

NBER WORKING PAPER SERIES

COPYRIGHT PROTECTION, TECHNOLOGICAL CHANGE, AND THE QUALITY OF NEW PRODUCTS:
EVIDENCE FROM RECORDED MUSIC SINCE NAPSTER

Joel Waldfogel

Working Paper 17503
<http://www.nber.org/papers/w17503>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2011

The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2011 by Joel Waldfogel. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Copyright Protection, Technological Change, and the Quality of New Products: Evidence
from Recorded Music since Napster

Joel Waldfogel

NBER Working Paper No. 17503

October 2011

JEL No. K11,L82

ABSTRACT

Recent technological changes may have altered the balance between technology and copyright law for digital products. While file-sharing has reduced revenue, other technological changes have reduced the costs of bringing creative works to market. As a result, we don't know whether the effective copyright protection currently available provides adequate incentives to bring forth a steady stream of valuable new products. This paper assesses the quality of new recorded music since Napster, using three independent approaches. The first is an index of the quantity of high-quality music based on critics' retrospective lists. The second and third approaches rely directly on music sales and airplay data, respectively, using of the idea that if one vintage's music is better than another's, its superior quality should generate higher sales or greater airplay through time, after accounting for depreciation. The three resulting indices of vintage quality for the past half-century are both consistent with each other and with other historical accounts of recorded music quality. There is no evidence of a reduction in the quality of music released since Napster, and the two usage-based indices suggest an increase since 1999. Hence, researchers and policymakers thinking about the strength of copyright protection should supplement their attention to producer surplus with concern for consumer surplus as well.

Joel Waldfogel

Frederick R. Kappel Chair in Applied Economics

3-177 Carlson School of Management

University of Minnesota

321 19th Avenue South

Minneapolis, MN 55455

and NBER

jwaldfog@umn.edu

Creative products, such as movies, music, and books, have high fixed costs and low marginal costs.¹ Private firms have traditionally been able to profitably bring them to market because these products are excludable, through a combination of technology and the complementary legal framework provided by copyright law. Physical media products are sufficiently difficult to copy that purchasing them has been the easiest means to their acquisition. Moreover, copyright grants legal monopoly rights to creators, assisting them in appropriating returns from their works. While this arrangement gives rise to monopoly's usual harm to consumers, this harm is thought to be offset by copyright's incentive effects on the creation of new works.²

Recent technological changes may have altered the balance between technology and copyright law. First, file sharing reduces the revenue available for any particular digital product, with recorded music as a leading example. On its own, this would tend to reduce the flow of new products, particularly if creators are motivated by economic factors. Organizations representing the recorded music industry have voiced concern that weakened effective copyright protection will undermine the flow of new recorded music products. The International Federation of the Phonographic Industry (IFPI) describes music as “an investment-intensive business” and worries that “piracy makes it more difficult for the whole industry to sustain that regular investment in breaking talent.”³ The Recording Industry Association of America's (RIAA) explains that its anti-piracy efforts seek “to protect the ability of the recording industry to invest in new bands and new music...” And: “this theft has hurt the music community, with

¹ See Caves (2000) for extensive discussion of the nature of media products.

² See Boldrin and Levine (2008) for a discussion of the net benefits of intellectual property rules.

³ See http://www.ifpi.org/content/section_news/investing_in_music.html , accessed October 20, 2010.

thousands of layoffs, songwriters out of work and new artists having a harder time getting signed and breaking into the business.’’⁴

At the same time that file-sharing has weakened effective copyright protection, other technological changes have reduced many of the costs of bringing digital creative works to market. Production, promotion, and distribution of music have all been made less expensive by new computing and information technologies. As a result, the revenue needed to cover costs to maintain the traditional flow of products may have declined. It is possible that despite being weakened by Napster, the effective copyright protection still available may be sufficiently strong to facilitate a continued flow of valuable new recorded music products. Making this determination requires understanding of whether consumers continue to face a steady stream of valuable new products in the face of the compound experiment of weakened copyright protection in conjunction with new technologies for bringing products to market in the post-Napster era. This paper seeks to address that question by, first, creating indices of the quality of recorded music over time and, second, by asking how these indices have fared since Napster. While reductions in revenue are comparatively easy to document, quantitative assessment of the volume of consequential new music products is more challenging. It is natural to point, for example, to the number of products released each year, but the distribution of consumption is skewed, and most products are rarely if ever purchased.⁵ Thus, most products contribute little to consumer and producer surplus; and the number of products, while interesting, is not particularly informative about the welfare generated by products. A second impulse is to quantify the number of products whose sales pass some threshold (e.g. 5000 copies). But in an era of

⁴ See <http://www.riaa.com/physicalpiracy.php>, accessed October 20, 2010.

⁵ See Handke (2006, 2009) and Oberholzer-Gee and Strumpf (2009) for discussions of the increased volume of music released in recent years.

increasing theft, 5000 copies is an increasingly difficult target. A work of equal quality appearing in, say, 1998 and 2008 would sell fewer copies in 2008, so this method will not work.

Against the backdrop of this challenge, this paper presents three independent approaches to quantifying the evolution of music quality over time. First, I develop an index of the quantity of high-quality music based on critics' retrospective lists of the best works of multi-year time periods. In particular, I assemble data on album quality from 88 Anglophone sources, chiefly from retrospective lists (e.g. Rolling Stone's 500 best albums, Pitchfork Media's 200 best albums of the 1990s, etc.). Each of these rankings allows us to create an index of the number of albums released each year meeting the criterion. I combine the indices statistically to create an overall index of the volume of high-quality music since 1960.

My second and third approaches to quantifying the evolution of music quality are more tightly linked to the service flow of recorded music by vintage, making use of the following insight: If one vintage's music is better than another's, its superior quality should generate higher sales or greater airplay through time, after accounting for the time elapsed since release. Using data on the airplay and sales of recorded music by calendar time and vintage, I am able to construct two separate indices of the mean utility or "quality" of music from different vintages. The approach evaluates vintages by the extent to which whether they continue to be played – or continue to sell – at above-normal rates after accounting for their age. I create these usage-based indices of vintage quality for the period since 1960. I can then ask whether these indices track the critical index, as well as whether they track each other. Moreover, I can ask how all three of the indices evolve, absolutely or relative to pre-existing trends, since the major technological changes following Napster.

