Is Tom Cruise Threatened?

Using Netflix Prize Data to Examine the Long Tail of Electronic Commerce

Tom F. Tan

Serguei Netessine

Operations and Information Management, Wharton Business School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

fangyun@wharton.upenn.edu, netessine@wharton.upenn.edu

Abstract

We analyze a large data set from Netflix, the leading online movie rental company, to shed new light on the causes and consequences of the Long Tail effect, which suggests that on the Internet, over time, consumers will increasingly shift away from hit products and toward niche products. We examine the aggregate level demand as well as demand at the individual consumer level and we find that the consumption of both the hit and the niche movies decreased over time when the popularity of the movies is ranked in absolute terms (e.g., the top/bottom 10 titles). However, we also observe that the active product variety has increased dramatically over the study period. To separate out the demand diversification effect from the shift in consumer preferences, we propose to measure the popularity of movies in relative terms by dynamically adjusting for the current product variety (e.g., the top/bottom 1% of titles). Using this alternative definition of popularity, we find that the demand for the hits rises, while the demand for the niches still falls. We conclude that new movie titles appear much faster than consumers discover them. Finally, we find no evidence that niche titles satisfy consumer tastes better than hit titles and that a small number of heavy users are more likely to venture into niches than light users.

Keywords: the Long Tail effect; movie rental; product variety; product rating; purchase frequency; Internet, e-commerce.

1 Introduction

Chris Anderson, editor-in-chief of *Wired Magazine*, coined the term "Long Tail effect" (Anderson, 2004) suggesting that, due to the introduction of the Internet, niche products will comprise higher and higher market share, while the demand for hit products will continue to decrease. As a result, he predicted that the old Pareto rule, stating that 20% of all the products generate 80% of the revenues, will no longer hold: hit movies will constitute a smaller and smaller proportion of demand. His predictions of the Long Tail effect were motivated by observations in the media, entertainment and other industries. For example, Anderson (2006) finds that the top 50 best-selling albums of all time were produced in the 70s and 80s; none of them were recorded in recent years. He also observes that the ratings of the top TV shows have gradually decreased and that the top show today would not have ranked among the top ten in 1970. Part of the reason, according to Anderson, is that niche products will better and better satisfy consumer preferences because consumers will continue to have more and more varying preferences while the Internet will make even the most obscure products available to the masses.

The potential for the existence of the Long Tail effect is of great importance for product assortment decisions in a variety of industries, for advertising dollars spent on supporting this variety, and for supply chain management of these products on the Internet. For example, Blockbuster stocks 3,000 DVDs per store on average, while 20% of Netflix rental revenues come from outside the top 3,000 titles (Anderson, 2004). In addition, Ecast, a digital jukebox company, sold 98% of its 10,000 albums available online at least one track per album per quarter (Anderson, 2006), while brick-and-mortar music stores only stock a fraction of this variety. If demand is indeed shifting toward more obscure titles, managers should ensure that these titles are available and that they are advertised properly. Further, Anderson explains that the new online recommendation systems help the niche products quickly find their demand in the market once they are made available. As a result, he asserts that "the tail of available variety is far longer than we expected", and that the combined market share of the niches can outgrow the hits (Anderson, 2006). This comment about the increasing demand for the niches seems to be consistent with Varian's opinion in light of the cheaper technology in the media industry. Specifically, Varian (2006) notes that this "creative, inexpensive and compelling semiprofessional content available via the Internet" has an increased demand particularly among young people, so that the salaries of celebrities, such as Tom Cruise, may decrease.

Although arguments and evidence in favor of the Long Tail effect appeared pervasive at first, there are also indications that hits still drive some markets, and may even become more popular over time, whereas the rising demand for niches is, at best, overestimated. In particular, some evidence suggests that new products appear so quickly that consumers have no time to discover them. Gomes (2006) discloses that at Ecast, the quarterly no-play rate increased from 2% to 12% as product variety has grown. Ignoring this increasing number of products with no demand is known to cause a biased estimation of the sales distribution (Schmittlein et al., 1993). In addition, an even stronger demand for hits is found in the motion picture industry, where both the number of movies that generate box-office revenues of over $50 million and their percentage of the total revenues increased from 14 and 14% in 1998 to 19 and 22% in 2003, respectively (Eliashberg et al., 2006). Finally, Orlowski (2008) reports on an industry study which discovered that 80% of the digital song inventory sold no copies at all - and the 'head' of the frequency distribution was far more concentrated than expected. Given this conflicting evidence, whether or not the Long Tail effect exists remains a hotly debated issue among practitioners.

