal to one. In either case revenues (and profits) fall when the price is lowered.¹² The effect of sampling (more music-listening services at a constant CD price) is to lower the price of music-listening services. The net effect should be to lower the revenues generated by music-listening services. With a price per CD that is independent of the sampling effect, this implies that the quantity of CDs will fall due to sampling. This analysis assumes, as does the candy bar problem, that the price of CDs is exogenous to the change in music-listening services created by CDs.

This is quite the opposite of the claim that is most frequently made for sampling.

B. Substitution Effect

The substitution effect is quite simple to analyze. The copy is treated as a substitute for the original. If the copy is identical or close in quality to the original, and if the cost of making the copy is low, the copy for a price of zero dominates the original at its positive price.

This is most appropriately modeled as a decrease in demand. The substitution effect can only work to reduce the price and quantity sold in the market.

¹² Boldrin and Levine (2003) assume that elasticity is greater than one in order to generate their results, discussed below.

C. Network Effects

Some products have network effects. These occur for a product when consumers' values of the product change depending on the number of other users there are of the product. Telephones and fax machine are two examples of products where the value of those products depends on the number of individuals using those products.

It has sometimes been claimed that network effects might be important to understand the impact of copying. Conner and Rumelt (1991), Takeyama (1994) and Shy and Thisse (1999) each examine models where the existence of unauthorized users creates additional value to the purchasers of legitimate copies and thus might increase the profits of the seller. These models are usually put forward in the context of software, although they might also be applied to file-sharing.¹³

There are several issues to be addressed in the context of possible network effects for sound recordings. The first is whether there are network effects at work. Contingent on network effects existing for sound recordings, a different issue is whether those network effects work merely to shift demand among different sound recordings or whether they work to alter the size of the market. A third issue is the nature of the impact of network effects—are they local or global—does the totality of usage matter, or just the number of individuals using the product who personally know one another?

The network effects story applied to file-sharing is relatively straightforward. File-sharing is likely to increase the number of music listeners since file-sharing provides access even to those who are unwilling to pay for it.¹⁴ If there are more listeners to the music, even if the additional users were

¹³ It is unclear how strong network effects are for different categories of software. For some categories, such as personal finance software, network effects might be expected to be weak or nonexistent. In other cases, such as spreadsheets network effects are thought likely to be large. Although there have been attempts to measure the strength of network effects for spreadsheets, these attempts were marred by using Lotus 1-2-3 file compatibility as measure of network effects when such compatibility was also important for upgrading spreadsheet users wanting to remain compatible with their old work independent of any network effects.

¹⁴ It is possible that listening to file-shared music might merely replace listening to the radio or listening to purchased music. File-sharing, however, does allow more precise choice of music than could be achieved by radio listening.

merely among the file-sharers, the value of music for non-file-sharing individuals might increase and the non-file-sharers would then purchase more music.

Before we accept this possibility of network effects, however, we need to understand possible mechanisms that might lead to network effects. The question is: do consumers derive value from other individuals listening to the same music they listen to? Unlike telephones, where network effects are obvious, or software, where the ability to transfer files might be important, the linkage of values between different music listeners is far less clear.

i. Local or Global Network Effects?

As normally modeled, network effects depend on the number of other users of the product. This would be a global network effect where only the total number of users counts, not the identity or locations of specific users. If you are in the habit of selling products over the telephone, network effects from telephones might well be global—you don't care who or where your customers are, as long as there are a lot of them.¹⁵

On the other hand, most potential network effects are likely to be local—you care mainly about calling certain individuals, you plan to transfer files with a few particular colleagues, you discuss movies and music with your friends. In these cases the total number of users of a product is largely irrelevant to you.

The first question is whether the activity of music-listening is subject to either local or global network effects to any important extent. This is a difficult question to answer. When a musical composition/performance becomes popular it might be due to some innate quality within the music, combined of course with some form of information, be it radio play, newspaper and magazine reviews, or word of mouth. On the other hand, it might be due some self-reinforcing network effect.

¹⁵ It is likely that for indirect network effects, where the price of some secondary item (such as repair service) is impacted by the number of users, that the identity of the users doesn't matter much and that the network effects are global. In the case of file-sharing, however, network effects, to the extent that there are network effects, would seem to be of the direct variety.

If network effects are global in nature then it isn't clear how file-sharing could impact sales since file-sharing doesn't usually show up in global measures of music popularity, which are normally based upon purchased CD sales or airplay.¹⁶ If consumers care about the total universe of users they can only estimate the number of users from these publicly available statistics.

Local network effects, on the other hand, suggest that CD purchasers interact with friends who engage in file-sharing and then are influenced by their friends file-sharing musical tastes.

There is also the question of whether network effects, if they exist for musical works, have the impact mainly of shifting output from one musical composition to another or whether they have the impact of changing the overall size of the market. Normally, network effects are thought to increase the value of products as networks get larger. Products have some autarky value independent of network effects which can then be enhanced as the size of the network increases, as would be the case for computers and the Internet. In this instance network effects can only increase the value of products and the size of markets. It is possible, however, that network effects can lower utility below the autarky value—e.g., an unpopular restaurant that loses patrons who do not want to be seen as 'uncool'. In this latter case network effects need not increase the overall value or size of markets. If network effects work mainly to allocate utility from one CD to another depending on shifts in popularity, but do not increase the overall utility of CDs in general, then they can have a neutral or negative impact on overall sales.

ii. The Nitty Gritty of Network Effects in Music

What might be the nature of network effects in music? Although it is almost certainly a common event that one person hears music at a friend's house and decides to buy it, that is not sufficient for a network effect to exist. Just hearing new music is merely a form of sampling. For it to be a network

¹⁶ Websites such as BigChampagne.com purport to measure popularity of downloads (although it is apparently uploads that they measure) and it is possible that such measures might become part of the public information about popularity.

effect, the value of the music once it has been heard must be higher because friends like it than would otherwise be the case. Music that is part of group activities, such as dancing, are good candidates for network effects since the value of one person likely depends on the enjoyment of other persons. Even so, music listened to as a part of group activities implies that one of the group members already has that music. Whether the other members are likely to purchase the music themselves to enhance group consumption would seem to be questionable. In the case of file-sharing, a network effect would exist when a person who doesn't file-share purchases music that his file-sharing friends can enjoy when they come over to visit. These do not appear to be powerful network effects, so as a practical matter, therefore, network effects from file-sharing would seem likely to be weak.

Although it is conceivable that network effects could play a role in the overall sale of CDs, and it is conceivable that file-sharing could increase overall usage sufficiently that CD sale might increase, it seems unlikely to be more than a minor impact.

The only test of network effects on record sales that I am aware of is indirect—looking at the impact of radio on record sales. Again, network effects from radio listening need to be separated from the sampling effects from radio, which is a more obvious impact

One might argue that radio play is more likely to exhibit global network effects. Global network effects are possible with radio since radio play is one form of communicating what 'everybody' is listening to. If individuals derive value from being part of the crowd, then there would be global network effects from radio play. Second, local network effects are also possible since groups can sit around listening to the radio and commenting on the songs. This might increase sales for those individuals wishing to make their friends happy in exactly the same manner as it might for file-sharing.

