

Does Chatter Matter?
The Impact of User-Generated Content on Music
Sales

Vasant Dhar
Elaine Chang
Leonard N. Stern School of Business
New York University
May 2007

Abstract:

The Internet has enabled the era of user-generated content, potentially breaking the hegemony of traditional content generators as the primary sources of “legitimate” information. Prime examples of user-generated content are blogs and social networking sites, which allow easy publishing of and access to information. In this study, we examine the usefulness of such content, consisting of data from blogs and social networking sites in predicting sales in the music industry. We track the changes in online chatter for a sample of 108 albums for four weeks before and after their release dates. We use linear and nonlinear regression to identify the relative significance of online variables on their observation date in predicting future album unit sales two weeks ahead. Our findings are as follows: (a) the volume of blog posts about an album is positively correlated with future sales, (b) greater increases in an artist’s Myspace friends week over week have a weaker correlation to higher future sales, (c) traditional factors are still relevant – albums released by major labels and albums with a number of reviews from mainstream sources like *Rolling Stone* also tended to have higher future sales. More generally, the study provides some preliminary answers for marketing managers interested in assessing the relative importance of the burgeoning number of “Web 2.0” information metrics that are becoming available on the Internet, and how looking at interactions among them could provide predictive value beyond viewing them in isolation. The study also provides a framework for thinking about when user-generated content influences decision making.

1. Introduction

An increasing portion of information on the Internet is becoming user-generated. It is becoming more commonplace to find user-generated information on virtually any topic of interest on the Internet, which is increasingly becoming the repository for global information. Understanding how people are using this information and are influenced by it is an active area of research.

In this research, we have investigated the impact of user-generated content, defined as the conjunction of blogs and social networking sites, on music sales.¹ The open question, one that is also being investigated in a number of areas, is whether user-generated content provides any predictive value, or whether it is largely retrospective in nature, or just plain noise. Financial markets are an obvious area of interest, as are books, movies, music, travel, and a host of other areas including politics² where large numbers of people weigh in with opinions or data. The research question we ask in this research is the following: *does user-generated content provide any predictive value for online music sales?* It also assesses the significance of user-generated content as a predictor compared to the more traditional information sources such as professional reviews in print or electronic media.

The rest of the paper is organized as follows: we begin by reviewing prior research in

¹ Blogs are short for “web logs” – websites that tend to be written in the style of a journal, with posts arranged in reverse chronological order. Some blogs are little more than online diaries, while others are almost collections of essays on weighty topics like politics, religion, or health. The growing importance of blogs has been well-noted by the media, with surveys estimating that a staggering 30% of the American population considers blogs an important source of information. Social networking sites enable users to create profiles and make connections to other users who live in the same area, share similar interests, or simply seem interesting. These networks generally allow users to create a public list of other users that are mutual friends – that is, both users have listed each other as a friend. One of the most important social networks to the music industry is Myspace (www.myspace.com).

² South Korean President Roh Moo-hyun was taken from virtual obscurity into the presidency based mainly on the support of young voters and their use of blogs and Internet chat rooms. See: http://curlyjoe.gnn.tv/headlines/9334/South_Korean_Anti_Americanism

Section 2. We then describe our data sample, the variables included in the data analysis, and the sources used to gather the data. Section 4 discusses our methods of analysis and the empirical results. We conclude the paper in Section 5 with a summary and analysis of our results, including questions that this study leaves unanswered. We end by proposing topics worth further research in Section 6.

2. Background and Prior Research

People rely on opinions of others for “experience goods” – products whose quality is difficult to observe or sample adequately before purchase. For such products, consumers often rely on others as an input into their decision. Experience goods include a fairly broad range of areas including books, movies, physical locations, and even various kinds of advice, including medical and financial.

A couple of studies have analyzed Internet chatter in Finance. Tumarkin and Whitelaw³ looked for a relationship between Internet postings and stock prices on the Raging Bull discussion forum. They found that high discussion activity was correlated with abnormal market returns, but did not predict future returns. In a similar study, Antweiler and Frank⁴ found the level of activity to be correlated with volatility, but of no predictive value on prediction market direction. These preliminary studies suggest that the *level* of activity on discussion groups in financial markets has no predictive value, but they leave open the possibility of a relationship between the *content* (i.e. semantics) of postings and future returns. This is currently an open area of research, with an early study looking at the relationship between news content, where a new story is classified into a type

³ Tumarkin, R., and R.F. Whitelaw. News or Noise? Internet Postings and Stock Prices. *Financial Analysts Journal*, pages 41-51, 2001.

⁴ Antweiler, W., and M.Z. Frank. Is all that talk just noise? The Information Content of Internet Stock Message Boards. *Journal of Finance*, 59(3): 1259-1295, 2004.

such as earnings announcement or SEC investigation, and stock prices.⁵

There is a fair amount of research examining the influence of traditional and user-generated content in the movie and book industries. Movie ratings tend to be quite structured, in terms of numerical scores or like/dislike data. Book reviews tend to be more complex for obvious reasons.

Eliashberg and Shugan (1997) found statistical evidence that the percentage of positive and negative professional movie critics' reviews were predictive of total cumulative box office sales, though not of early box office sales (sales in the first four weeks).⁶ In addition, they found that the number of reviews is a significant predictor of the first week's box office sales, but not of the movie's performance over its entire lifecycle. Moul (2006) examined a sample of theatrical movie admissions from 1990-1996, and found that approximately 10% of the variation in implied consumer expectations among movies is due to the effect of word of mouth among moviegoers.⁷

While the previous studies used data before the Internet became ubiquitous in the late 1990s, other researchers have since examined online user-generated content about movies. Dellarocas et al. (2004) explored more specifically the usefulness of online product ratings in predicting movie revenues.⁸ The researchers first conducted a survey asking a random sample of people who did not rate movies online to rate a number of movies, and then compared the results to the online movie review data they gathered. The correlation between the ratings given by the two

⁵ Mc Macskassy, S, Hirsh, H., Provost, F., Sankaranarayanan, R., Dhar, Vasant., Intelligent Information Triage, The 24th Annual International Conference on Research and Development in Information Retrieval (SIGIR), September 2001.

⁶ Eliashberg, Jehoshua and Steven Shugan (1997). "Film Critics: Influencers or Predictors?" *Journal of Marketing*, Vol. 61, No. 2, 68-78.

⁷ Moul, Charles (2006). "Measuring Word of Mouth's Impact on Theatrical Movie Admissions." *Journal of Economics and Management Strategy*, forthcoming.