The Long Tail effect has also recently generated widespread interest in academic circles (more on these and related papers later). Brynjolfsson et al. (2006) present plausible factors that may drive the Long Tail effect, including both supply-side and the demand-side effects. On the supply side, they suggest that the Internet reduces the production and distribution costs of niche products. On the demand side, they note that both the active and the passive search tools of the Internet lower the search costs and hence facilitate finding niche products. Moreover, Tucker and Zhang (2009) suggest that product popularity information, such as the number of people who have browsed the product, can increase the appeal of niche products disproportionately, thus causing the Long Tail effect. On the other hand, Fleder and Hosanagar (2008) suggest that sales diversity can be reduced by selection-biased recommendation systems because these systems tend to recommend products with sufficient historical data (i.e., hits), while Park and Tuzhilin (2008) propose an algorithm that can promote recommendations for the tail items. Bockstedt and Goh (2008) analyze the data of consumer-created custom CDs to examine whether people tend to bundle the hits or the "long tail" music and suggest that managers should sell unbundled information goods to meet the demand from the mainstream consumer. Elberse and Oberholzer-Gee (2008) find further evidence that online retailing triggers demand to shift toward the tail of the distribution, although they also find that a substantive part of demand is concentrated on an even smaller portion of products.

So far, both academic theories and the empirical evidence provide what can probably be described as conflicting evidence for the existence and the magnitude of the Long Tail effect: while there are many anecdotal examples of its presence, there are fewer than a handful of rigorous studies. At the same time, whether or not the Long Tail exists is a fundamental question for decision-makers in marketing, operations, and finance who face the prospect of further penetration of the Internet channel, which offers expanding product variety and new recommendation systems to help manage it.

In this paper we provide empirical evidence to shed additional light on the existence of the Long Tail effect from a different perspective. We use a novel longitudinal data from Netflix that contains 100 million online ratings of 17,770 movie titles by 480,000 users from 2000 to 2005. Netflix is the key example in Anderson's evidence for the Long Tail effect and he primarily refers to the popularity of products in absolute terms, e.g., the top 10 or the top 100 for hits, and the bottom 10 or the bottom 100 for niches. In his own words, "number one is still number one, but the sales that go with that are not what they once were" (Anderson, 2006). Following this example, we first study the number of ratings for movie titles over time and find that, when movie popularity is measured in absolute terms, there is only partial evidence to support the Long Tail effect: demand for hits decreases over time but demand for niches decreases too.

The above definition of the Long Tail effect and movie popularity is static, which implicitly excludes the impact of an increasing product variety. This definition would certainly reflect product popularity in a channel where product variety is relatively stable and where all products are consumed, such as in a brick-and-mortar store. However, product variety has skyrocketed during the Internet age, and more products than ever are not being discovered by consumers. For example, in our data the number of rated movies increased by a factor of four over five years while the number of unrated movies exceeds the number of rated movies by a factor of two in 2005. Such a dramatic increase in product variety is likely to create demand diversification. For example, given a choice set of only five movies, people may tend to concentrate their demand on one movie whose popularity rank is number 1 or equivalently in the top 20%. However, out of a wider choice set of 500 movies, the demand may be concentrated on 100 movies whose popularity ranks in the top 100 or also in the top 20%. This example causes a conflicting definition of hits and niches amid different sizes of product variety at different points in time. Should we classify the top 20%, which is respectively the top one out of five movies and the top 100 out of 500 movies as the hits, or should we restrict the label of hits to only the top one movie no matter the total variety?