Undoubtedly, radio play influences the sales of records. Its main impact, however, seems to be to shift sales from one recording not receiving airplay to other recordings that do receive airplay. This conclusion is reached in Liebowitz (2004a) who examined two instances where radio broadcast did not

increase overall record sales. In the first instance, the introduction of radio in the US in the 1920s, the sale of sound recordings fell dramatically. In the second instance, the introduction of private radio stations in the UK during the last three decades of the twentieth century led to no increase in the sales of records. Although these results were not meant to be tests of network effects *per se*, they are certainly inconsistent with the possibility of strong network effects.

D. Indirect Appropriability

The final impact of copying that might apply to file-sharing is indirect appropriability. This is a concept coined in Liebowitz (1985) and analyzed for the case of file sharing in Liebowitz (2002) which I summarize below. It has recently been brought up by Boldrin and Levine (2003) whose work was then critiqued by Klein and Murphy (2003). The basic idea is that originals from which copies are made might undergo an increase in demand as those making copies of originals capture some of the value from those receiving the copies and transfer this value into their demand for the originals that they purchase.

If, for example, everyone who purchased a CD made one cassette to play in their automobile, then the demand for the original CD would increase by the value of being able to make the tape and the sellers could capture some of this higher value by increasing the market price, as would happen when demand increased. This value is captured indirectly since there is no direct payment made for the copy.

In order for indirect appropriability to work, however, one of two conditions must hold. First, the variability in the number of copies made must be small, as in the example above. Or else the seller needs to be able to charge higher prices for those originals from which the most copies are made as in the real world example of photocopies. The most heavily photocopied copyright materials are journals, and most photocopying of journals takes place in libraries. Thus publishers of journals were/are able to charge higher prices to libraries than they charge to individual subscribers to take account of

photocopying and indirectly appropriate some of the value from copying. Support for this thesis came from empirical work revealing that such price discrimination was practically unheard of prior to the advent of the photocopier, and the most heavily copied journals were also those with the greatest price differential.

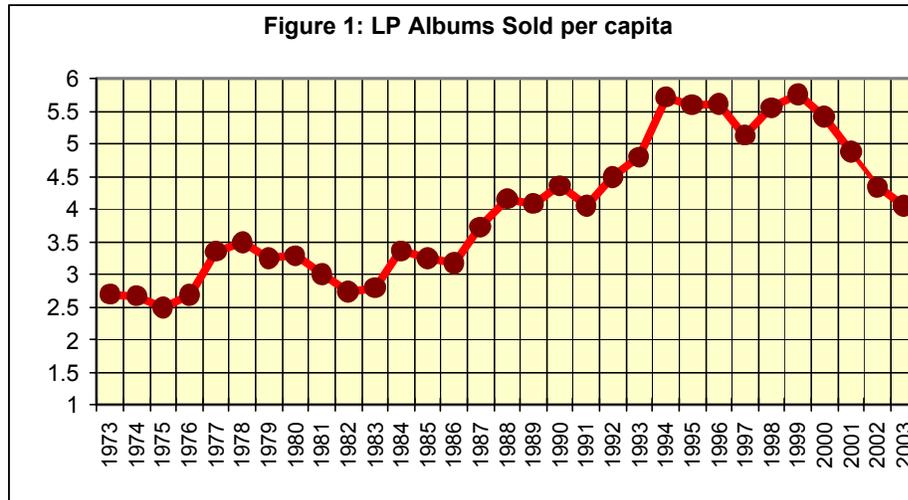
Contrary to the claims of Boldrin and Levine, indirect appropriability will not work for file-sharing, however. Because there is great variability in the copies made from each original, and sellers of originals cannot identify which originals are going to be used on file-sharing systems, the mechanisms that allow indirect appropriability to function will not work.

III. The Prima Facie Case against File-sharing

Data on the sales of recorded music tend to be available on a yearly basis from organizations of record producers, such as the Recording Industry Association of America (RIAA) and the International Federation of the Phonographic Industry (IFPI). The United States is by far the largest single market.

Figure 1 represents the per capita sale of full-length albums sold in the United States since 1972.¹⁷ The recent decline in sales is readily apparent. Clearly, there are yearly fluctuations in this series. These statistics indicate a somewhat smaller decline than usually suggested by the industry since I have stripped out the impact of ‘singles’ because their decline appears to be part of a much longer secular decline quite separate from file-sharing. Nevertheless, the current decline stands at thirty percent.

¹⁷ Data on quantities were reported beginning in 1973. For prior years only the industry revenues were reported. The data on revenues tend to be estimates based on the list price of recordings. I use full-length albums to avoid measurement problems as, for example, with singles, as they have been in a twenty year secular decline.



This recent decline is clearly the largest that has occurred in the last thirty years. In absolute terms (units per capita) it is more than twice as large as the next closest decline and in percentage terms it is nearly fifty percent larger than the next largest decline, which may also have been impacted by copying.¹⁸ This is sufficiently striking that it would appear that something unusual has occurred in the last few years. Such a large change would be caused by either an unusually powerful but established factor, an unusual confluence of established factors, or some completely new factor.

Napster, which was the forerunner of modern file-sharing, came into existence in 1999, which also happens to be the peak year in sales. At the time, CD burners were still relatively expensive and did not yet have a large market penetration. Although Napster was closed down in 2001, file-sharers soon migrated to other file-sharing services.

The question, in a nutshell, is whether the decline in sales that began in 2000 is due to file-sharing. This is the file-sharing hypothesis.

It is incumbent on us to respect the facts that we have. The confluence of file-sharing's birth and the decline in record sales, the rapid growth to an immense size of file-sharing activities, and the

¹⁸ The second largest decline occurred in 1978-82. Long playing vinyl records were in decline at that time, being replaced by prerecorded cassettes. This decline coincided with a fairly deep recession and the possibility that blank cassettes were being used to copy records. As I explain in Liebowitz (2004), cassettes ultimately allowed, for the first time, portability of prerecorded music, leading to a large increase in sales that overcame any negative impacts of copying.

unusually large decline in the sound recording market all support a claim that file-sharing is responsible for the decline in sales. Add to this the predictions of economic theory that file-sharing should lead to a decline in sales and we have what appears to be a very strong case that file-sharing is the cause of the decline.

Economists, however, are trained to eschew straightforward answers. Nevertheless, the file-sharing explanation for the decline in sales is sufficiently compelling that only extremely powerful empirical evidence to the contrary should allow us to dismiss the file-sharing explanation. As the following review of the evidence shows, compelling and consistent evidence against the file-sharing hypothesis has not been forthcoming.

IV. Economists to the Rescue?

If we could run controlled experiments it would be easy to determine the impact of file-sharing. Because we cannot run controlled experiments and must use the less than ideal data at hand, we might not be able to answer this question in a convincing manner. Various analyses have been undertaken using different data and different approaches, and more are sure to come.