⁸ Dellarocas, Chrysanthos, Neveen Farag Awad, and Xiaoquan (Michael) Zhang (2004). "Exploring the Value of Online Product Ratings in Revenue Forecasting: The Case of Motion Pictures." Working paper, University of Maryland. May 2004.

populations was significant, suggesting that online consumer reviews are a valid proxy for opinions about a movie in the larger population. Dellarocas et al. also found that weighted averages of user ratings were more informative in predicting future movie sales than professional critic reviews.

Duan et al. (2005) also explored the relationship of online user ratings to movie revenues.⁹ They found that the average online user review rating does not have a significant impact on box office revenues. The number or volume of online postings, however, is significant. The researchers suggest that online ratings do not influence consumers because consumers can make an independent judgment about the movie's quality by drawing upon many different sources. Increases in the number of postings, however, do matter. A possible explanation is that the volume of postings is indicative of the interest level in a movie, regardless of what the ratings say.

Chevalier and Mayzlin (2006) analyzed book sales, examining the effects of online consumer ratings on book sales at Amazon.com and Barnes and Noble's website, BN.com.¹⁰ They found that additional positive reviews at Amazon.com would increase sales at Amazon.com relative to BN.com and vice versa. Additional evidence suggested that very negative reviews had a greater impact on sales than very positive reviews.

Gruhl et.al (2005) analyzed blog mentions to books in from IBM's WebFountain project¹¹ that contains a very large number of blogs, which grow at roughly 200K blogs per day.¹² They found that it is difficult to predict the direction of sales rank for books based on blogs, but it is often

⁹ Duan, Wenjing, Bin Gu and Andrew B. Whinston (2005). "Do Online Reviews Matter? – An Empirical Investigation of Panel Data." Working Paper, the University of Texas at Austin. January 2005.

¹⁰ Chevalier, Judith and Dina Mayzlin (2006). "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research*, Vol. 43, No. 3, 345-354.

¹¹ Gruhl, D., L. Chavet, D. Gibson, J. Meyer P. Pattanayak, A. Tomkins, and J. Zien. How to Build a Webfountain: An Architecture for Very-Large Scale Text Analytics. *IBM Systems Journal*, 43(1):64-77, 2004.

¹² Gruhl, D, R. Guha, Ravi Kumar, Jasmine Novak, and Andrew Tomkins. The Predictive Power of Online Chatter. *KDD 2005*, Chicago, IL, August 2005.

possible to predict *spikes* in sales ranks based on the volume of blog mentions.

What should we expect to find in the music space relative to the various findings in the literature on books, movies, or financial advice? Which of these should we expect it to be similar to and why?

Consider books and movies. In both cases recommender systems have been successful. Movie recommender systems are able to accurately classify an individual into group that has similar tastes and make recommendations based on what the tastes of that cohort. For movies, a deeper analysis of the movie isn't typically warranted. For books it often is, in part because of the higher level of commitment involved and even the need for a close fit to what one is looking for. For these reasons, it is important to find a good fit with the reviewer.

As with books, interpreting blogs and content on social networking sites about music is hard. Without extensive prior knowledge about the raters, the reader doesn't know which opinions are from people who are "similar" and should therefore be weighted more heavily. On the other hand, if there is a large volume of blogs writing about an album, chances are that the album is creating some buzz.

Our hypothesis is that blogs and social networks matter independently for the following reasons. First, we believe that a considerable amount of effort goes into writing good blogs, and their authors feel passionate enough about the topic to spend the time writing and sharing them with others. Readers recognize good blogs and pay attention to them. Secondly, while blogs are largely unfiltered, some blog sites tend to have more authority than others. In other words, reputation helps them attract traffic that is in turn influenced by their content.

We also hypothesize that social networks matter. In the music industry, Myspace (www.myspace.com) provides a special music category that allows artists to create profile pages

including band biographies, to upcoming tour dates and streaming music tracks. Through these band profiles, individual Myspace users can advertise artists that they like to their friends, simultaneously bookmarking¹³ the artist's work for later perusal and promoting the artist to their other friends. The number of friends a band has is displayed on its Myspace page is like a public badge of popularity. We would expect, on average, that a band with thousands of friends on Myspace to be more popular with Myspace users than a band with a handful.¹⁴

The methodology used in this study is as follows. We gathered data tracing the changes in user-generated content for an album by tracking the volume of blog chatter, the number of friends an artist has on Myspace, and album reviews for four weeks before and after the release date. We control for the influence of external differences in promotion budgets and so on by introducing a dummy variable for whether an album is released by a major or independent label. We construct measurable indicators for user-generated content as well as traditional content with the intention of understanding their relative significance on music sales. We use blog chatter and the extent of social network connectivity as the proxy for user-generated content. We then use a linear regression and a regression tree to examine the relative significance of the variables in predicting album unit sales two weeks ahead.

3. Data

3.1 The sample

Our data consist of album statistics and data collected from publicly available information on websites. The data were either manually recorded or downloaded through automated Javascript

¹³ Bookmarking on the Internet is becoming increasingly common with several sites devoted exclusively to it, such as del.icio.us. Bookmarks can point to traditional or user-generated content.

¹⁴ There are obvious exceptions. Some established artists have no Myspace presence at all.

programs. We compiled the sample of music albums by collecting the names of albums released in the United States between January 16 and March 6, 2007 from Pause & Play (www.pauseandplay.com), a website devoted solely to listing upcoming album releases. Old material, such as reissues and compilations, were excluded from the sample. We further cross-checked the release date given by Pause & Play with Amazon's page for the album in order to verify that the record label had not moved the release date, and if the album did not have a corresponding page on Amazon, it was eliminated from the sample. The final sample consisted of a total of 108 albums.

The complete data set is 864 observations (108 albums x 8 weeks). The observation dates are 23, 16, 9, and 2 days before release and 5, 12, 19, and 26 days after release. Within this dataset, there are therefore two observations for each album for a total of 216 observations of pre-release chatter and post-release sales where it is safe to say that sales after release are not influenced by *post release* chatter or result in increased post release chatter. These 216 points observe chatter 9 and 2 days before the release date and corresponding forward sales 5 and 12 days after release respectively. The results from the subset of 216 observations represent the true relationship between chatter and subsequent sales. We present our analysis on both sets of data and comment on some of the differences and their potential implications.