Naturally, when the product variety is large, the demand for any one product tends to be smaller than when the product variety is small. Likewise, when the consumer base is large, learning about new products is faster than when the consumer base is small. In this case, two competing effects might be observed: 1) consumers discover the obscure products as they appear and 2) new products appear, possibly so quickly that most consumers have no time to discover them. Which effect dominates is an empirical question that we aim to address in this paper. Therefore, we argue, the definitions of hits and niches should vary with time as both product variety and consumer base vary. In this paper, we propose a dynamic definition of product popularity which adjusts for active product variety over time (which excludes titles that have no current ratings). The active product variety reflects the dynamics of both product variety and consumer base. We find that, if we define the popularity of a movie in relative terms, the Long Tail effect is absent – in fact, the demand for hits increases, whereas the demand for niches decreases. Specifically, we find that demand for the top 0.1% of movies increases five times as fast as demand for the top 10%, indicating that demand for the "hits of the hits" continues to skyrocket. The same finding is manifested by changes in the Pareto principle over time: while Anderson argues that the 80/20 rule will weaken (the top 20% of products will constitute less than 80% of demand), we find that the opposite is true: the share of demand for the top 20% of movies increases over time from 86% in 2000 to 90% in 2005. Furthermore, Anderson (2004, 2006) has argued that more and more consumers will choose niche products because they will tend to satisfy consumer preferences better. We, however, find that, contrary to Anderson's suggestion and independent of how popularity is measured, consumers tend to be less satisfied with niche movies than with hit movies and moreover, it is mostly heavy movie watchers, who constitute a small fraction of all consumers, that venture into niche movies.

To summarize, the contributions of this paper are three-fold. First, we provide new empirical results suggesting that there is only partial evidence for the existence of the Long Tail effect when it is measured in an absolute sense. Second, we propose to delineate two effects: demand diversification due to expanding product variety on the Internet and consumers learning about new products. We suggest that, when measuring product popularity, one has to adjust for instantaneous active product variety. With this definition, we find no evidence of the Long Tail effect. Third, we study demand at the consumer level and find that new movies appear so quickly that most consumers have no time to discover them, and that niche movies do not satisfy consumer tastes better than hit movies.

The remainder of this paper is organized as follows. We review the related literature in section 2 and develop our hypotheses in section 3. In section 4, we describe our data set and methods of research. We present the results of our empirical analysis in section 5. We conclude with a discussion of our results, limitations, and future research opportunities in section 6.

2 Related Literature

Traditional theories argue that the presence of popular items (i.e., hits) is quite persistent in the market. Rosen (1981) suggests that buyers tend to concentrate their demand on the "superstars" because of their imperfect substitutability and the joint consumption effects. The imperfect substitution causes small differences in the "talent" of the sellers to be "magnified in larger earnings differences". The joint consumption effect further ensures that demand concentrates on the few most talented sellers. Frank and Cook (1995) add that consumers tend to demand similar products such as movies or music in order to have a common language in their social interactions.

Although the traditional "superstar" theories have compelling arguments to support the hit-driven market, they are restricted to the pre-Internet era when the shelf space in brick-and-mortar stores was limited. Anderson's (2004) explanation about the Long Tail effect echoes Brynjolfsson et al.'s (2003) theory regarding the value of increased product variety on the Internet. This important and timely paper not only finds evidence of a significantly larger product variety at online retailers (such as Amazon) than at brick-and-mortar stores, but it also suggests that the Internet significantly lowers search costs so consumers can find more products. Consistent with this view, Cachon et al. (2008) find that the lowered search costs have a market-expansion effect, which encourages firms to enlarge their assortment. As a result, consumers are more likely to find niche products, thus causing demand for them to increase.

Although it is widely accepted that search costs are lowered as information technology proliferates, this does not necessarily reduce demand concentration. For example, Ghose and Gu (2006) suggest that search costs are even lower for popular products than for niches, which may limit the Long Tail effect. In addition, Hervas-Drane (2009) provides a model to argue that different search processes have mixed impacts on demand concentration.

Brynjolfsson et al. (2007) use a unique data set from a retailer operating both Internet and catalog channels and provide evidence of the Long Tail effect. Their paper is one of the few rigorous studies to directly compare demand concentration in both Internet and brick-and-mortar (cataloger) channels. Brynjolfsson et al. (2007) assume that the number of available products is the same in both channels. They further randomly select 100,000 items from the cataloger channel (selling 7,725,574 items in total) and another 100,000 items from the Internet channel (selling 702,659 items in total, about 10% of the total sales of the cataloger channel). The comparison of these sales data suggests that the Internet channel exhibits less concentrated sales, which is in line with the Long Tail theory. Smith et al. (2008) argue that the long tail titles, i.e., the niches, not only increase consumer surplus but also enhance producer surplus. Our study differs in that we do not possess data on sales in the Internet and brick-and-mortar channels; rather, we study what happens in the Internet channel over time.