There are some preliminary points that I should mention before discussing these approaches in detail, however.

First, although measuring the impact of file-sharing on the sound recording industry provides important information, it is only part of the analysis that would be required to answer the question about file-sharing's impact on social welfare or even its impact on the industry's ability to appropriate value. Even if it were the case that file-sharing had no effect on recording industry revenues, it might still have a large impact on appropriability. In the simplest case, assume that half the population purchases CDs and that these individuals do not change their behavior in the face of file-sharing. Then, due to file-sharing, the other half of the population listens to music downloaded from file-sharing sites.

The producers appropriate none of this additional value created by their product, so appropriability, which can be defined as producer surplus relative to total surplus, has fallen.

The new consumption of music by this latter half of the population is an additional surplus for society, but that doesn't mean that this change is optimal.¹⁹ Total surplus is maximized not when static consumer surplus is maximized for a fixed output, but when the sum of consumer and producer surplus is maximized in a system where producers receive payment so as to induce greater output. Some level of appropriability other than zero is usually optimal and it is likely that efficiency would require that some revenues be generated from file-sharers.

Nonetheless, economists have limited themselves to the more prosaic question of whether or not file-sharing has decreased revenues to the sound recording industry. This is a useful first step, but we need to remember that it is only a first step.

A. Examining Alternative explanations

One straightforward method to help determine whether the prima facie argument is correct is to examine other possible explanations for the decline in record sales that has occurred. Although this indirect form of testing might appear unusual, it is not. The large literature on discrimination is almost entirely based on indirect testing—differentials in black/white wages, for example, that cannot be explained by other factors are attributed to discrimination. This approach has been widely accepted within the profession.

Examination of other factors that might impact record sales is the approach taken in Liebowitz (2004). Liebowitz addresses four main alternative factors that might impact record sales:

- a. Price or Income Changes
- b. Changes in Substitute/Complement Markets

¹⁹ I hesitate slightly to call it 'consumer's surplus' since these individuals do not pay a positive price and do not fit the normal definition of a paying consumer.

- c. Changes in the Quality of Music
- d. Changes in the Supply of music.

I address these points in turn.

i. Price and Income changes

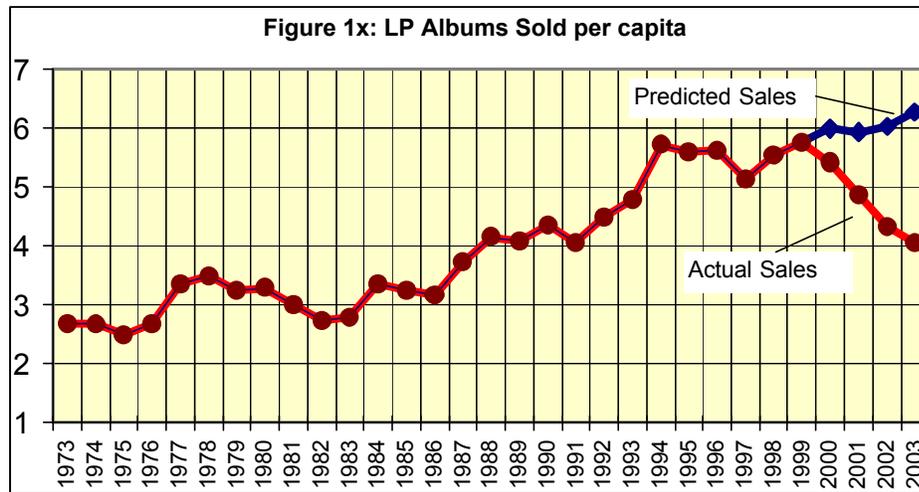
Although the claim is often made that price increases are responsible for the decline in CD sales the evidence does not support this view. Real list prices have been virtually constant for the last decade, and the increasing share of the market generated by discount sellers has probably lowered the average transaction price. Prices, therefore, are ruled out as an explanation of the sales decline.

We are interested in the reduced form estimate of the quantity of recordings sold. Regressing first differences in albums per capita on real per capita GDP and movie box office revenues per capita over the period 1973-1999 gives the results found in Table 1.²⁰ Movie box office receipts appear to have little relationship to album sales. The impact of income on sales over the last thirty years is statistically significant, but a change in per capita disposable income of \$1000 would alter per capita sales by only .28 units and the largest sustained decline in per capita GDP (1982) during the last thirty years was approximately \$661. Although the US experienced a recession in 2001, it was far too mild, given the historical impact of income on sales, to be able to explain the decline in sales that occurred.

Table 1: 1973-1999				
	Coefficients	Std. Error	t	Sig.
(Constant)	-0.0269604	0.089802	-0.3	0.767
Per Capita Change in GDP	0.0002794	0.0001245	2.25	0.035
Change in Movie Box Office	0.0421614	0.059508	0.71	0.486
R Square	Adjusted R Square	Dependent Variable: Change in Albums per Capita		
0.2122	0.1437	N=26		

²⁰ It is obviously inappropriate to use the post 1999 data in these regressions since file-sharing started in 1999 and we wish to determine the impact of income and movies on record sales in order to judge whether the post 1999 sales were negatively affected by file-sharing. If file-sharing had an impact on sales, it would distort regression results ignoring file-sharing. If we use the entire sample period with a file-sharing dummy for the years after 1999 the dummy is significant with a value of -.54.

In fact, since 1999 real GDP per capita has increased by \$1883, which, according to the regression, should have led to an increase in sales of .52 units per capita. The year-by-year sales figures after 1999, based on changes in income, are provided in Figure 1x as the line with diamond markers. Of course, changes in GDP are not capable of explaining much of the prior declines so there are obviously other important factors at work.



ii. The Impact of Substitutes, such as the DVD Market

One claim that is frequently encountered is the claim that record sales are down because alternatives, such as videogames, movie viewership, and DVDs are up.

We have already seen that there doesn't appear to be much of a relationship between movie box office receipts and record sales. Table 2 lists correlations between yearly changes in the sound recording industry and changes in these other industries. These forms of entertainment seem to be impacted by similar factors since yearly changes in each are positively related to one another. In some instances the number of observations is quite small, reducing our confidence in the applicability of these measurements.

Table 2: Correlation of change in per capita record sales with change in		
movie revenue per capita	0.199	1973-1999
videogame revenue per capita	0.265	1991-1999
units of prerecorded video per capita	0.177	1990-1999

Whether these positive correlations indicate that these products are complements with one another or whether they are due to a simultaneity problem we cannot be sure, although income appears to have different impacts on these markets.²¹ Certainly, videogames are often based on movies (and occasionally vice versa) and soundtracks of movies also sell very well, but it is difficult to believe that these complementarities are prevalent. Nevertheless, just as we saw that movies had an insignificant impact on record sales in Table 1, the positive coefficients reported in Table 2 should not be taken to indicate much of relationship between these variables. The point, of course, is that there is no evidence that these products have *negative* impacts on one another.