3.2 Dependent variable: album sales

Since information on digital music sales is difficult to obtain, we focused on physical CD sales. Although digital music sales have grown tremendously over the past few years, physical

album sales still account for 85% of the music market.¹⁵ Nielsen SoundScan would have been the ideal source for album sales data, as it is the industry standard tracking system for sales of music products in the US, but its data are proprietary and very expensive to obtain.

We compute album sales based on Amazon.com sales ranks. Amazon is one of the largest online CD retailers and its sales ranks are easily observed.¹⁶ Each product page on Amazon lists the product's sales rank relative to all products in its category. Sales ranks order albums according to quantity: the highest-selling album is ranked #1, the next highest #2, and so on. In addition, Amazon allows consumers to preorder, or purchase products that are announced far ahead of the actual release date. Products available for preorder are also given sales ranks; in notable instances, such as the Harry Potter series, products have made it to #1 on Amazon without even being released.¹⁷ The preorder feature on Amazon allows us to separate chatter prior to the product being experienced and evaluated from the chatter that follows after release. It should be noted that if an album's sales are lower than a certain cutoff point, Amazon does not even calculate the sales rank, simply giving the album a ranking of "none". The cutoff rank for each category is not published, but the ranks in our sample ranged from 1 to 603,638 and followed an exponential decay pattern with fewer albums receiving very high ranks. For the purposes of data analysis, we set such "none" rankings to 700,000.

We converted weekly sales ranks into weekly sales quantities using a technique previously developed by Brynjolfsson et al. (2003) and Chevalier and Goolsbee (2003).¹⁸ Both studies found

¹⁵ "Can Music Survive Inside the Big Box?" *Wall Street Journal*. April 27, 2007. B1.

¹⁶ "Apples iTunes Slams Into Top Ten Music Retail Chart." *The Register*. November 22, 2005. http://www.theregister.co.uk/2005/11/22/itunes_us_retail_top_ten/.

¹⁷ "Where Harry Most Enchants." *The Washington Post*. April 17, 2007. B01.

¹⁸ Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith (2003). "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers." *Management Science*, Vol 49, No. 11, 1580-96; Chevalier, Judith and Austan Goolsbee (2003). "Measuring prices and price competition online: Amazon and Barnes and Noble." *Quantitative Marketing and Economics*, Vol. 1, No. 2, June 2003, 203-222.

ample evidence of the following Pareto relationship between the ordinal sales rank of a book and its weekly sales quantities:

$$\ln Quantity = \alpha + \beta * \ln Rank$$

For the α parameter, we used the 10.526 calculated for book sales on Amazon by Brynjolfsson et al. (2003). For the β parameter we used the value of -1.61 that has been calculated for DVDs.¹⁹ Since these parameters are not calibrated specifically to the music category on Amazon, the actual estimates for the parameters could be different. As we describe later in the paper, however, varying the parameters does not have a material impact on the results.

3.3 Independent variables

3.3.1 Days since release

In general, we would expect that music sales would follow the pattern of movie revenues, with sales highest right after the release date. As the pool of potential consumers decreases, we would expect sales to fall.

To account for the effect of time on sales, we calculated “days since release” by subtracting the release date from the observation date. Interestingly enough, virtually all of the CDs in the sample were released on a Tuesday, which may be related to when Billboard chart information is released or when music stores receive their shipments.

3.3.2 Major label versus independent label releases

Four major record groups dominate the music industry: Universal Music Group, Sony

¹⁹ Ghose, Anindya, Michael Smith, and Rahul Telang. “Internet Exchanges for Used Books: An Empirical Analysis of Product Cannibalization and Welfare Impact.” *Information Systems Research*, Mar. 2006, pp. 3-19.

BMG Music Entertainment, EMI Group, and Warner Music Group. Combined, these four groups account for an estimated 82% of the music market.²⁰ Independent labels make up the remaining 18%. Since albums released by major labels are more likely to have larger promotion budgets and are often by more popular and mainstream artists, we accounted for the “major label effect” by introducing a dummy variable where 1 corresponded to a major label release, and 0 to an independent label release.

3.3.3 Average number of reviews and average rating

A major source of information for consumers about an album is CD reviews. The level of buzz around an album can be measured by the volume of reviews an album received, and its quality can be measured by how positive the ratings are on average. We gathered review data by compiling a number of review websites that posted either numerical (e.g. 5 star scale) or letter grade ratings, and checking those websites for new album reviews every Sunday. We converted the rating systems of the different websites to a 100 point scale. If albums were not rated, then they received a 0 rating. While doing so does decrease the mean rating significantly, a 0 rating takes into account the fact that the album is not receiving any attention. Most of the albums did not receive consumer reviews until after the album was released, which explains the relatively low mean rating in Figure 1.

To examine the relative influence of different types of reviews, we sorted review sources into three different categories: consumer reviews, online media reviews, and mainstream media reviews.

²⁰ “The Net is a Boon for Indie Labels.” *The New York Times*. December 27, 2005.
<http://www.nytimes.com/2005/12/27/arts/music/27musi.html?ex=1293339600&en=7d783101017430f0&ei=5088&partner=rssnyt&emc=rss>.

Figure 1: Descriptive Statistics for Review Data

Variable	N	Mean	SE Mean	StDev	Minimum	Q1	Median	Q3	Maximum
No of consumer reviews	864	4.85	0.575	16.911	0	0	0	3	208
No of online media reviews	864	0.4063	0.0281	0.8257	0	0	0	0	3
No of mainstream media reviews	864	0.8009	0.0292	0.857	0	0	1	1	3
No of total reviews	864	6.057	0.598	17.566	0	0	1	5	211
Average consumer rating	864	35.44	1.49	43.91	0	0	0	90	100
Average online media rating	864	15.45	0.957	28.136	0	0	0	0	84
Average mainstream media rating	864	40.91	1.25	36.72	0	0	60	73.33	91.67
Average total rating	864	50.38	1.32	38.93	0	0	70	83.19	100

Consumers. Consumer reviews represent the truest form of word of mouth. The reviewers are usually individuals who have already purchased or listened to the album in question, and are not being compensated for writing the review. In our study, consumer reviews were represented by Amazon.com’s “customer reviews” feature, which allows individual users to submit reviews on a product and rate it on a 5 star scale. We recorded the total number of customer reviews and the average number of stars given for each album and converted the star rating to a 100 point scale by multiplying the Amazon rating by 20.