The paper proceeds as follows. Section II lays out a simple theoretical framework illustrating the importance of the long-run supply question. Section III describes the critics' data and the resulting index. Section IV describes our sales and airplay data in detail, along with our empirical approach for extracting vintage quality from data on sales or airplay by time and vintage. Section V presents statistical results on the changes in these indices since Napster. Our indices are consistent with each other, and with other historical accounts of recorded music quality, and we find no evidence of a reduction in the quality of music released since Napster. Indeed, the two usage-based indices suggest that the quality of music has increased fairly substantially since 1999. Section VI presents a discussion, and a brief conclusion follows.

II. Theory

Like any product, music generates surplus for two parties, buyers and sellers. While recorded music is durable in some senses – the recordings can last forever and can be reproduced digitally without degradation in quality – it is subject to taste depreciation. Obviously, there are exceptions. Many people still listen to classical music that is hundreds of years old. But for the most part, consumers prefer new music, as we will see in the data below: While roughly one seventh of music on the radio in a particular year was released in the same year, less than 10 percent was originally released 5 years earlier, and less than 2 percent was originally released 10 years ago.

The fact that music depreciates is important for a welfare analysis of supply disruptions. If it did not, then the consumer losses from a slowdown in new product introductions would be

of only second-order importance. If the amount of music available increased a few percent in a normal year, then a complete cessation of new production would still leave consumers with nearly as much variety as they would have faced if new products had continued to arrive. But because most music does seem to depreciate for most users, disruptions to supply are potentially important for the welfare that this product delivers.

The welfare analysis of sharing zero-marginal-cost digital products has both static and dynamic components. The static analysis describes music that already exists. Putting aside all of the usual problems with theft – such as costs incurred preventing theft from occurring – it is easy to see that sharing files for music that already exists increases welfare on balance. Producers lose, but their losses – when consumers steal things they used to pay for – are all transfers to consumers, who now enjoy greater surplus (the price they had formerly paid plus the former consumer surplus). In addition to the transfers from producers to consumers, file sharing also turns deadweight loss – circumstances in which consumers valued music above zero but below its price and therefore did not consume – into consumer surplus. In a purely static analysis – again, ignoring problems associated with theft – eliminating intellectual property rights benefits consumers more than it costs producers and is therefore beneficial for society.

Of course, the static analysis above is valid only for works that already exist. The dynamic analysis is different. If developing products requires investments of time or money, then producers may only make these investments in the hopes of obtaining returns. If the returns are eliminated, then producers may stop investing, as the above statements of the industry associations suggest. If music fully depreciates in one period, then no valuable products are available in the second period; and there is no surplus for either party. In contrast to the welfare-

improving static effects of file sharing on welfare, the dynamic impact is potentially devastating. This focuses attention on the paper’s goal, a quantification of the flow of high-quality new recorded music in the past decade.

III. A Critic-Based Quality Index

The basic data for constructing the critic-based index are professional critics’ retrospective rankings of songs and albums from multiple years, such as “best-of-the-decade” lists.⁶ For these lists, the staff of a magazine or website produce a list of the best albums (or songs) of the past decade, or quarter century, or all time. That is, experts evaluate music from different years, subjecting all of it to a time-constant quality threshold for list inclusion. I have been able to assemble data from 88 different rankings (and ratings), 64 covering albums and the remainder covering songs. All of the rankings are from Anglophone countries (the US, England, Canada, and Ireland).⁷

These rankings generate data of the form: $\mu_1 > \dots > \mu_N$, where μ_i is the quality of work i . If T_k is a quality threshold such that $\mu_k > T_k > \mu_{k+1}$, then each of these rankings allows us to calculate the number of works above a constant T_k released in each year. These rankings allow ready creation of indices showing the volume of works released in each year that pass some threshold. Figure 1 displays the sources and their respective chronological coverage periods.

⁶ This material is described more extensively in Waldfogel (2011).

⁷ We discovered rankings in a variety of places. The Acclaimed Music website lists many of these, including the majority of the lists we use for the period since 1999. See, in particular, the lists of the top albums and songs of the 2000s at <http://www.acclaimedmusic.net/>, accessed December 21, 2010.

Prominent examples include Rolling Stone's 2004 list of the 500 best albums or Pitchfork Media's list of best 200 albums of the 2000s. Entries on the Rolling Stone list "were chosen by 273 of the world's pre-eminent musicians and critics ranging from Fats Domino to Moby" (Levy 2005). Figure 2 depicts the index derived from Rolling Stone's list, and a few things are immediately evident. First, perhaps because Rolling Stone was founded in 1967, its editors are very fond of 1960s music. Second, the index trails off toward the year that the list appeared (2004).

Indeed, the process of producing long-term retrospective lists appears biased against recent works. For example, Pitchfork Media produced a list of the top 100 albums of the 1990s in October 1999, then another list covering the same period in November 2003. The latter list was introduced with a statement contrasting it with their 1999 ranking, "...looking back at that list a lot has changed: our perceptions of the decade are different now, our personal tastes have expanded, our knowledge of the music has deepened..."⁸ And, indeed, the later ranking includes a greater emphasis on the last years of the decade. Ten percent of the albums on the 2003 list were released in the last two years of the decade, compared with only seven percent for the 1999 list. Hence, we can use the retrospective rankings but exclude the year the ranking appeared as well as the previous year to avoid a bias against recent works. Together, the 64 album lists cover the period 1960-2007 and include 15,158 entries. The 24 song lists also cover 1960-2007 and include 1806 entries. While critic-based data are unconventional, we can provide a few pieces of evidence of their legitimacy. First, we find that they track well-known historical trends in music. For example, historians of contemporary popular music believe that the late 1960s was a period

⁸ See <http://pitchfork.com/features/staff-lists/5923-top-100-albums-of-the-1990s/>, accessed October 18, 2010.

of unparalleled creative output in recorded music.⁹ And the indices – such as Rolling Stone, reported above – reflect that. Second, the various indices are highly correlated with each other. Of the five indices that extend back to the 1960s, all but one of their pairwise correlations exceed 0.7.

Because the period following 1999 is crucial to this study, it is important to provide evidence of the reasonableness of the rankings and resulting indices for the post-1999 period. We have 56 professional critics' album lists – and 22 professionals' songs lists – covering this period (beginning in 2000). To determine whether these lists contain a common signal rather than simply noise, we examine overlap across lists. We see a great deal of concordance across these lists: Two albums – *Funeral* by Arcade Fire and *Kid A* by Radiohead appear on 47 of the 56 lists covering the 2000s. *Is this It?* by the Strokes and *Stankonia* by Outkast appear on 45 and 37 lists, respectively. One hundred albums account for 40 percent of the entries on decade-best lists, 250 albums account for over 60 percent, and 500 albums account for over three quarters of the 4202 entries on 56 publications' best-of-the-2000s lists. At least 300,000 albums were released during the decade. Yet, 500 albums – less than 0.2% of the decade's new releases – account for three quarters of the entries on 56 critical best-of-the-2000s lists.