Another way of looking at this issue is to examine the average usage of time taken up by these activities. Published estimates indicate that in 1999 listening to recorded music (as opposed to radio) took approximately 48 minutes of a person's time per day, whereas going to movies took 2 minutes, watching prerecorded movies took 9 minutes, and playing videogames took 10 minutes.²² If these measurements are at all accurate then the relatively low time usage of prerecorded movies or videogames does not appear capable of usurping much of the time spent listening to prerecorded music.

Since theatrical movies take up a trivial amount of time, and videogames do not preclude listening to prerecorded music, we can focus on prerecorded videos which have grown at a rapid rate since 1999 largely in the form of DVDs. We can take the most extreme case possible. Assume that each minute of watching prerecorded videos replaces a minute of listening to prerecorded music (given

²¹ The correlation of income changes with changes in movie revenues is .06, with videogames is .41, and with prerecorded videos is -.37.

²² US Statistical Abstract, Table No. 1102. Media Usage and Consumer Spending: 1996 to 2005. <http://www.census.gov/prod/2003pubs/02statab/infocom.pdf>. Original source is Veronis, Suhler, Stephenson.

the much larger potential time-targets of television and radio listening, this is quite an extreme assumption). How much of a decline would this have caused to the sound recording industry?

year	VHS	DVDs	sum	per capita
1999	561	91.3	652.3	2.390954
2000	565	174.4	739.4	2.685739
2001	539.6	350	889.6	3.196482
2002	407.5	650.6	1058.1	3.761384
2003	240.4	985.3	1225.7	4.3112

The growth in per capita prerecorded video was 80% from 1999 through 2003. An 80% increase would raise the time spent watching VHS tapes and DVDs by 7.2 minutes per person per day. Reducing time spent listening to prerecorded music by 7.2 minutes, if it translated one-to-one into lost sales, would reduce sales of sound recordings by 15%. Yet per capita sales of sound recordings fell by 30% during this period of time. Thus, even in the most extreme case imaginable, a time substitution could have caused no more than half of the actual decline, and this is an implausibly large estimate. Further working against this hypothesis is the fact that the decline in record sales in 2003 was only half that of the prior two years whereas the increase in prerecorded video was as large in 2003 as the prior two years.

We must also remember that record sales might have been expected to increase after 1999, as discussed in the last section. This stretches even further the credibility of these alternative explanations.

Although this is an area for further research, it seems most unlikely that under more precise measurements these variables could have the type of negative impact required to explain more than a relatively small portion of the fall in record sales.

iii. Has the Music Changed?

When the record industry has a decline, critics of the industry will often claim that sales fell because the music was not interesting. Thus, it is not surprising that this has emerged as an answer to the question: “if MP3s are not hurting the industry, what is?”

I looked at two pieces of evidence related to the ‘quality’ of music. One was the financial success of concerts from 1990-2001. Although there are defects with this measure, the years 2000 and 2001 had the largest real increases in concert revenues—at the same time the record industry was experiencing unusually large decreases in revenues.

Group	Decline in Listening
Persons 12+	9.30%
Teens 12 - 17	11.67%
Persons 18-34	13.48%
Persons 25-54	8.79%
Persons 55-64	6.74%
Persons 65+	5.75%

A second potential source of data is time spent listening to radio. If music is losing its luster, then radio listening should decline. Data from the last 5 years, shown in Table 4, do reveal a fairly serious decline in listening—9%. The decline, however, is across the board, not just for groups that listen to new music. It is implausible that those over the age of 65, for example, would find their attraction to radio diminished because modern musical compositions were not to their liking.

Further doubt about any claimed doldrums in current music composition comes from examining the listening shares of radio formats. If current music were inferior to past music, its share of listening should fall relative to old music and non-music categories. Yet the category of ‘News and Information’ is only up 1.8% over this period. The largest decline in radio listening (51.3%) is for the ‘Adult Standards’ category which consists of music from the pre-rock era, hardly a repository of new music.

The categories of ‘Contemporary Hits’ and ‘Urban,’ both of which have a great deal of new music, each increased their shares (5% and 23% respectively) and ‘Alternative’ (modern rock) was down only 2%. It is true that ‘Country’ music had a 9% decline, but this was largely offset by an increase of 43% in the ‘Religious’ category (which contains two categories of contemporary music).

The evidence on musical quality is far too weak and scant to explain the very large decrease in sales that has occurred.

iv. Changes in the Supply of Music

It has been claimed by some that the number of new titles has fallen in recent years and this decline in new titles is responsible for the decline in sales.²³ Of course, the number of new releases is not exogenous and we would expect the quantity of releases to fall if demand fell because of file-sharing. Unless we are willing to entertain the possibility that the supply of songwriters and musicians has suddenly dried up at the same time that file-sharing has started, we would expect that changes in demand are more likely to influence the number of new titles than would changes in underlying supply conditions.

Further, this discussion has suffered from an imprecision in the measured numbers of new releases.²⁴ One factor that is clear, however, is that the large majority of new releases, between 65% and 80%, are from independent record labels which account for merely 10-15% of industry sales. Looking at industry-wide new releases overweights independent albums at the expense of major labels and is like having the tail wag the dog. Thus the total number of new releases is a misleading statistic as regards the overall sales in the industry and really cannot be a useful barometer of supply even if it

²³ See for example George Ziemann "RIAA's Statistics Don't Add Up to Piracy" available at <http://www.azoz.com/music/features/0008.html>. The analyses contained on these pages are probably best described as ‘rants’ but they have nevertheless been taken seriously by some. Understanding the differences between retail sales volume measured by SoundScan’s barcode reading technology and wholesale shipment numbers reported by major labels would be a useful exercise, but instead the difference tends to be treated in these analyses as some sort of industry conspiracy.

²⁴ Apparently the RIAA at one time provided the number of new releases and then stopped. A general statement about new releases on the RIAA was then taken by Ziemann and used in his analyses, although that was apparently an error. See <http://www.theinquirer.net/?article=9048>

were exogenous. It would be like comparing the total footage of home movies shot year-by-year to explain the yearly economic performance of motion pictures. They are not really part of the same market.

v. The Perfect Storm?

The conclusion from examining these alternative explanations is that they do not appear capable of explaining the decline in sound recordings that has occurred. It would take a remarkable confluence of events, a perfect storm if you will, to explain the large drop that has occurred in the sound recording market. That doesn't mean that it could not have happened. But in a choice between file-sharing as an explanation and the confluence of various disparate factors all perfectly aligned to harm the sound recording industry, Occam's razor requires that we accept the file-sharing hypothesis.

B. Using Countries as the Unit of Analysis

There are several papers that take this approach. The idea is straightforward enough: compare changes in sales of sound recording in different countries over time using some measure, such as the share of Internet users, to proxy for the impact of file-sharing across countries.