Online media. We categorized “online media reviews” as those posted by media websites that only have a presence on the Internet and do not publish “dead tree” magazines or other materials. We chose websites that are, at least anecdotally, considered influential on the music scene: Pitchfork Media (www.pitchforkmedia.com), PopMatters (www.popmatters.com), and Stylus Magazine (www.stylusmagazine.com). Pitchfork and PopMatters both rate albums on a 1-10 scale, so we multiplied their ratings by 10. Stylus Magazine rates albums on an A+ to F scale that includes + and – gradations for all letter grades besides F. To obtain a numerical rating, we set A+ to equal 100 and F equal to 0, leaving equal divisions between each possible rating. We averaged the three ratings with equal weights to obtain the “average online media rating.”

Mainstream media. “Mainstream media reviews” differ from the other two in that they are reviews that are posted online by major print media publications. We chose Rolling Stone (www.rollingstone.com), Entertainment Weekly (www.ew.com), and Allmusic (www.allmusic.com). Allmusic is known primarily to consumers as an online music information and review database, but it also publishes All Music Guide reference books and licenses its data and reviews to Billboard so that an artist search on Billboard’s website will return an Allmusic review. For our study, this was considered a mainstream media site. Rolling Stone and Allmusic grade music albums on a 1-5 scale, and Entertainment Weekly uses an A+ to F scale like Stylus Magazine’s. The ratings for the three were averaged, with each source weighted equally.

3.3.4 Blog chatter

A vast array of blogs exists on the Internet, spanning a variety of subjects. Even within a subject, there are many different types of blogs. A music blog might focus on only one specific genre of music, or post only concert reviews, or offer MP3s for readers to download. Many music blogs might post CD reviews, but not all of them rate CDs on any sort of scale. This makes capturing blogs’ average rating of an album extremely difficult, but it seems plausible that any mention of an album on a blog still has the potential to generate consumer awareness. The more blog posts about an album that are online and accessible to a consumer, the more likely a consumer will have information about the album to make a purchase decision.

To measure “blog chatter,” we looked at blog post volume. We did searches on Technorati, a search engine for blogs, for each album in the sample and noted the total number of blog posts that came up. We limited the search to only blogs in English with some authority – in Technorati

terms, the more other blogs link to a blog post, the more authority of that post. This constraint was necessary to narrow the results down to legitimate blogs written by humans instead of spam sites that list all sorts of keywords in hopes of getting hits.

We also recorded the “change in chatter” from week to week. Our hypothesis was that a rapid increase in awareness may signal a critical amount of buzz surrounding an album, which would imply that sales may increase faster than normal.

3.3.5 Social Networking Intensity: Myspace

One measure of the popularity of an artist is the number of friends he or she has on a social network. Myspace, a social networking site recently bought by News Corporation, has a strong reputation for promoting music artists. Approximately 80% of the artists in the sample – which ranged John Mellencamp to American Idol runner-up Katherine McPhee - had an official Myspace page. We gathered Myspace friends data by visiting each artist’s Myspace page weekly and noting the number of Myspace users who had added them as friends, which is publicly displayed on each profile. We also calculated the percentage change in Myspace friends from week to week.

4. Analysis and results

To analyze our data, we began by regressing the independent variables on weekly unit sales of an album two weeks after the observation date. We first examined the correlations between the variables. Similar to the results in Dellarocas (2004), we found that average mainstream media ratings showed low correlation with consumer ratings (see Figure 2).²¹ Online media ratings

²¹ Dellarocas, Chrysanthos, Neveen Farag Awad, and Xiaoquan (Michael) Zhang (2004). “Exploring the Value of

showed an even lower correlation with mainstream media ratings. Most significantly, of all the variables, blog chatter showed the greatest correlation (0.495) with sales two weeks after the observation date.

Figure 2: Pearson correlation matrix for variables data

	Days since release	Number of consumer reviews	Number of online media reviews	Number of mainstream media reviews	Total number of reviews	Average consumer rating	Average online media rating	Average mainstream media rating	Average total rating	ln (sales 2 weeks after obs date)	ln (blog chatter)	ln (number of Myspace friends)	ln (% change in Myspace friends)
Number of consumer reviews	0.280 (0.000)												
Number of online media reviews	0.370 (0.000)	0.255 (0.000)											
Number of mainstream media reviews	0.513 (0.000)	0.459 (0.000)	0.548 (0.000)										
Total number of reviews	0.312 (0.000)	0.997 (0.000)	0.319 (0.000)	0.516 (0.000)									
Average consumer rating	0.682 (0.000)	0.318 (0.000)	0.349 (0.000)	0.488 (0.000)	0.347 (0.000)								
Average online media rating	0.389 (0.000)	0.188 (0.000)	0.839 (0.000)	0.449 (0.000)	0.242 (0.000)	0.314 (0.000)							
Average mainstream media rating	0.496 (0.000)	0.187 (0.000)	0.371 (0.000)	0.795 (0.000)	0.237 (0.000)	0.446 (0.000)	0.339 (0.000)						
Average total rating	0.621 (0.000)	0.234 (0.000)	0.321 (0.000)	0.664 (0.000)	0.273 (0.000)	0.700 (0.000)	0.370 (0.000)	0.814 (0.000)					
ln (sales 2 weeks after obs date)	0.067 (0.050)	0.328 (0.000)	0.196 (0.000)	0.343 (0.000)	0.341 (0.000)	0.234 (0.000)	0.110 (0.001)	0.184 (0.000)	0.177 (0.000)				
ln (blog chatter)	0.252 (0.000)	0.332 (0.000)	0.520 (0.000)	0.511 (0.000)	0.369 (0.000)	0.319 (0.000)	0.406 (0.000)	0.371 (0.000)	0.321 (0.000)	0.495 (0.000)			
ln (number of Myspace friends)	0.062 (0.112)	0.189 (0.000)	0.145 (0.000)	0.210 (0.000)	0.201 (0.000)	0.123 (0.000)	-0.008 (0.830)	0.169 (0.000)	0.094 (0.017)	0.359 (0.000)	0.544 (0.000)		
ln (% change in Myspace friends)	0.000 (0.995)	0.143 (0.001)	-0.009 (0.834)	0.145 (0.001)	0.145 (0.001)	0.077 (0.074)	-0.012 (0.784)	0.094 (0.029)	0.097 (0.024)	0.250 (0.000)	0.127 (0.003)	-0.166 (0.000)	
Major/ Independent label	0.001 (0.980)	0.220 (0.000)	-0.013 (0.704)	0.122 (0.000)	0.217 (0.000)	0.063 (0.063)	-0.100 (0.003)	0.058 (0.086)	0.012 (0.732)	0.458 (0.000)	0.210 (0.000)	0.291 (0.000)	0.265 (0.000)