The relationship between critical acclaim and sales provides another source of validation for the critical data. If the designation of being an acclaimed album is relevant to whether the album's existence and consumption generated extra satisfaction for consumers, then critically acclaimed albums should sell more. And, indeed, critical acclaim and sales are linked. Of the

⁹ For example, Larkin (2007) writes, "The 60s will remain, probably forever, the single most important decade for popular music."

50 most acclaimed albums of the 2000-2009 decade, half sold at least half a million copies in the US. This is highly atypical: less than one percent of albums sell more than half a million copies.

If we define y_{it} as the number of works on list i that were originally released in year t , then we can describe the time pattern of new works supply with a regression of the log indices on index dummies and flexible time dummies: $\ln(y_{it}) = \mu_i + \theta_t + \epsilon_{it}$, where μ_i is an index fixed effect, θ_t is a time effect common across indices for year t , ϵ_{it} is an idiosyncratic error. Figure 3 shows time series plots of the annual values of θ_t , along with a vertical line in 1999. Because the regression dependent variable is in logs, the index is in percent terms.

The index rises from 1960 to 1970, then declines to about 1980. The index then rises in the mid-1990s and declines to 1999. Following 1999, the index is stable. Although more formal statistical characterizations follow at section V, this is our first glimpse of results; and a few things are evident. First, while the index had been declining prior to Napster's appearance in 1999, the decline did not continue past 1999. Second, this approach gives no indication of a reduction in the quantity of high-quality music following Napster.

IV. Usage-Based Approaches

While interesting, the critic-based index has some weaknesses. First, despite its apparent relationship with sales, the critical data are not themselves reflective of consumer behavior. Second, critics' best-of lists may include only a handful of albums from each year whose subsequent critical acclaim does not faithfully reflect the service flow of music from that year generally. Concerns of this sort lead us to an alternative indices based more directly on the

service flow from the music of each vintage. These usage-based data for this study, on sales and airplay by calendar time and vintage of music, are drawn from two sources.

1. Airplay Data

For the five years 2004-2008, I observe the share of songs aired on radio that were originally released in each prior year. The airplay data are from a major firm monitoring airplay. The firm monitors the songs played on 2000 radio stations and maintains data on, among other things, the original release date of each song.¹⁰ Each year's data are based on observing over a million "spins," so the vintage shares, even for vintages as early as 1960, are estimated with precision.¹¹

The distribution of a year's airplay across vintages clearly demonstrates depreciation: recent songs make up the largest share of what's played, and older songs are played steadily less. As Figure 4a indicates, about 13 percent of songs on the air in 2008 were released in 2008. About 16 percent of the songs aired in 2008 were released in 2007, and going farther back, the share then declines almost monotonically in vintage: 12 percent for 2006, 9 percent for 2005, 7 for 2004, and so on. The decay pattern includes some curious deviations from smooth decline – for example, 1995 appears to be above the pattern defined by the other vintages – and this is suggestive about vintage effects. I observe an analogous vintage distribution based on airplay in 2007, 2006, 2005, and 2004. Figure 4b highlights the share of music aired in 2008 originally released in the years 1960-1990. Even for early dates, the decay pattern is smooth.

¹⁰ I am grateful to Rachel Soloveichik of the Bureau of Economic Analysis for sharing the airplay data she employed in Soloveichik (2011).

¹¹ For example, of the songs aired in 2005, suppose 1 percent were originally released in some year. Given that the proportion is calculated with over a million spins, the 95 percent confidence interval surrounding that year's proportion would be no larger than 0.04 percent.

2. Sales Data

Ideally, I would observe sales of recorded music by time and vintage of recorded music products. That is, I would like to observe the sales of 1975 music in 1990, and so on. Moreover, I would like to observe actual sales, so that I could accurately characterize the entire sales distribution by vintage. My sales data, from the RIAA's Gold and Platinum Certification database (<http://riaa.com/index.php>), approximate this ideal. The RIAA announces when each single or album's sales pass 0.5 million ("gold"), 1 million ("platinum"), as well as multiples of one million.¹² The timing of these successive certifications allows me to create a rough measure of each album's sales over time. The measure is crude in that I only observe whether its sales pass each of these thresholds and, if so, when. Still, because I am not interested in particular albums but rather in the total sales of music from each vintage in each calendar year, some of the measurement error may average out.

I obtained all of the certifications awarded between 1958 and 2010. This is a total of 17,935 album certifications, 4,428 single certifications, and 2,341 certifications for other products. Each certification includes the work's original release date and certification date (month, date, and year), the artist, the album title, the label, the type of certification (gold, etc), and whether the artist is a soloist, part of a duo, or part of a group. Prior to 1987, many certifications are missing release dates. Excluding those observations leaves me with 15,866 album and 3,556 single certifications with complete data. If an observation is a Gold certification, I code it as 0.5 million in sales in the year of certification. I code a Platinum

¹² Prior to 1989, a single was certified "gold" only when its sales reached 1 million. Since then, singles have received gold certifications with 0.5 million sales.

certification as an *additional* 0.5 million sales in the year of certification. Finally, I code a multi-platinum certification as an additional 1 million in sales in the year of certification.

While the certification data cover only a small fraction of albums released, they cover a relatively large fraction of music sales. That is, sales are heavily concentrated in a small number of high-selling albums. For example, the RIAA reported CD shipments of 292.9 million units in 2009.¹³ Sales calculated from certifications awarded in 2009 total 155.5 million, which is roughly half of the total reported physical album shipments. While sales data derived from certifications are imperfect, they appear nevertheless to cover a large share of total sales. Moreover the sales implied by certifications reflect time patterns known to hold for music sales in the aggregate. For example, certification-based sales rise to a peak around 2000, then decline.

Although the certification database reports release dates and certification dates by day, the data are sufficiently sparse that I aggregate by year. The resulting database is organized certification year by release year. That is, for each year I can calculate the total certification-based sales of albums released in any year since, say, 1960.