Peitz and Waelbroeck (2004) use data from 16 countries for the period 2000-2001. They find a rather large decline for this one year period—11% or 12%, for the world or US respectively. Using survey data from the US they then conclude that the growth in Internet users couldn't have been responsible for more than a quarter of the decline that occurred in 2002.

Alejandro Zentner (2004) uses international cross section data from 97-98 to 2001-2002. He finds a worldwide decline of 15% and a US decline of 30%. Using an analyses of European survey data he finds that file-sharing reduces the probability of buying music by 30% and a drop in music sales of 7.8%.

Difficulties with the approach used by these authors include the fact that many of these countries are very different from one another. They have different levels of organized piracy, speak different languages, and have very different levels of per capita income, CD sales, Internet use, and stereo equipment. This great heterogeneity across countries makes it less likely that file-sharing will have the same impact in each country.

C. Using Cities as the Unit of Analysis

A paper by Eric Boorstin, a student at Princeton, used cities as the unit of analysis. This paper has received some attention because one of his advisors, Ed Felton, is a very influential computer scientist who has declared on his widely read web page that Boorstin's analysis provides the best explanation of the impacts of file-sharing.

The paper is quite clever. Boorstin uses Census data on the number of individuals in each city with and without Internet access in his attempt to determine the impact of file-sharing on the sales of CDs within a city. Internet use is taken as a proxy for file-sharing. In this sense his paper is similar to those performing international comparisons except that his comparisons are based across cities instead of countries. He uses SoundScan data for the sales of CDs in each metropolitan area.

This approach has certain advantages over international comparisons. Socioeconomic variations, as related to music are likely to be smaller across cities in the US than across the countries used in some other analyses. The language is common throughout the country, and the popularity of songs is fairly uniform within the country.

A further advantage is that Census data on Internet usage is broken down by age group so that, in principal, the analyst can determine the impact of file-sharing for different age cohorts of individuals.

The data used by Boorstin cover the years 1998, 2000, and 2001. By ending in 2001, his data contain only two of the four recent years of declining sales, so that we should not expect the analysis of

that data to be able to explain the full drop in sales. Additional controls used by Boorstin are two (unusual) measures of income within a metropolitan area and dummy variables for the years 2000 and 2001.

Boorstin provides regression results (using a fixed-effects model) with CD sales as the dependent variable. These results are partially reproduced in Table 5 (only the results related to file-sharing are reproduced with other variables not reported). Those results have a certain intuitive appeal. The impact of file-sharing is measured as the difference between the coefficients of the cohort of an age group with Internet access and those in the cohort without Internet access. For youthful individuals, file-sharing is deemed harmful to sales as indicated by the negative coefficient in the Internet effect column, whereas for older individuals file-sharing is found to be beneficial. Boorstin attributes the positive impact of older individuals to the sampling hypothesis, where he followed other writers in accepting that sampling would have a beneficial impact on sales. He concludes that the overall impact of file-sharing is not negative since the positive impacts of the older age cohorts outweighs the negative impacts from the younger age cohorts.²⁵

Total Sales	coefficient	t-statistic	internet effect	significance
Age 0-4	-0.3018421	-0.2		
5-14 no Internet	4.831504	3.71		
5-14 with Internet	1.245786	0.76	-3.585718	0.1196
15-24 no Internet	7.134016	5.41		
15-24 with Internet	2.897334	2.13	-4.236682	0.0311
25-44 no Internet	-2.225841	-2.17		
25-44 with Internet	2.921416	2.34	5.147257	0.0016
45+ no Internet	-0.5976357	-0.63		
45+ with Internet	7.614472	6.23	8.2121077	0

Upon closer inspection, however, several problems arise with regards to his specification. One important problem has to do with differentials between Internet users and non-Internet users in a city.

²⁵ Boorstin multiplies these Internet effects by the number of individuals in the cohort. I believe it would have been more appropriate to perform this multiplication using the share of CD purchases made by each cohort, as I do. It would not have changed his conclusion, however.

There are no controls in the regression to account for socioeconomic differences between these Internet users and non-Internet users. Surely, we expect that Internet users are likely to be more affluent and better educated than non-Internet users. Isn't that what the whole debate on the digital divide was about? The data used by Boorstin supports this view since cities with higher incomes also have a greater share of individuals in each age group who have Internet access, as Table 6 makes clear.

Age Group	Share of individuals without Internet access	Share of individuals with Internet access
10-24	-0.3321	0.4063
25-44	-0.2866	0.4875
45-64	-0.3534	0.3797
65+	-0.4483	0.1433

Greater income should lead to greater record sales for the groups with Internet access, although I do not know which way education would likely cut. Because Boorstin's regressions do not account for income differentials by Internet access, we should expect that Internet users will have higher coefficients on the purchase of sound recordings merely due to their higher incomes. This will bias upward (in a positive direction) the impact of file-sharing as measured by Boorstin. This is a potentially serious defect. It is also a possible explanation for the positive impact of file-sharing that Boorstin finds for the older age groups when we have seen the economic theory predicts no such positive impact from 'sampling'.

For this reason, one might argue that the groups with the oldest individuals should be the benchmark for zero impact of file-sharing since it seems very unlikely that those over 65, for example, would engage in file-sharing to any serious extent. If we were to use Boorstin's oldest age group as a benchmark, the other groups would then have various negative impacts from file sharing. I come back to this issue later.

The single most glaring problem, however, is his inclusion of dummy variables for the years 2000 and 2001. Since these years are the years when file-sharing was coming into existence, some of

the impacts of file-sharing are likely to be included in the dummies, as Boorstin freely admits.²⁶ Boorstin suggests that this was a difficult time for the music industry although the industry has claimed that the difficulty was due to file-sharing. But even if the industry were having difficulty for reasons other than file-sharing, it would be inappropriate to include these dummies which will interfere with the file-sharing variables. The pure file-sharing variables are unlikely to be impacted by factors such as a decline in music quality, so the inclusion of these dummies would bias upward the file-sharing variables and serve little extra purpose.²⁷

Making the single change to Boorstin's results of excluding the dummy variables for the years 2000-2001 provides the following results.

Total Sales	coefficient	t-statistic	internet effect	significance
Age 0-4	-4.75724	-2.54		
5-14 no Internet	1.630779	0.99		
5-14 with Internet	-5.3804	-2.73	-7.011182	0.0187
15-24 no Internet	10.41141	6.22		
15-24 with Internet	-5.27322	-3.61	-15.684629	0
25-44 no Internet	1.485543	1.2		
25-44 with Internet	6.496813	4.73	5.01127	0.0065
45+ no Internet	0.629566	0.52		
45+ with Internet	7.082242	4.96	6.4526759	0

This table now looks to tell quite a different story. The coefficients indicate larger negative impacts for young people and smaller positive impacts for older individuals compared to Boorstin's original results. In fact, we can use the information in Table 8 which lists the age distribution of record purchasers²⁸ to estimate the overall impact of file-sharing from the new analysis.²⁹

²⁶ On page 54 of his paper Boorstin states: "The music industry has been complaining about a decline from its 1999 peak, and these dummy variables will pick up the market factors not included in my model. The increasing popularity of substitutes and theorized decline in overall music quality would be captured by these dummy variables. Also captured would be a change in the level of piracy that cannot be captured by Internet usage statistics."