Pearson correlation
(P value)

4.1 Linear regression

In order to normalize the distributions of *BlogChatter* and *%ChgMyspace*, we took the logs of the variables. We then ran an ordinary least squares regression and estimated a model of the resulting form:

$$\begin{aligned} \ln(\text{Sales2WeeksAfter}) = & \alpha + \beta_1(\text{DaysSinceRelease}) + \beta_2(\text{MajorIndieLabel}) + \\ & \beta_3(\text{NoMainstreamReviews}) + \beta_4(\text{AvgMainstreamRating}) + \\ & \beta_5(\text{AvgConsumerRating}) + \beta_6 \ln(\text{BlogChatter}) + \\ & \beta_7 \ln(\text{\%ChgMyspace}) \end{aligned}$$

As stated in Section 3.2, the parameters in the Pareto relationship we used to convert sales rank into sales quantity are not calibrated to the CD category at Amazon. In order to ensure that the regression results were not distorted by this estimation, we ran a sensitivity analysis. We found that the overall R^2 of the linear regression equation remains unchanged no matter how the β parameter is adjusted (see Figure 3). Changes in the β parameter simply change the regression coefficients monotonically. Since the R^2 of the regression equation is not sensitive to the Pareto relationship parameters, we used the previously stated α parameter of 10.526 and β parameter of -1.61 to obtain the results shown in Figure 4.

Figure 3: A sensitivity analysis of the β parameter

B parameter	-1.61	-1.8	-2	-2.2	-2.4
R sq	0.56	0.56	0.56	0.56	0.56
constant	-5.56	-7.45	-9.45	-11.5	-13.4
days since release	-0.0889	-0.0994	-0.11	-0.121	-0.133
no of mainstream reviews	0.976	1.09	1.21	1.33	1.46
avg consumer rating	0.0168	0.0188	0.0209	0.023	0.0251
avg mainstream rating	-0.0207	-0.0231	-0.0257	-0.0283	-0.0308
ln blog chatter	1.39	1.55	1.73	1.9	2.07
ln % change in friends	0.329	0.368	0.409	0.449	0.49
major/indie label	2.62	2.93	3.26	3.58	3.91

The findings strongly suggest that higher blog post volumes and higher percentage changes in Myspace friends correspond to increased weekly sales in the future. This is interesting because both these variables tend to have usable information *before* an album is released, while the majority of reviews in all review source categories only begin to appear within the first week surrounding the album release date.

Figure 4: Results for the linear regression

Variable	Coefficient	T Statistics
Days since release	-0.0904	-7.62
Number of mainstream reviews	1.13	4.29
Average mainstream review rating	-0.0241	-4.44
Average consumer review rating	0.0179	4.69
ln (Blog chatter)	1.27	13.92
ln (% change in Myspace friends)	0.363	2.78
Major or independent label	2.54	8.99

* all significant at the 1% level

Interestingly, the coefficient for *BlogChatter* is larger than that for *%ChgMyspace*. Therefore, an additional 1% increase in blog post corresponds to a greater increase in sales than an additional 1% increase in Myspace friends. Figures 5 and 6 are based on a median album with median values for each independent variable. They indicate that higher percentage changes in Myspace friends produce diminishing returns, which is not true for blog chatter where the relationship is more linear, or even slightly concave.

There are a few possible explanations for the differences in figures 5 and 6 and for the relative importance of blogs and social networks. We posit that Myspace is considered a less credible source of user-generated content because its content involves less commitment from reviewers. Adding a friend on Myspace is a fairly simple process – a user need only click “add to friends” on the artist’s profile page and confirm it for the friend to be added. Therefore, it is possible that an artist’s Myspace friends count is bloated with individuals who only passed by the

artist’s profile once, and are not very interested in the artist’s new album or in spreading word about the album. In addition, being a Myspace friend is a very passive process; no information, aside from the increased friend count, is produced for other potential consumers to view. On the other hand, blog writers have to spend time and make an effort in writing a blog post, and are unlikely to do this without adequate thought and commitment, and even future credibility in mind.

Figure 5: Blog chatter vs. weekly sales quantities two weeks after the observation date

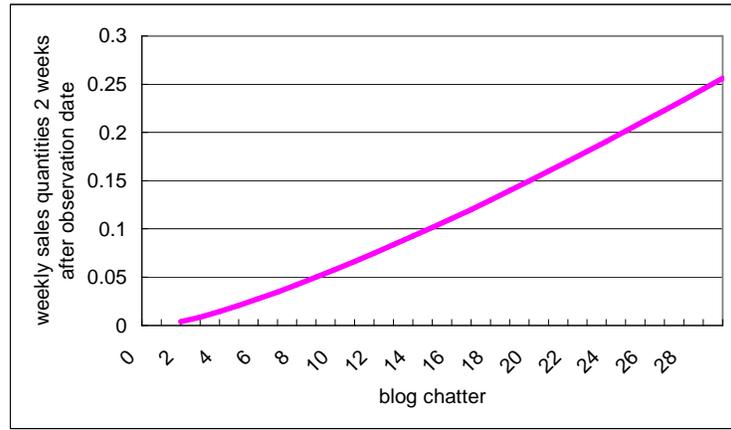
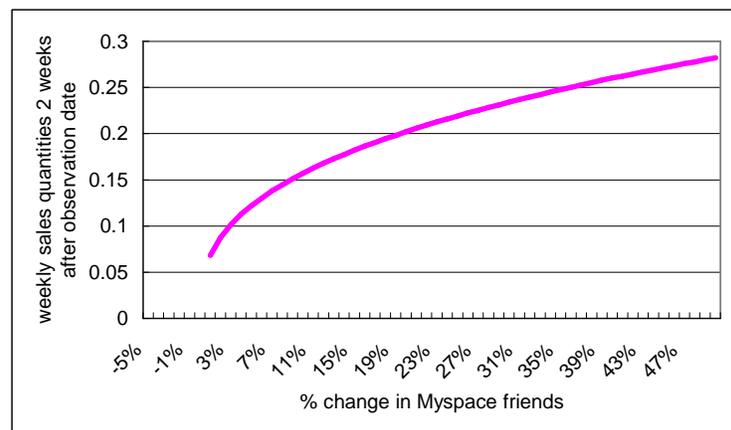


Figure 6: Percentage change in Myspace friends vs. weekly sales quantities two weeks after observation date



The coefficients of the other independent variables are mostly as expected. The negative coefficient for *DaysSinceRelease* confirms our hypothesis that sales of an album fall over time

after the release date. The positive coefficients for both *AvgConsumerRating* and *NoMainstreamReviews* suggest that album sales rise with higher consumer ratings and more mainstream reviews. Based on the coefficient for *MajorIndieLabel*, we estimate that a release by a major label has approximately twelve times the sale of a release by an independent label. Other sources have estimated that major label releases sell about six times more than independent label releases.²² The discrepancy could be due to the small size of the sample we used – 108 albums in comparison to the approximately 35,000 albums released in a single year.²³ Nevertheless, the general trend is consistent with prior research.