Figure 5 shows the album sales distribution by vintage, averaged over all years in the data, and a few patterns are evident. First, sales tend to be concentrated around the time of release. Second, there is relatively steady decay in sales over time. Roughly 45 percent of certification-implied sales occur in the same year as release. Another quarter occur in the year following release, while about 8, 7, and 5 percent occur two, three, and four years after release. This figure shows smooth decline, in part because of averaging. If we examine the analogous album sales distribution for a particular certification year, the certifications are sufficiently

¹³ See <http://76.74.24.142/548C3F4C-6B6D-F702-384C-D25E2AB93610.pdf>, accessed July 22, 2011.

sparse that the data are somewhat bumpy and include many zeroes. Figure 6a shows the distribution of release years for certifications in 2000, and Figure 6b highlights the shares of year 2000 sales for albums originally released prior to 1990. Due to relatively sparse samples arising from the lumpiness of certification-based sales data, the figure does not decline nearly smoothly as the airplay distributions.

The certification data begin in 1958, but the data are quite sparse prior to 1970. In what follows, we focus on the period since 1960 with the airplay data and the period since 1970 for the certification data.

3. Empirical Approach for Usage Data

Our goal is to derive an index of the importance of the music from each vintage. To this end define $s_{t,v}$ as the share of vintage v music in the sales or airplay of music in period t . Suppose that we observe this for V vintages and T years. For a given year t , s varies across vintages for two reasons. First, music sells less, and is played less, as it is older, an effect akin to depreciation. Second – and this is the effect of interest – vintages are used differently because they differ in quality. Our goal is to control for depreciation and to ascertain an index reflecting the quality of each vintage.

A simple way to measure the evolution of vintage quality, in a way that controls for depreciation, is to compare different vintages' market shares in years that occur equally long after the respective vintages' original appearances. To this end, define $s(k,v)=s_{t,v|t-v=k}$. The term $s(k,v)$ is the share of vintage v music among airplay or sales k years later ($t=v+k$). Because the

airplay data cover spins from 2004 to 2008, $s(0,v)$ – current-year music’s share among this year’s airplay – can be calculated for $v=2004,\dots,2008$. More generally, the share for k' -year-old music – $s(k',v)$ – can be calculated for $v=2004-k',\dots,2008-k'$. By contrast, the certification data cover many more calendar years, effectively from 1970 to 2010.

The airplay data – with $t=2004,\dots,2008$, and $v=1960,\dots,t$ – support the calculation of a family of vintage quality indices. For example, $s(44,v)$ (the market share for 44-year-old music) can be calculated for $v=1960-1964$, $s(20,v)$ for $v=1984-1988$, and so on. There is vintage overlap across adjacent indices: $s(0,v)$ – the market share of music released this year – is available 2004-2008; and $s(1,v)$ – the market share of music released last year – is available 2003-2007. Thus, for 2004-2007, both are available, and both series should measure the evolution of the quality of the overlapping vintages. Indeed, their movements should track one another. Their levels, on the other hand, should not. Because of depreciation, a given vintage’s share will generally be higher when the vintage is recent. That is, generally, $s(k,v) > s(k+q,v)$, for $q>0$.

Figure 7 displays all of the 45 adjacent $s(k,v)$ series, in 9 separate panels for $k=0,44$. The figure clearly shows two things. First, vintages’ qualities manifest themselves in correlated series. In any given vintage year, when one series is rising, the others tend – overwhelmingly – to be rising as well. Second, for any particular vintage v , the series levels tend to fall, the longer is the retrospective period k .

Because t runs back only to 2004, we of course lack a continuous series covering the entire period since 1960. Still, we have a set of series covering overlapping 5-vintage periods, and the within-series percent changes provide a measure of the change in quality between one vintage and another. For each vintage between 1960 and 2004, there are four separate series

$s(k,v)$ covering the vintage.¹⁴ This allows us four measures of the proportionate change in quality since the previous vintage from 1961 to 2005 (3 for 2006, 2 for 2007, and 1 for 2008).

We can define $\Delta(k,v)$ as this proportionate change, and we can calculate it from log first differences: $\Delta(k,v) = \ln[s(k,v)/s(k,v-1)]$. We estimate the change in quality between adjacent vintages by averaging Δ 's. Because these averages show the percent change, we need to accumulate them to create an index of the level of music quality in each vintage. That is, $I(v) = \sum_{\tau=1960}^v \Delta(\tau)$. This index provides a simply calculated measure of the evolution of vintage quality.

A regression approach generates an analogous index. We can regress the log share on terms in age and vintage dummies. That is, $\ln(s_{t,v}) = f(t-v) + \mu_v + \epsilon_{tv}$, where $f(t-v)$ is a flexible function of the elapsed time between the release date of the music and the calendar year t , μ_v is a vintage effect, and ϵ_{tv} is an error term. In particular, if we define $t-v$ as the age of music in integer years (a), then given that we have multiple years of sales data, we can operationalize $f(\cdot)$ as a full set of age dummies. The index of vintage quality is then the sequence of vintage effects (μ_v). By including a full set of age dummies, the approach identifies the evolution of vintage quality from variation in the log share $s_{t,v}$ among observations of equal age.

4. Structural Interpretation

In addition to its intuitive interpretation, our approach also has a structural random utility interpretation, although our context differs in some respects from the standard product choice model. Normally, one models a consumer choosing among imperfect substitutes – such as

¹⁴ Because the airplay data end in 2008, we have only four series covering 2005, 3 for 2006, 2 for 2007, and one value – $s(0,2008)$ for 2008.

related varieties within some product category – along with an outside good. In this context the inside goods are different vintages of music. The difference here is that because of piracy, we don't observe the total size of market for inside goods. Although sales of recorded music are falling, we don't believe this is because recorded music is falling in utility relative to alternatives. Instead, sales are falling because of increased stealing. Because the data on overall music sales are not informative about the value of music relative to its alternatives, we employ the normalizing assumption that the overall value of music relative to the outside good is constant over time. In the case of airplay this normalizing assumption has the behavioral justification that music stations fill an essentially fixed amount of time with music.

Specifically, in each period t , consumers can choose among music from different vintages v ($v=1960, \dots, t$) or an outside good. The utility of choosing vintage v music at time t is given by: $U_{t,v} = f(t - v) + \mu_v + \epsilon_{t,v}$, where $f()$ is a function describing the depreciation of music as it ages, μ_v is a vintage-specific utility-shifter, and $\epsilon_{t,v}$ is an extreme-value error. The outside good has utility equal to 0; that is, $U_{t,0}=0$. Given this setup, choice probabilities are given by:

$$s_{t,v} = \frac{e^{f(t-v)+\mu_v}}{1+\sum_{v=1960}^t e^{f(t-v)+\mu_v}}, \text{ where } s_{t,v} \text{ is the share of vintage } v \text{ music in period } t \text{'s consumption.}$$

Because $s_{t,0} = \frac{1}{1+\sum_{v=1960}^t e^{f(t-v)+\mu_v}}$, there is a closed-form way to “invert” the market shares. That is, $\ln(s_{t,v}) - \ln(s_{t,0}) = f(t-v) + \mu_v$. Our assumption of a constant utility of music relative to the outside good is equivalent to assuming that $\ln(s_{t,0})$ is constant. This, in turn means that we can rewrite the log shares of inside goods – the shares of each vintage in each year – as

$\ln(s_{t,v}) = A + f(t-v) + \mu_v$, where $A = \ln(s_{t,0})$. Thus, our regression of $\ln(s_{t,v})$ on terms in the age of music and vintage dummies recovers the evolution of “mean utility” with vintage. This is our rationale for describing our vintage dummies as an index of quality.