²⁷ Pre-recorded video usage may or may not be related to Internet/income levels, although there is no compelling reason to believe that there is such a link. Even if there were, however, the period 1993-1996 saw a strong upsurge in the sale of prerecorded videos but no decline in the sales of music albums.

²⁸ Boorstin uses the numbers of individuals in each age group.

Table 8: Consumer Information 2001		
age group	share of purchases	Purchase Intensity
10-14	8.5	115.93%
15-19	13	182.68%
20-24	12.2	176.54%
25-29	10.9	164.02%
30-34	10.3	141.84%
35-39	10.2	130.60%
40-44	10.3	128.80%
45+	23.7	64.77%

In 2001, for example, 41.7% of purchases were made by consumers between 25 and 44 years of age and consumers older than 45 purchased 23.7% of the total. Applying the coefficients from the regression to these purchase share numbers implies an overall decline in record sales of 92 units, which is just slightly more than the decline that had occurred as of 2001 (.68 from 1998 and .88 from 1999). Given that this is likely to be an underestimate due to the differential incomes of Internet/non-Internet users, this would imply that the entire decline (plus some forgone increase) in the sales of albums was due to file-sharing.

There is one further change to Boorstin's analysis that I would suggest. In his regressions the number of individuals in various groups in each city was regressed on the number of CDs sold. A more natural and more exacting specification is to use the relative size of these groups (measured as a share of the population) in each city to explain per capita sales in each city. This removes the strong correlation that would be expected merely from the large size differentials between cities. The results are reported below. I also rearranged the groupings into more even natural and ones.

²⁹ This information comes from the RIAA and is based on polls taken by Peter Hart Research. More information is available at the RIAA website www.riaa.com. I have not seen the surveys used to generate this information so I can not form an opinion about their likely veracity. For our purposes, the implications of the results in the table will hold as long as any imperfections in the survey remain consistent over time.

Table 9: Age Groups Based on Share of Population					
Per Capita Units	Coef.	Std. Err.	t	Internet Effect	Significance
10-29 no Internet	2.863291	1.81591	1.58		
10-29 Internet	0.2094592	1.570257	0.13	-2.6538318	0.1024
30-44 no Internet	1.751894	1.8017	0.97		
30-44 Internet	2.658693	2.121496	1.25	0.906799	0.606
45-64 no Internet	3.471715	1.528949	2.27		
45-64 Internet	5.396905	1.691543	3.19	1.92519	0.2006
65+ no Internet	4.859013	1.651844	2.94		
65+ Internet	5.962898	2.43968	2.44	1.103885	0.6637

Under this specification the only Internet effect that is significant is the negative impact for the 10-29 age group. The coefficients are somewhat more reasonable since the Internet effect for the 10-29 age group now implies that a complete shift from no Internet usage to full Internet usage by this group would lower sales of sound recordings by 2.7 units per person, which is roughly a drop of one third for this age group which has average per capita purchases in the range of 9 units.

Of course, the Internet effect is still understated due to the income differential between groups. The 65+ group seems very unlikely to be impacted by file-sharing and thus this may be taken as a benchmark for the impact of higher income on sales of sound recordings. If so, the Internet effect changes to -3.75 for the 10-29 age group, implying that Internet usage would lower sales to this group by approximately 40% if everyone switched or about 24% at the average Internet penetration of about .57 in 2001.³⁰ With this age group responsible for almost one half of all sales, this would lead to a decline of about 12% due to file sharing, compared to the actual decline of about 15% by 2001. These two values are so close that we can say that this evidence is consistent with a view that all of the decline is due to file sharing.

This is a charming story, but it isn't clear how reliable these results are.

³⁰ If we apply this logic to Boorstin's specification we find that when consumers over 65 are put in a separate category the other groups all indicate various negative impacts from file sharing. This would then imply a considerably larger (and implausible) decline in sales than has actually occurred (it would essentially imply that file-sharing would have reduced sales by 8 units per person, which is larger than sales per person before file-sharing occurred).

First of all, the magnitudes of the coefficients in the initial specification are troubling. Actual per capita sales are in the range of 5 or 6 CDs per year, yet according to the initial regression the impacts of Internet access are in the vicinity of 5 to 15 CDs. These large magnitudes are difficult to believe although the specification in Table 9 seems to address some of these concerns, although the high coefficients for older age groups is still troubling.

Second, if this regression is capable of measuring the impact of Internet use, it should also be capable of using population differentials across cities to determine relative purchase rates of different age groups, for which we have statistics from other sources. We have already noticed the disturbing results that older age groups appear to have larger coefficients than younger age groups when broken down by Internet access. Running a regression without Internet groupings provides the results found in Table 10.

group	Predicted Purchase Coef.	Std. Err.	t	Actual Purchase Intensity
Age 10-24	0.8066711	1.510765	0.53	157.78%
Age 25-44	2.022539	1.513401	1.34	140.34%
Age 45+	4.179985	1.192174	3.51	64.77%

Based on these regressions, individuals over 45 would appear to purchase records with a greater frequency than younger individuals. We know, however, that individuals in the 45+ age category actually purchase records with the slightly less than half the intensity of younger individuals according to purchase statistics reported by the RIAA. Unfortunately, the inability of these regression results to generate results consistent with known behavior reduces any confidence that we can have in the reliability of this approach.

It is not clear how much confidence we can place in these results. But they certainly do not show file-sharing to increase sales, or even a neutral result once the year dummies are removed. The stronger

negative impact for younger individuals appears quite robust, however, providing some possibility that this technique provides some useful information.

D. Using Records as the Unit of Analysis

A recent paper by Oberholzer and Strumpf (O&S) uses records as the unit of analysis. Unlike almost everyone else, Oberholzer and Strumpf were allowed access to actual downloads logs on a server that was part of a file sharing system. They then matched music downloads to the recordings where the songs appeared. Using data from SoundScan on recordings they then compared the downloads of songs to the sales of songs over a seventeen week period. The creation of this data set was a substantial effort requiring much work and creativity. Whether it can be used to determine the impacts of file-sharing, however, is another question,.

There are two potential problems. First, can an analysis using records as the units of observation provide information about the impacts on the entire industry, as opposed to the impacts on individual recordings? Assuming that this first problem is not an issue, a second obstacle exists. There is a serious simultaneity problem because popular songs are going to be both heavily purchased and heavily downloaded. This simultaneity problem would need to be overcome before the results would be accurate.

First, however, I discuss their results.

- i. What they conclude versus what I think they should conclude

O&S suggest that their results indicate that file-sharing has no impact on the sales of sound recordings. I read their paper differently, and believe their results actually indicate that file-sharing has a positive impact on sound recordings. Ignored in this discussion is the aforementioned possibility that that impact of file-sharing on individual sound recordings might be different than its impact on the sound recording industry.