The negative coefficient for *AvgMainstreamRating*, on the other hand, is surprising, since a higher critical rating would be expected to correspond to higher sales. A closer examination of the data reveals a number of possible explanations for such a coefficient. One is simply that critical acclaim may not correspond with popular appeal, and this seems to be supported by the relatively low correlation between the average mainstream media rating and average consumer rating shown in Figure 2. In addition, when we examined the five albums with the lowest mainstream media rating, all of them had high sales ranks (see Figure 7), indicating high popularity. It is difficult to say if this is representative of a trend. However, we speculate that there may be some sort of selection process occurring. Albums that will not be well-received by critics are released by labels if the label believes the album will sell anyway; albums that are not surefire hits based solely on the artist's notoriety or other factors will only be released if they meet a certain minimum quality standard. Whatever the reason, the results seems to agree with the finding in Dellarocas et al. (2004) that average consumer ratings are better predictors than professional critics' ratings.

²² Hull, Geoffrey. *The Recording Industry*. Routledge: New York, 2004, 136.

²³ Hull, 136.

Figure 7: The five albums with the lowest mainstream media rating

Artist	Album	Average Mainstream Rating	Sales Rank 2 Weeks After
Celtic Woman	A New Journey	41.67	10
John Mellencamp	Freedom's Road	50.00	16
Kidz Bop Kids	Kidz Bop 11	50.00	136
Katharine McPhee	Katharine McPhee	53.33	48
Art Garfunkel	Some Enchanted Evening	54.17	149

4.2 Prediction Versus Causality

As we noted earlier, one might question whether chatter is truly predictive of subsequent sales, or whether increased sales lead to increased chatter which in turn leads to increased sales! In order to test this, we partitioned the dataset as described earlier such that only pre-release chatter was considered and paired with post-release sales. The results are as shown below for the 216 such observations that we described at the outset of this section.

Figure 10: Results for pre-release chatter only

Variable	Coefficient	T Statistics
Days since release	-0.1011	-1.29
Number of mainstream reviews	1.2432	1.66
Average consumer review rating	-0.0322	-2.48
Average mainstream review rating	0.0252	1.77
ln (blog chatter)	1.4272	7.75
ln (% change in friends)	0.2319	0.84
Major or independent label	2.4230	4.07

The R-square for this model was 0.5. Interestingly, the coefficients are similar, but only blog chatter and the major or independent label are significant at the 1% level. Again, blog chatter is the most significant variable.

Based on both results, it seems clear that user-generated content as measured by blog chatter matters in subsequent sales for music. Interestingly, the increase in size of the social network is not significant in the reduced dataset, suggesting that it may have no predictive value

before release or that it may only matter after release.

Finally, it is natural to ask whether it is reasonable to conclude that increased blog chatter really *causes* an increase in sales since it does, after all, precede sales). It is not possible to make such a conclusion based on this study. It may be the case, for example, that it is another, unobserved variable that causes *both* increased chatter *and* sales with the increased chatter occurring first. One might posit that it is in fact “the quality of the artist” that “causes” both increased blog chatter and sales, where high quality is somehow recognized in the marketplace by some mechanism, which in turn has its effect on what we are observing. Without a strong prior model that includes such a variable, it is not possible to draw any causal connection. This is important not just theoretically but also practically because it means that it may be futile to engineer an increase in blog posts with the expectation that this will lead to higher sales!

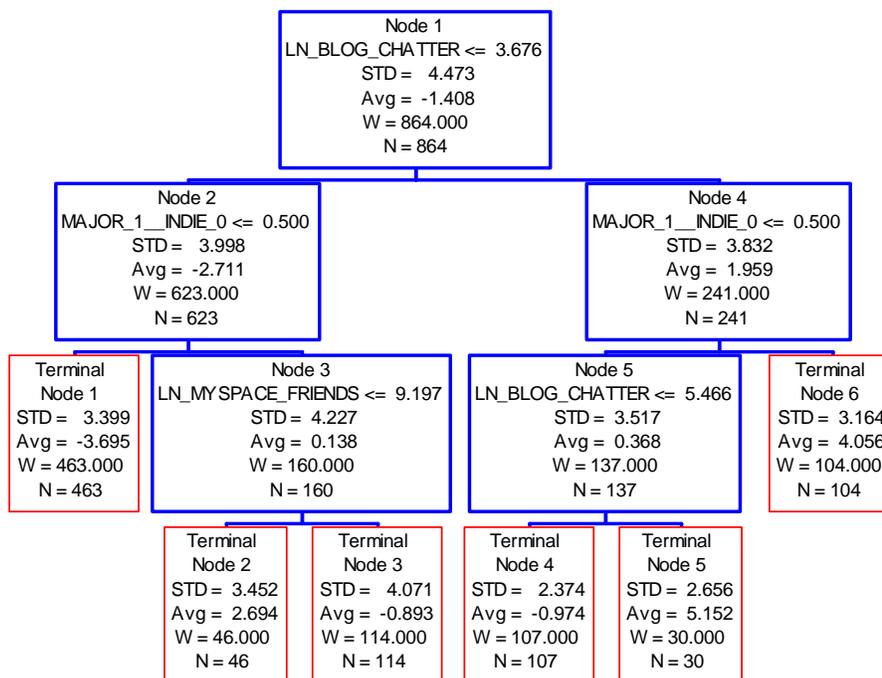
4.3 Nonlinear Regression (Tree)

An interesting question here is whether there are any significant interaction effects among the various metrics. We posit that these would be of particular interest to marketing managers interested in sifting through the burgeoning volume of Web 2.0 metrics becoming available on the Internet. Which ones, when considered simultaneously, provide insights not derivable by looking at them in isolation? It is difficult to have well formulated hypotheses about this at the current time considering the recent nature of such metrics. However, it is worthwhile working bottom-up using inductive pattern discovery methods to find the interesting interactions that can be tested further in future studies.