V. Results

This section presents two groups of results. First, we present our estimates of vintage quality indices, based on both airplay and certification data, which we compare with the critical index. We then use all three of the indices to evaluate whether quality has changed since Napster. It is quite difficult to know how vintage quality would have evolved following 1999 in the absence of both Napster and the other technological changes, so we can’t estimate the effect of Napster *per se*. However, we can quantify the post-Napster experience relative to various counterfactuals. These include: a) relative to levels defined by 5, 10, or N years prior to Napster, b) the trends implied by the 5, 10, or N years prior to Napster.

1. Airplay Data Results

Before turning to regressions we first report the indice $I(v)$ calculated from airplay data (as described above). Figure 8 reports the resulting index, and it rises sharply from 1960 to 1970, then fall as sharply until the mid-1980s. It then increases slightly in the mid-1990s, followed by a decline through 1999. Following 2000, the index rises sharply, reaching a level last experienced in the mid-1970s.

Table 1 reports regressions of $\log(s_{t,v})$ on terms in age and a full set of vintage dummies. The first column includes first and second order terms in age. The second column adds a cubic

term. The age coefficients from a regression with a spanning set of age dummies is reported in Figure 9a. The coefficients give rise to a smooth and monotonic depreciation pattern. After 10 years, songs receive roughly a quarter as much ($0.25 \approx e^{-1.5}$) airplay as during the year they are released.

Figures 10a-10c show the vintage indices derived from the coefficients on the vintage dummies in the regressions. All three of these indices strongly resemble Figure 8. Quality rises from 1960 to 1970, then falls to at least 1985. In all three specifications, the vintage quality index rises substantially after 1999.

2. Certification Data

The latter three columns of Table 3 report regressions of $\log(s_{t,v})$ of terms in age along with vintage effects using album certification data, and Figure 9b shows the flexibly estimated age effects. As expected – given the lumpy and sparse nature of the certification data – the album certification depreciation pattern is less smooth than the airplay pattern. Because of data sparseness, we include all formats (albums, singles, and other media) to increase precision. Figures 11a-11c show the resulting vintage quality indices for 1970-2010, based on quadratic, cubic, and flexible specifications. All three show relatively steady decline from 1970 to about 2000. The vintage quality indices then rise until about 2008.

3. Post-Napster Changes

In order to ascertain the effect of the changes in technology surrounding Napster on the volume of high quality music brought forth by the industry, we would ideally compare the world experiencing the changes to an otherwise similar environment not experiencing the same shocks

to demand and supply. Unfortunately, we lack such a “control” for comparison with our “experiment.” We can still pursue the more modest goal of asking whether the volume of high quality music has changed since Napster, using a few different benchmarks.

First, we can ask whether the level of the index changed following Napster. This comparison is, of course, sensitive to the amount of pre-Napster time included in the calculation, so we perform the calculation with various starting times. Second, we can ask whether the time trend following Napster deviates from the time trend defined prior to Napster. This approach, too, depends on the number of pre-Napster years used for defining the pre-existing time trend. While such approaches do not allow us to ascertain the causal impact of even the compound experiment brought about by the various technological changes surrounding Napster, they do allow us to quantify what has happened to the amount of consequential new music brought to market. Particularly against the backdrop of the music industry’s stated concerns about piracy threatening its ability to bring music to market, even this more modest goal can shed useful light on our understanding of whether the sharp reductions in revenue to recorded music have undermined the flow of new products.

Tables 2–4 report these regressions, for critic-, airplay-, and certification-based indices, respectively. In each table the first set of columns compares the post-Napster level of an index to its level for various durations prior to Napster (1995-1999, 1990-1999, etc). The second set of columns compares the post-Napster trend to the time trend defined for various pre-Napster periods. Not surprisingly, in light of the figures already reported, the critical index gives a somewhat different result from the usage-based indices.

As Table 2 indicates, relative to the entire pre-Napster period, the post-Napster level of the critical index is 23 percent below, and this difference is statistically significant. The post-Napster critical index is below all pre-Napster periods, although this difference is statistically significant only for comparisons with pre-Napster periods beginning in 1970 or earlier. The deviations between the post-Napster trend and the pre-Napster trends (defined with various starting points) are all statistically insignificant. While the post-Napster critical index is lower than the level prior to Napster, this is largely attributed to the peak that occurred in 1970. Relative to various pre-Napster trends, there is no evidence of a decline in quality in the period since Napster.

Table 3 repeats this exercise using the airplay-based index. While column (1) indicates that the post-Napster level is below the average for the entire pre-Napster period (1960-1999), the remaining columns of the first half show that the post-Napster airplay index is statistically significantly above the averages for the decade immediately prior to Napster. The latter half of Table 3 shows that relative to all pre-Napster trends, the airplay-based quality index has a positive and statistically significant time trend.

Finally, Table 4 reports results of this exercise using the certification-based index. Relative to the various pre-Napster periods, the post-Napster level of the certification-based index is generally above its pre-Napster level, and the difference is statistically significant relative to all pre-Napster periods (except the period beginning in 1970). The latter half of the table shows that, relative to all pre-Napster time trends, the post-Napster trend deviates positively and significantly

VI. Discussion

The indices derived in this paper, both based on critics and those from airplay and sales data by vintage and time, are rather similar to one another. All show increases in vintage quality through the 1960s to 1970, declines to the mid-1980s, followed by relatively flat periods. Finally, none show declines – and the two usage-based indices show substantial increases – following 2000. The lack of decline is somewhat puzzling against the backdrop of the sharp decline in revenue since Napster. It is costly to bring new music to market, and one might share the recording industry's expectation that a sharp reduction in revenue would reduce the amount of new music brought to market. A possible resolution to the puzzle is the observation that as some new technologies have reduced revenue, other new technologies have reduced the cost of bringing new music to market.