O&S begin their regression analysis by looking at their universe of approximately 650 recordings and the relationship of sales to downloads (their table 11). As a benchmark they run a regression with sales as the dependent variable and downloads as an independent variable (using musical genres as additional independent variables). The coefficient on downloads is 1.07 and statistically significant. Of course, this coefficient is infected by the simultaneity problem mentioned above, implying that the coefficient is biased upward.

In an attempt to solve the simultaneity problem, O&S use several instrumental variables related to the length of songs on the albums, which they argue are appropriate instruments. In the second stage of their two stage regression the coefficient on fitted downloads is significantly positive, at 1.467. Strangely, the instruments that were supposed to control for the upward bias in the coefficient on downloads actually increased the coefficient. It would appear that the instruments are not controlling for simultaneity.

Given parameters used elsewhere in their paper, this coefficient translates into additional sales of 198 million albums that were caused by downloading. This is rather difficult to accept as being even remotely plausible. O&S tend to skip quickly by this result, however.

Their Table 12 adds weekly information into the mix, allowing O&S to use a fixed effects model. Since the fixed effects model includes a dummy for each sound recording, much of the popularity simultaneity is taken care of merely by the use of the fixed effects model and the coefficient drops from 1.193 to .28. However, there is still simultaneity going on week to week since records are more and less popular over time. Again, O&S introduce several instruments to solve the simultaneity problem. The second stage coefficient on downloads, in their two preferred models are close to zero, taking on values of -.014 and +.088. The -.014 implies that it would take 5000 downloads to decrease sales by 1 album, a number they have repeated several times when describing their results. The

“preferred” measurement (.088) would translate into a gain due to downloading of 12 million albums, which is quite small.

There is an important conceptual error in using these regression to calculate the overall impacts, however. These regressions give each record album equal weight in the analysis. This would not be appropriate to answer a question about the overall impact of downloads on sales unless the impact of downloads on sales was the same for high selling records as it was for low selling records. O&S accidentally test this assumption in their Table 13 where they find that the impact of downloads is higher for albums in the top quartile of sales than it is for less popular albums. Since the top quartile of albums represent the large majority of total sales (85%), giving equal weight to all albums causes O&S to come to erroneous conclusions regarding the overall impact of downloads.

In their table 13 the quartile of largest sellers has a coefficient of .468 ($t=1.52$). This coefficient implies an increase of 63 million albums due to downloading. Although the level of statistical significance is slightly below traditional norms, this is not a ‘zero’ impact.

Again, as was the case with their table 11, their instruments appear to be incapable of doing the job as I discuss these in more detail below. This review of their results, however, indicates to me that their results actually imply that file-sharing increases the sales of sound recordings.

ii. The Unit of Analysis: Potential Fallacy of Composition

To make this concern concrete, imagine that we wished to learn what the elasticity was for the wheat industry. What would happen if we estimated the elasticity of demand for each of the firms within the industry and then assumed that the elasticity for the industry was the average of the elasticities of its constituent parts? We all know, and teach our students, that industry elasticity is much lower than the elasticity facing the firms because there are fewer choices for consumers at an industry level than there are at the firm level. This would be an example of the fallacy of composition that sometimes occurs when going from constituent parts to the whole.

A similar compositional problem would occur with respect to advertising at the firm level versus advertising at the industry level. The advertising analogy is particularly useful apropos record sales due to the fact that downloading, sales, and popularity are all related in some sense to the ‘buzz’ that surrounds a recording, part of which is directly due to advertising.

Assume that downloading increases the sales of a record. Perhaps downloading creates a buzz for record x and people buying records buy those records with buzz. Or perhaps downloading allows consumers to sample a high quality recording such as x, leading to an increase its sales. In this example, heavy downloading makes recording x more successful than lightly downloaded recording y.

A regression using downloads to explain sales would return a positive coefficient, assuming that all other simultaneity problems were overcome. After all, increases in downloads, by assumption, lead to an increase in the sales of the downloaded recordings in this example. But that does not mean that downloads increase overall record sales. A positive coefficient could be entirely consistent with record sales being severely harmed by downloads and thus couldn’t answer the question about the overall impact of downloading.

The sampling story, discussed above, is a case in point. Sampling might lead to greater sales for those recordings providing higher quality, but it is likely to lead to an overall decrease in record sales as discussed above. Any informational aspect of downloading separate from sampling, if there were one, has a similar result. Political advertising clearly increases the market share of individual candidates (increases their ‘buzz’), but it need not, and probably has not historically increased total voter turnout. Yet a regression on individual advertising campaigns would show that when candidates increase their spending they generate a higher share of votes (or higher poll numbers).

My paper on the impact of radio broadcast on record sales shows a similar result. Although radio play obviously has a positive impact on market share of individual records, and would expected to in a

regression to have positive coefficient on the sales of individual records, its impact on overall record sales appears to be zero or negative.

Still, the possibility of compositional difficulties is not proven. Nevertheless, the methodology of O&S is suspect and should be treated with caution.

iii. Problems with Instruments

In Table 11 each instrument takes on 670 different values, one for each album, although we have seen that the instruments appeared unable to solve the simultaneity problem. In Table 12, each instrument only has 17 observations, one for each week.

One of the instruments is the number of German youth not attending school. The claim is that Germany provides almost 15% of downloaded files to Americans and that additional files available when German youth are on vacation from school will influence downloading in the US. This seems like quite a stretch.

First, we must assume that German school kids leave their computers on when they are at home but turn them off when at school. Also, we must disregard the possibility that German university students make fewer files available when they come home for vacation since their computers presumably stay at school turned off. Since the time sample used by O&S includes the Christmas holidays, going home for the holidays is a real possibility.

Second, we must assume that having these files available during the school day in Germany has a non-negligible impact on files available to American downloaders. Since Germany is 7 or so hours ahead of most American time zones, most Americans are sleeping during the period when these additional files are made available.

Third, we must assume that a decrease in files from German kids shifts the supply curve for downloaders, making downloads more difficult. But the supply of MP3s might be sufficiently large

that there is no scarcity for most files. In that case the impact of files available from German school children may not impact downloads any more than everyone in Germany taking a deep breath would impact the air available for breathing in the US. O&S do get a positive coefficient on the German holiday variable, but that may just be due to the fact that some holidays are common between the countries, and Christmas is a big holiday in both countries.³¹

Finally, it is crucially important that German holidays are not linked to the error term on record sales. Yet the fact that the major holiday of Christmas is common between the US and Germany might indicate a relationship between German school kids on vacation and (the error term for) US record sales. If this were the case then usage of this instrument might suffer from the problem of ‘weak instruments’ which is known to cause very serious estimation problems.³²

Although the congestion variables seem less troublesome there are questions here as well. The congestion variable is measured on a weekly basis. Is Internet usage like phone usage – heaviest during the day when businesses are in operation? If American music downloading takes place at home during the evening (off-peak usage for American and European business) wouldn’t overall congestion largely be irrelevant?