We ran a tree induction algorithm on the data to uncover the significant interaction effects. Interestingly, the tree induction algorithm identifies blog chatter as the most important variable by partitioning the data first on that variable (see Figure 8). If an album has more than 40 blog posts

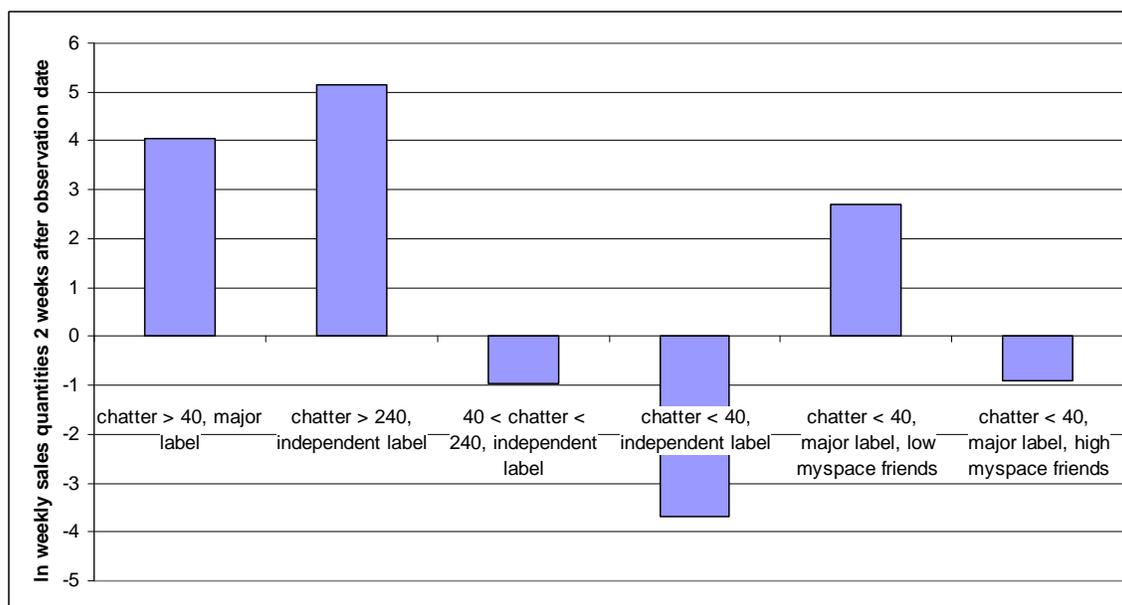
($\ln(40)$ is roughly 3.7 which is the threshold for the first split in Figure 8), it will have an above average level of sales (almost one standard deviation above average). If an album has more than 40 blog posts *and* is released by a major label, then it is likely to have very high sales (almost one and a half standard deviations above average as indicated by node 6 in Figure 8). This is no surprise, as a large number of blog posts indicate a high level of buzz, and being released by a major label means it is more likely that there will be significant promotion of the album through channels other than the Internet. Interestingly, though, if blog chatter is extremely high – above 240 posts – it is possible for an album to overcome the disadvantage of being released by an independent label. In fact, albums with such extreme highs in chatter correspond to sales even higher than major label, high chatter albums (Figures 8 and 9). However, even if chatter is relatively high for an independent label (above 40 posts), sales will be higher than the average for the sample, but still relatively low if the 240 post level is not breached. An independent label with low blog chatter has very low sales (node 1), which is as expected.

Figure 8: Regression tree analysis



Finally, the tree indicates that major label releases with low blog chatter (less than 40 blogs) and low numbers of Myspace friends have higher sales than major label releases with low blog chatter and high numbers of Myspace friends. This seems counter-intuitive at first, but in the sample, major label releases without a Myspace page were considered to have 0 Myspace friends, which could explain the result. In addition, major label releases that had a Myspace page but few Myspace friends were from artists such as John Mellencamp and Art Garfunkel; we would presume that the majority of their audiences, who are “older,” do not generally use Myspace.

Figure 9: Regression tree nodes



5. Conclusion

The Internet provides consumers with a powerful word of mouth channel for information on upcoming music releases. Previous research has found evidence that online consumer reviews can predict book and movie sales (Dellarocas et al. 2004, Duan et al. 2005, and Chevalier 2006), but no study prior to ours that we know has explored the effects user-generated content including blogs and social networking sites as well as traditional sources of reviews for predicting online music sales. We analyzed the usefulness of blogs and social networks, as well as reviews in consumer, online media, and mainstream media, in predicting album sales in the four weeks before and after the album's release date. We found that the most significant variable is blog chatter or the volume of blog posts on an album, with higher numbers of posts corresponding to higher sales. Higher percentage changes in Myspace friends may also be significant, although the results here were not consistent. Contrary to the conclusions in Duan (2005), we find that the average consumer rating is significant, while the number of consumer reviews is not. Our results also

confirm the findings in Dellarocas (2004), with average consumer ratings better predicting sales than average mainstream media ratings.

Although we found that user-generated content is a good predictor of music album sales, our analysis showed that traditional factors cannot be ignored. While independent label releases with extremely high blog chatter can sell even more units than major label releases, our findings estimated that the average major label release sold approximately twelve times more than the average independent label release. We also found that the higher the number of mainstream media reviews, the greater the sales.

The results of this study suggest that user-generated content should be considered seriously by record labels. Most notably, since blog chatter and Myspace friend information is available *before* an album releases and ships, record labels can examine these two variables to predict future sales well in advance of when the album is available in stores.

At the same time, we caution against assumptions of causality for reasons discussed in the last section. It is conceivable, for example, that if blog posts start becoming manipulated *because* people think they have an impact on sales, that the predictive power might disappear because the underlying reasons for it disappear. There is a crude analogy here to efficiency in financial markets where predictive models lose their power over time as the relationships become recognized and exploited by people who seek to benefit from their existence.

6. Future research

This research has generated a number of interesting topics worth further research. One would be a closer examination of the relative effects of user-generated content by genre of music. Presumably, genres of music more popular with the younger generations – those most likely to use

the Internet heavily – would be more influenced by online user-generated content than those popular with older generations, such as classical music. Alternatively, one might find that genres such as Scandinavian death metal that are more difficult to find information on through mainstream sources may be more strongly correlated with user-generated reviews.