Bringing new music to market has three major activities: creation, promotion, and distribution. New technologies have sharply reduced the costs of each of these. Creation entails both composition activity as well as recording, mixing, engineering, and manufacturing. Many aspect of creation were traditionally expensive, but new technologies have changed this. As Kalmar (2002) notes, with the development of digital audio tape in 1987, “a label can set up their own recording studio for about five grand.”¹⁵ Costs have continued to decline in the last decade: Software such as Pro Tools, which sells for roughly \$100, turns an inexpensive personal computer into a home recording studio .¹⁶

¹⁵ See Kalmar (2002), p. 73.

¹⁶ See Donald Bell, “Avid Introduces new Pro Tools Studio Bundles.” CNET, Oct 1, 2010 (http://news.cnet.com/8301-17938_105-20018292-1.html, accessed October 28, 2010).

Music is an experience good, and consumers need to become aware of music to be interested in purchasing it. Record companies have traditionally made consumers aware of their products by promoting their new releases on radio. Even prior to the Internet, the labels produced more music than radio stations could air, so the labels paid the stations to promote their music. While the literal practice of “payola” was outlawed in 1960, labels continued to pay for airplay through independent promoters, and payments for their services were substantial: in 1985 the record labels collectively paid \$65 million for airplay when the industry’s pre-tax profit was \$200 million. The cost of promoting a hit single was about \$150,000.¹⁷

In the past decade, the way that consumers learn about new music has changed substantially. Where radio used to be the main means for discovering new music, consumers now learn about new music from a variety of web sources, including Pandora, MySpace, and YouTube. Over half of young consumers (age 12-34) use the Internet for learning about new music, while only 32 percent use radio, according to the 2010 “Infinite Dial” study conducted by Edison Research and Arbitron. Just over a quarter (27 percent) of the population 12 and over had used Internet radio in the previous month, and Pandora was the most recognized Internet radio site. Among those who had ever listened to Internet radio, 28 percent named Pandora, followed by Yahoo Music (9 percent), AOL Radio (6), and Last.fm (4).¹⁸ The Internet appears to have undermined the scarcity of terrestrial radio stations as music promotion channels.

The Internet has also substantially changed music distribution. Many factors, including the need to get a large quantity of physical product into many stores before popularity waned,

¹⁷ Caves (2000), p. 292, provides the source for the quote and the data cited in this paragraph.

¹⁸ See “The Infinite Dial 2010: Digital Platforms and the Future of Radio.”

(http://www.fmqb.com/goout.asp?u=http://www.edisonresearch.com/home/archives/2010/09/the_american_youth_study_2010_part_one_radios_future.php accessed October 28, 2010).

avored large-scale enterprises prior to the Internet. Music can now be distributed electronically, eliminating inventory and transportation costs. Using TuneCore's service, for example, an artist can make his song available on iTunes for \$9.99.¹⁹

Some observers have argued that these reductions in cost have made it possible for smaller-scale organizations to bring music to market. Even if major recording labels are now less able to recoup returns from their investments, independent labels may now play a larger role in bringing music to market. Do the data support the contention that independent labels are bringing forth more of the supply following Napster? Pitchfork Media's ranking of the top albums of the 1980s, 1990s, and the 2000s includes each album's issuing label or, more commonly, a less recognized entity that may be either an independent label or a sub-label of one of the majors. Using mostly Wikipedia entries, I have been able to code each of the labels on the top 100 albums of each decade as either a major or an independent. This is not a trivial task, as the major owners produce records under a long list of label imprints. The data provide support for the idea that independent labels are playing an increasing role (see Figure 12). While the share of the top 100 on independent labels was 50 percent in both the 1980s and the 1990s, it rose to 60 percent in the period since 1999.²⁰ This difference (between the 2000s and the previous two decades) is significant at the 5 percent level in a one-sided test ($p\text{-val} = 0.04$). The ascendance of independent labels has been noted elsewhere.

¹⁹ See <http://www.tunecore.com/>, accessed October 28, 2010. At the site: "What Does Worldwide Distribution Cost" \$9.99 per single, \$9.99 per ringtone, \$49.99 per album."

²⁰ Pitchfork's focus on artists they view as interesting likely explains the high share of independent label releases among their most highly rated albums. According to Leeds (2005), independent labels' collective share of recorded music revenue rose to 18 percent (27 percent including indie albums distributed by majors) in 2005, its highest share in 5 years.

Pitchfork has is disproportionately focused on independent, rather than mainstream, music. It would be useful to see how the independent share has evolved for music reaching larger and more mainstream audiences. To this end I calculate the independent share among the top-selling 200 US albums, on the yearend Billboard 200, for 2002-2010 (see Figure 13). The share of albums on independent labels increases from 1.5 percent in 2002 to 7.5 percent in 2010.²¹ The share among the top 100 has risen from 4 to 12 percent. It thus appears that independent labels are accounting for a growing share of successful albums, using various measures of success.

VII. Conclusion

We have presented evidence, from three independent approaches, showing clearly that the quality of new recorded music has not fallen since Napster. While it may well be true that the recording industry has experienced substantial declines in its revenue and perhaps its profitability as well, there is no evidence that consumers have suffered from a withdrawal of creative effort. The flow of products appears to as strong as before, if not stronger. Despite these emerging conclusions, two important caveats are in order. First, it is entirely possible that absent the weakening of effective copyright protection, the other changes in technology might have ushered in an era of even greater creative output. It is impossible to say whether creative output is as high as it would have been without piracy. However, it is clear that creative output in recorded music is as high, or higher, than it was prior to Napster. While the period since Napster may be a period of unusually low revenue to recorded music (relative to history), it is

²¹ In addition to reporting yearend top-200 albums by sales, Billboard also reports separate lists of the top-selling albums from independent labels, making it possible to calculate the share of top-selling albums from independent labels.

not a period of unusually low quantities of consequential output. A second important caveat is that while new music supply appears robust despite changes in technology, it is difficult to say whether this finding would carry over to other contexts, such as motion pictures, where bringing products to market is far more costly.

Much of the debate over appropriate copyright policy in the digital era has focused on the effect of Napster on firms' ability to appropriate revenue. Revenue is, to be sure, important for financing the flow of new products; but revenue is a means toward the end of assuring continued production of new creative works. Emerging results on the continued availability of new recorded music products suggests that researchers and policymakers thinking about the strength of copyright protection should supplement their attention to producer surplus in creative industries with a concern for consumer surplus as well.