There is also a potential simultaneity problem here as well. If some new songs are very popular, sales will increase and so will downloads. With downloads taking up nearly half of total Internet bandwidth, congestion might be related to (the error term of) sales and the ‘weak instrument’ problem becomes a possible problem.

³¹ I am surprised that the coefficient on vacation kids was even positive. According to a web page listing German school holidays there are usually 12 days in October plus the typical Christmas holiday. Yet according to O&S’s Table 3, October is when downloads were lowest.

³² Reference to Woolridge.

E. Using Genres as the Unit of Analysis

I had high hopes for this methodology, although no one has attempted to use it to my knowledge. With data such as that created by Strumpf and Oberholzer, one could determine the variation in file-sharing proclivity by genre, measured as the number of files downloaded as a share of total sales. It is likely that certain groups of users (rock, hip hop?) would be far more likely to engage in file-sharing than other groups (classical).

If yearly sales by genres was reasonably smooth, then we could look to see if those genres with the greatest incidence of file-sharing suffered relative to other genres. As a control, we could use the share of radio listenership by genre.

I had high hopes for this methodology until I purchased data on genre sales from SoundScan. First, they had only seven genres plus a catchall genre called “current.” Second, and more important, sales by genre had extremely large changes over very short periods of time. The chart below, showing two genres, has a 170% increase in sales in the Metal genre in 1999 and a 100% increase in the R&B genre in 1997. Without more information that might help explain such sudden jumps, this data seem unlikely to prove reliable enough to use for any important empirical tests.

	Metal	R&B
1994	38,739	80,819
1995	31,101	80,718
1996	26,409	74,035
1997	28,983	141,613
1998	30,086	166,379
1999	82,698	175,339
2000	89,924	197,141
2001	88,158	195,498
2002	74,677	160,183
2003	74,629	149,972

V. Conclusions

The analysis of file-sharing is a relatively new phenomenon. Economists are just getting started. Nevertheless, progress is being made.

The theory underlying the analysis of file-sharing has not received the attention that it deserves. It has always been clear that some possible aspects of file-sharing would harm copyright owners, such as the substitution of copies for the purchase of originals. What has not been understood is that the use of file-sharing to sample products is also likely to lead to harm to copyright owners. Although one can still construct theoretical conditions under which file-sharing might benefit copyright holders, these conditions seem quite far fetched. A broad analysis of the various theoretical factors at work supports a view that file-sharing is likely to cause serious damage to the owners of copyright materials that are so shared.

The last few years have provided economists a laboratory of sorts to analyze the impacts of file sharing. The basic evidence in the United States over the last few years—the birth of file-sharing and the subsequent decline in CD sales—makes for an extremely compelling and simple explanation in spite of the protestations to the contrary from a large and vocal group of individuals supportive of file-sharing.

Empirical examinations by economists must be undertaken against the backdrop of this simple explanation. That these basic facts fit nicely with the economic theory only adds support to this explanation. With such a strong prior, empirical studies need to meet a higher hurdle than normal before they might be considered to overturn this expectation.

All of the empirical works to date suffer from various imperfections. Nevertheless, the majority of studies find results supportive of the thesis that file-sharing is causing harm. As more evidence accumulates, both in the world at large and in the pages of academic publications we can expect to learn more about the impacts of file-sharing. The political arena is calling for an answer now, however, and is impatient to wait for academics to reach unanimity, which academics almost never do anyway. If an answer is needed, the answer that would appear to have the greatest likelihood of being correct given our current state of knowledge, is that file-sharing hurts copyright owners.

- Boorstin, Eric “Music Sales in the Age of File Sharing” Senior Thesis, Princeton University, April 2004.
- Conner K. R. and R. P. Rumelt, ‘Software Piracy-An Analysis of Protection Strategies,’ *Management Science* 37 (2): Feb. 1991, pp. 125-139.;
- Gopal, Ram D., Sudip Bhattacharjee, and G. Lawrence Sanders “Do Artists Benefit From Online Music Sharing?” *Journal of Business*, forthcoming.
- Hirshleifer, Jack "Suppression of Inventions," *Journal of Political Economy*, March/April 1971, 79, 382—83.
- Lessig, Lawrence *Free Culture*, The Penguin Press, New York 2004.
- Liebowitz, Stan J. “Will MP3 downloads Annihilate the Record Industry? The Evidence so Far” *Advances in the Study of Entrepreneurship, Innovation, and Economic Growth*, V. 15, 2004, pp. 229-260.
- Liebowitz, Stan J. “The Elusive Symbiosis: The Impact of Radio on the Record Industry” *Review of Economic Research on Copyright Issues* Vol. 1, 2004a, pp.20-45.
- Liebowitz, Stan J. *Rethinking the Network Economy*, Amacom, New York, 2002.
- Liebowitz, Stan J. “Copying and Indirect Appropriability: Photocopying of Journals,” *Journal of Political Economy*, 93-5 October 1985, Pp. 945-957.
- Liebowitz, Stan J. “Durability, Market Structure And New-used Goods Models,” *American Economic Review*, September 1982a, 72-4, Pp. 816-824.
- Liebowitz, Stan J. “The Impacts of Cable Retransmission on Television Broadcasters,” *Canadian Journal of Economics*, 15-3, August 1982, Pp. 503-524.
- Oberholzer, Felix & Koleman Strumpf “The Effect of File Sharing on Record Sales An Empirical Analysis” working paper, 2004.
- Peitzl, Martin and Patrick Waelbroeck “An Economist’s Guide to Digital Music” Working Paper, May 2004.
- Peitzl, Martin and Patrick Waelbroeck “The Effect of Internet Piracy on CD Sales: Cross-Section Evidence” Cesifo Working Paper No. 1122 Category 9: Industrial Organisation January 2004
- Romer, Paul “When Should We Use Intellectual Property Rights?” *American Economic Review*, May, 2002, 213-216.
- Shy Oz and Jacques-Francois Thisse, ‘A Strategic Approach to Software Protection,’ *Journal of Economics and Management Strategy*, 8, 1999, pp. 163-190.;
- Takeyama, Lisa N. ‘The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities,’ *Journal of Industrial Economics*, 42, 1994, pp. 155-166.

Varian, Hal “Copying and Copyright” *Journal of Economic Perspectives*, forthcoming.

Watt, Richard “The Past and the Future of the Economics of Copyright” *Review of Economic Research on Copyright Issues*, 2004, vol. 1 (1), pp. 1-11.

Weimann, Gabriel “Cable Comes to the Holy Land: The Impact of Cable TV on Israeli Viewers,” *Journal of Broadcasting & Electronic Media*, 40, 1996, pp. 243-257.

Zentner, Alejandro “Measuring the Effect of Online Music Piracy on Music Sales” 2003.

Zentner, Alejandro “Measuring the Effect of Music Downloads on Music Purchases” April 2004