Another direction would be to examine not only blog post volume on an album, but also blog sentiment – how positive or negative are the opinions expressed in the blog post. Various studies, including Das and Chen (2006) and Ghose et al. (2006), have experimented with different ways to quantify sentiment expressed in text.^{24 25}

Extracting sentiment and converting it to a numerical score would help describe the quality of reviews an album is getting. In addition, in our experience, most blog posts tend to be positive overall – rarely will a blogger feel the need to spend time and effort to write a scathing review. If this hypothesis holds true, then if blog sentiment is generally negative on an album and blog post volume is high, it could signal dim future prospects.

A final avenue of study could be on the relative effect of user-generated content on digital sales versus physical album sales if data for digital download sales could be obtained from iTunes or other similar digital download stores. Since digital downloads enable instant gratification and – by virtue of their lower price – are more likely to be purchased impulsively, user-generated content may have a more immediate effect on digital download sales. It seems plausible that a consumer could read a highly favorable blog post or review online, and then decide to pay the 99 cents or so to hear the song immediately. This is a fruitful area of future research.

²⁴ Das, Sanjiv Ranjan, and Mike Chen (2006). “Yahoo! For Amazon: Sentiment Extraction from Small Talk on the Web.” Working paper, Santa Clara University. January 2006.

²⁵ Ghose, Anindya, Panagiotis Ipeirotis, and Arun Sundararajan (2006). “The Dimensions of Reputation in Electronic Markets.” Working paper, New York University. February 2006.

Bibliography

- Antweiler, W., and M.Z. Frank. Is all that talk just noise? The Information Content of Internet Stock Message Boards. *Journal of Finance*, 59(3): 1259-1295, 2004.
- “Apples iTunes Slams Into Top Ten Music Retail Chart.” *The Register*. November 22, 2005. http://www.theregister.co.uk/2005/11/22/itunes_us_retail_top_ten/.
- Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith (2003). “Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers.” *Management Science*, Vol 49, No. 11, 1580-96.
- “Can Music Survive Inside the Big Box?” *Wall Street Journal*. April 27, 2007. B1.
- Chevalier, Judith and Austan Goolsbee (2003). “Measuring prices and price competition online: Amazon and Barnes and Noble.” *Quantitative Marketing and Economics*, Vol. 1, No. 2, June 2003, 203-222.
- Chevalier, Judith and Dina Mayzlin (2006). “The Effect of Word of Mouth on Sales: Online Book Reviews.” *Journal of Marketing Research*, Vol. 43, No. 3, 345-354.
- “Clap Your Hands Say Yeah: ‘This Home on Ice’ [track review].” *Pitchfork Media*. June 14, 2005. http://www.pitchforkmedia.com/article/track_review/30565-in-this-home-on-ice.
- “Clap Your Hands Say Yeah Sidesteps Label Deals.” *Billboard*. September 20, 2005. http://www.billboard.com/bbcom/news/article_display.jsp?vnu_content_id=1001140150.
- Das, Sanjiv Ranjan, and Mike Chen (2006). “Yahoo! For Amazon: Sentiment Extraction from Small Talk on the Web.” Working paper, Santa Clara University. January 2006.
- Dellarocas, Chrysanthos, Neveen Farag Awad, and Xiaoquan (Michael) Zhang (2004). “Exploring the Value of Online Product Ratings in Revenue Forecasting: The Case of Motion Pictures.” Working paper, University of Maryland. May 2004.
- Duan, Wenjing, Bin Gu and Andrew B. Whinston (2005). “Do Online Reviews Matter? – An Empirical Investigation of Panel Data.” Working Paper, the University of Texas at Austin. January 2005.
- Eliashberg, Jehoshua and Steven Shugan (1997). “Film Critics: Influencers or Predictors?” *Journal of Marketing*, Vol. 61, No. 2, 68-78.
- “To Generate Buzz, Clap Your Hands on the ‘Net!’” *NPR*. Nov. 29, 2005. <http://www.npr.org/templates/story/story.php?storyId=5023133>.
- Ghose, Anindya, Michael D. Smith, Rahul Telang (2006). “Internet Exchanges for Used Books: An Empirical Analysis of Product Cannibalization and Welfare Impact.” *Information Systems Research*, Vol 17, No. 1, 3-19.
- Ghose, Anindya, Panagiotis Ipeirotis, and Arun Sundararajan (2006). “The Dimensions of Reputation in Electronic Markets.” Working paper, New York University. February 2006.
- Godes, David and Dina Mayzlin (2004). “Using Online Conversations to Study Word-of-Mouth Communication.” *Marketing Science*, Vol 23, No. 4, 545-560.
- Gruhl, D., L. Chavet, D. Gibson, J. Meyer P. Pattanayak, A. Tomkins, and J. Zien. How to Build a Webfountain: An Architecture for Very-Large Scale Text Analytics. *IBM Systems Journal*, 43(1):64-77, 2004.
- Gruhl, D, R. Guha, Ravi Kumar, Jasmine Novak, and Andrew Tomkins. The Predictive Power of Online Chatter. *KDD 2005*, Chicago, IL, August 2005.
- Hull, Geoffrey. *The Recording Industry*. Routledge: New York, 2004.
- “JupiterResearch Finds That MySpace Music Community Activity Far Outpaces Online Music Sites.” JupiterResearch. May 31, 2006.

http://www.jupiterresearch.com/bin/item.pl/press:press_release/2006/id=06.05.31-myspace_music_community.html.

Lyman, Peter and Hal R. Varian, "How Much Information", 2000.

<http://www.sims.berkeley.edu/how-much-info>,

Moul, Charles (2006). "Measuring Word of Mouth's Impact on Theatrical Movie Admissions." *Journal of Economics and Management Strategy*, forthcoming.

"The Net is a Boon for Indie Labels." *The New York Times*. December 27, 2005.

<http://www.nytimes.com/2005/12/27/arts/music/27musi.html?ex=1293339600&en=7d783101017430f0&ei=5088&partner=rssnyt&emc=rss>.

"Online Survey Finds General Public, Media Conference Attendees Agree that Traditional News Outlets Could Do a Better Job." *PR Newswire*. February 17, 2007.

"Where Harry Most Enchants." *The Washington Post*. April 17, 2007. B01.

"You Will Pay for Your Excessive Charm." *Said the Gramophone*. June 9, 2005.

http://www.saidthegramophone.com/archives/you_will_pay_for_you.php.

Tumarkin, R., and R.F. Whitelaw. News or Noise? Internet Postings and Stock Prices. *Financial Analysts Journal*, pages 41-51, 2001.