References

- Blackburn, David. "On-line Piracy and Recorded Music Sales." Unpublished Manuscript, Harvard University, December 2004.
- Blashil, Pat (coordinator for Zagat Survey). Music Guide: 1000 Top Albums of All Time. 2003. New York: Zagat Survey.
- Boldrin, Michele and David K. Levine. Against Intellectual Monopoly. Cambridge University Press. 2008.
- Caves, Richard E. Creative Industries: Contracts between Art and Commerce. Harvard University Press: Cambridge, MA. 2000.
- Connolly, Marie & Krueger, Alan B., 2006. "Rockonomics: The Economics of Popular Music," Handbook of the Economics of Art and Culture, Elsevier.
- Grein, Paul (May 14, 1989). "New Golden Rule: 500,000 Sales Mark for All Singles". Los Angeles Times. http://articles.latimes.com/1989-05-14/business/fi-188_1_singles-riaa-platinum. Retrieved 12 November 2010.
- Handke, Christian. "Plain Destruction or Creative Destruction? Copyright Erosion and the Evolution of the Record Industry." Review of Economic Research on Copyright Issues, 2006, vol 3(2): 29-51.
- Handke, Christian, "Digital Copying and the Supply of Sound Recordings." Unpublished Manuscript, 2010, available at <http://www.acei2010.com/upload/acei/handke.pdf>.
- Kalmar, Veronika. Label Launch: A Guide to Record Recording, Promotion, and Distribution. St. Martin's Griffin: New York. 2002.
- Knopper, Steve. Appetite for Self-Destruction: The Spectacular Crash of the Record Industry in the Digital Age. Free Press: New York. 2009.
- Larkin, Colin. "A Brief History of Pop Music." In The Encyclopedia of Popular Music, 5th edition, edited by Colin Larkin, Omnibus Press/MUZE: London and New York. 2007.
- Leeds, Jeff. "The Net is a Boon for Indie Labels." New York Times. December 27, 2005.
- Levy, Joe and Editors of Rolling Stone, The 500 Greatest Albums of All Times. Wenner Publishing. November 16, 2005
- Liebowitz, Stan J. "File Sharing: Creative Destruction or Just Plain Destruction?" Journal of Law and Economics, vol. 49, no. 1, April 2006, pp. 1-28.
- Oberholzer-Gee, Felix and Koleman Strumpf. "The Effect of File Sharing on Record Sales: An Empirical Analysis." [Journal of Political Economy](#). Volume 115, Issue 1, Page 1–42, Feb 2007

Oberholzer-Gee, Felix and Koleman Strumpf. [File-Sharing and Copyright](#). *NBER's Innovation Policy and the Economy series*, volume 10. ed. Joshua Lerner and Scott Stern. MIT Press. 2009.

Mortimer, Julie Holland, Chris Nosko, and Alan Sorenson. "Supply Responses to Digital Distribution: Recorded Music and Live Performances." NBER Working Paper No. 16507. October 2010.

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

Rob, Rafael and Joel Waldfogel. "Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students." NBER Working Paper 10874, October 2004.

Shapiro, Carl and Hal Varian. *Information Rules*. Harvard Business School Press. Cambridge, MA. 1999.

Soloveichik, Rachel. "Music As a Capital Asset." Unpublished Paper. Bureau of Economic Analysis. 2011.

Waldfogel, Joel. "Bye, Bye, Miss American Pie: The Supply of New Recorded Music since Napster." NBER Working Paper . 2011.

Zentner, Alejandro, "Measuring the Effect of File Sharing on Music Purchases." *The Journal of Law and Economics*. Volume 49, Issue 1, Page 63–90, Apr 2006.

Table 1: Regression Estimates of Depreciation

	(1) airplay	(2) airplay	(3) certifications	(4) certifications
Age	-0.1897 (0.0311)**	-0.1557 (0.0616)*	-0.2515 (0.0119)**	-0.4049 (0.0224)**
Age squared	0.0020 (0.0006)**	-0.0001 (0.0026)	0.0036 (0.0004)**	0.0140 (0.0017)**
Age cubed		0.0000 (0.0000)		-0.0002 (0.0000)**
Constant	2.6590 (0.5323)**	2.6493 (0.5381)**	-3.3120 (0.5955)**	1.5280 (1.1789)
Observations	235	235	868	868
R-squared	0.99	0.99	0.77	0.80

Notes: Dependent variable is the log vintage share in a year. All regressions include vintage fixed effects (coefficients not shown). Robust standard errors in parentheses. * significant at 5% level; ** significant at 1% level.

Table 2: The Post-Napster Critical Album Index Relative to Pre-Napster Levels and Trends

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-Napster Level	-0.2295 (0.1099)*	-0.2026 (0.1628)	-0.2380 (0.1355)	-0.1790 (0.1177)	-0.2849 (0.1064)*				
Level since 1995		-0.0310 (0.1370)							
Level since 1990			0.0116 (0.1052)						
Level since 1980				-0.1064 (0.0914)					
Level since 1970					0.2631 (0.1064)*				
Post-Napster Trend						0.0450 (0.0697)	-0.0130 (0.0430)	-0.0196 (0.0315)	-0.0248 (0.0283)
Trend since 1995						-0.0498 (0.0406)			
Trend since 1990							-0.0101 (0.0160)		
Trend since 1980								-0.0050 (0.0067)	
Trend since 1970									-0.0029 (0.0045)
Constant	1.9271 (0.0458)**	1.9312 (0.0497)**	1.9240 (0.0540)**	1.9831 (0.0663)**	1.7194 (0.0945)**	1.9365 (0.0484)**	1.9308 (0.0518)**	1.9434 (0.0593)**	1.9510 (0.0718)**
Observations	46	46	46	46	46	46	46	46	46
R-squared	0.09	0.09	0.09	0.12	0.20	0.09	0.07	0.07	0.07

Notes: Dependent variable is critic-based vintage index. Standard errors in parentheses. * significant at 5% level; ** significant at 1% level.

Table 3: The Post-Napster Airplay-Based Sales Index Relative to Pre-Napster Levels and Trends

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-Napster Level	-0.2231 (0.2340)	0.4875 (0.3277)	0.4822 (0.2346)*	0.3790 (0.1126)**	0.0032 (0.1944)				
Level since 1995		-0.8151 (0.2814)**							
Level since 1990			-0.9484 (0.1873)**						
Level since 1980				-1.2359 (0.0899)**					
Level since 1970					-0.9806 (0.1944)**				
Post-Napster Trend						0.3223 (0.1350)*	0.2391 (0.0747)**	0.2239 (0.0393)**	0.2032 (0.0241)**
Trend since 1995						-0.2165 (0.0827)*			
Trend since 1990							-0.1170 (0.0301)**		
Trend since 1980								-0.0767 (0.0094)**	
Trend since 1970									-0.0595 (0.0043)**
Constant	2.8434 (0.1013)**	2.9479 (0.1007)**	3.0866 (0.0948)**	3.4772 (0.0644)**	3.5977 (0.1705)**	2.9097 (0.1005)**	2.9942 (0.0985)**	3.2408 (0.0822)**	3.5293 (0.0680)**
Observations	48	48	48	48	48	48	48	48	48
R-squared	0.02	0.17	0.38	0.81	0.37	0.13	0.25	0.60	0.81

Notes: Dependent variable is airplay-based vintage index. Standard errors in parentheses. * significant at 5% level; ** significant at 1% level.

Table 4: The Post-Napster Certification-Based Sales Index Relative to Pre-Napster Levels and Trends (All Recorded Music Products)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-Napster Level	0.0801 (0.1008)	0.3515 (0.1411)*	0.2817 (0.1112)*	0.2383 (0.0752)**				
Level since 1995		-0.3257 (0.1262)*						
Level since 1990			-0.3024 (0.0963)**					
Level since 1980				-0.4745 (0.0752)**				
Post-Napster Trend					0.1629 (0.0525)**	0.1099 (0.0299)**	0.0964 (0.0186)**	0.0938 (0.0151)**
Trend since 1995					-0.0923 (0.0340)**			
Trend since 1990						-0.0422 (0.0133)**		
Trend since 1980							-0.0263 (0.0053)**	
Trend since 1970								-0.0233 (0.0036)**
Constant	0.3248 (0.0504)**	0.3791 (0.0515)**	0.4256 (0.0556)**	0.6411 (0.0614)**	0.3630 (0.0480)**	0.3940 (0.0502)**	0.4987 (0.0531)**	0.6695 (0.0651)**
Observations	40	40	40	40	40	40	40	40
R-squared	0.02	0.17	0.22	0.53	0.22	0.27	0.44	0.56

Notes: Dependent variable is certification-based vintage index. Standard errors in parentheses. * significant at 5% level; ** significant at 1% level.

Figure 1




Figure 2




Figure 3




Figure 4a




Figure 4b




Figure 5




Figure 6a




Figure 6b




Figure 7

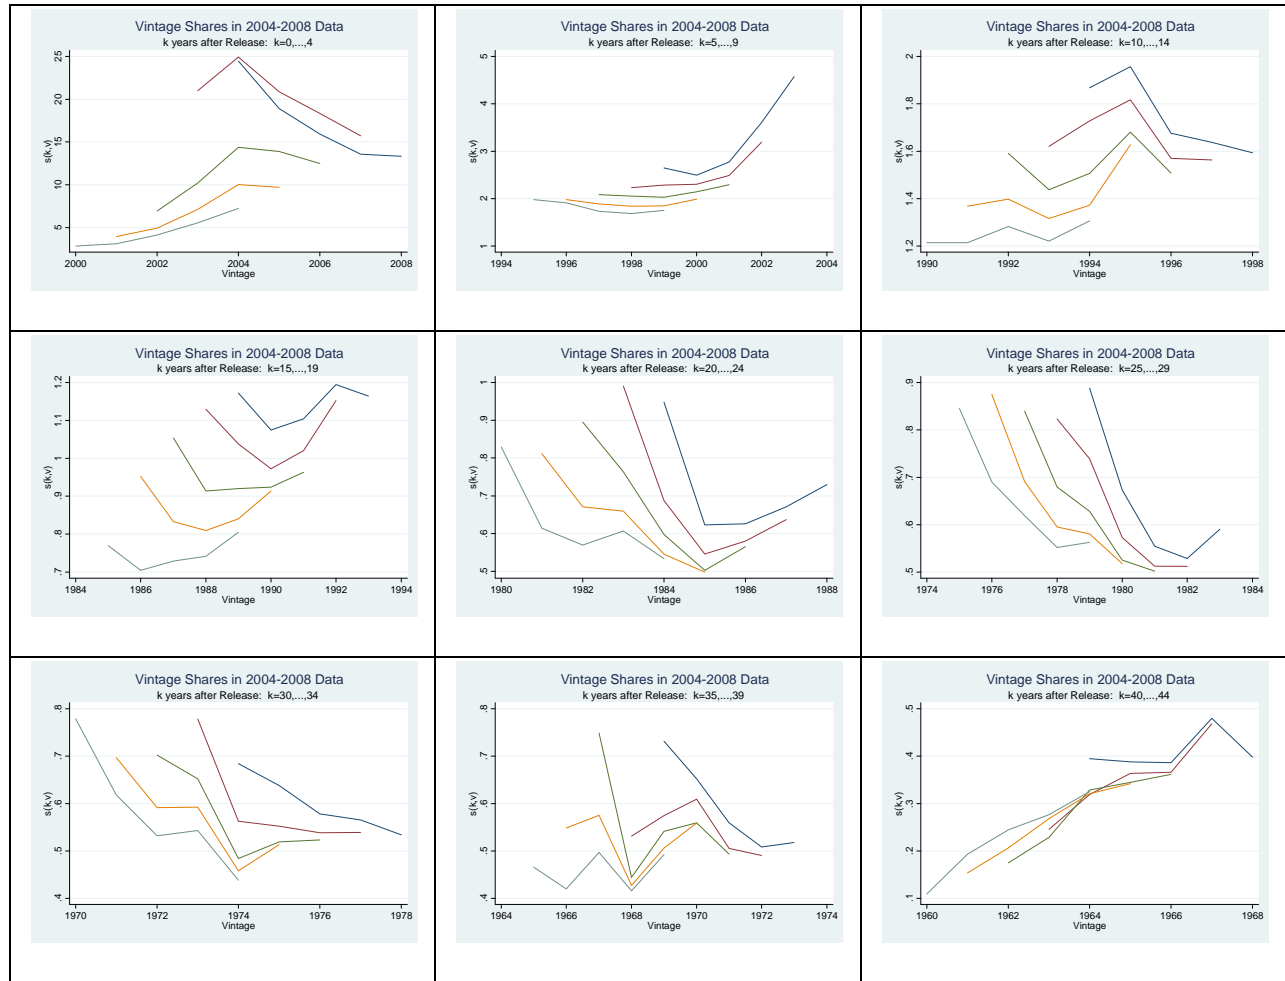


Figure 8




Figure 9a




Figure 9b




Figure 10a




Figure 10b




Figure 10c




Figure 11a




Figure 11b




Figure 11c




Figure 12




Figure 13