In another study exploring the Long Tail effect, Elberse and Oberholzer-Gee (2008) study the Nielsen VideoScan data, by tracking weekly video sales from 2000 to 2005. In this study, the authors rank all the movies across the six years according to their weekly sales and define the top percentiles as the hits and the lower percentiles as the niches. In other words, they create static definitions of hits and niches for all the movies over a span of six years, which implicitly excludes the impact of an increasing product variety. We use different definitions of hits and niches that vary with time. Elberse (2008) further reports analysis of the data from Nielsen SoundScan, which tracks weekly sales of music; from Quickflix, an Australian movie rental service similar to Netflix; and from Rhapsody, an online subscription-based music jukebox. The different data supports the same conclusion that the "long tail" and "super stars" coexist.

Since we study ratings submitted by consumers as part of our analysis, there are at least two other relevant studies. Oestreicher-Singer and Sundararajan (2009) use data from 200 distinct categories on Amazon.com to establish that categories whose products are influenced more by recommendations have significantly flatter demand distribution, which supports the existence of the Long Tail effect. Dellarocas and Narayan (2007) find empirically that online consumers are more likely to review popular products, and therefore, contrary to the Long Tail effect, online reviews may exhibit "tall heads" instead of "long tails".

In summary, the extant literature has found conflicting evidence of the Long Tail effect. The main differentiating features of our analysis are (1) dynamic definitions of product popularity and (2) detailed user-level analysis, both of which are possible because we study a unique and extensive data set which includes product ratings by consumers over the course of six years.

3 Hypotheses and Ranking Methodology

Even before the Internet, increased product variety was known to better satisfy customers' heterogeneous needs, thus increasing overall demand for the firm selling these products (Baumol and Ide, 1956; Kekre and Srinivasan, 1990). Consumers are found to have a propensity for seeking variety over time. Farquhar and Rao (1976) and Pessemier (1979) suggest that people seek a variety of products to dynamically balance and maximize the utilities obtained from the attributes of the different products. In addition, McAlister (1982) proposes that people consume new products because they are satiated with the attributes of the old products.

Product variety expands dramatically on the Internet, which allows for offering many more products economically. Internet retailers such as Amazon.com are able to utilize novel fulfillment strategies such as drop-shipping to sell products to customers without actually stocking them, thus dramatically increasing product variety without significant extra costs (Randall et al., 2006). In their pivotal paper, Brynjolfsson et al. (2003) show that the wider product variety offered by Amazon.com has increased the consumer surplus seven to ten times more relative to gains from competition and lower pricing alone. Further, most online companies now provide recommendation systems to lower search costs, so that consumers are able to discover even the most obscure products in the Internet era (Brynjolfsson et al., 2003; 2006). And if product variety expands and consumers are increasingly able to gain access to new products, it should also be the case that demand for most popular products decreases while demand for less popular titles increases. As Varian (2006) explains, "It is true that there is only one Tom Cruise, but it is equally true that there are only

24 hours in a day. The more time young people spend watching Lonelygirl15, the less time they will have to watch Mr. Cruise." These considerations lead us to hypothesize, similar to Anderson (2006), that the Long Tail effect should be present if we focus on hits and niches defined in absolute terms.

HYPOTHESIS 1a: *If popularity is measured in absolute terms, over time, demand for hits will decrease, while demand for niches will increase.*

HYPOTHESIS 2a: *If popularity is measured in absolute terms, over time, individual consumers will demand fewer hit products and more niche products.*

While there is little doubt that the Internet allows companies to offer a wider variety of products economically, it is less clear that consumers necessarily quickly discover these products. Niche products tend not to have associated advertising budgets, and there is typically no sales representative on the Internet helping the customer find the product. Both Anderson (2006) and Brynjolfsson et al. (2003) suggest that the role of the sales person on the Internet is played by the electronic customer relationship management systems which are increasingly utilized by Internet companies (Padmanabhan and Tuzhilin, 2003). These systems usually utilize data-mining techniques to help consumers discover products. However, Fleder and Hosanagar (2008) find that, although recommendation systems can guide users to new products, they often tend to lead "similar users toward the same products", thus causing the aggregate diversity of products consumed to decrease over time. Indeed, most product recommendation systems are based on so-called collaborative filtering techniques which would not recommend a product with no historical demand. Thus, such systems would most often recommend hits rather than niches. Further, research has shown that consumers only consider a few important choices when they shop online. For instance, Gu et al. (2008) find that, at some online retailers, consumers who consider only two or fewer alternatives are responsible for over half of all purchases, which implies that consumers might simply not have time to venture into obscure products.

In view of the limitations imposed by current recommendation systems, we can expect to see two competing effects. First, the Internet channel allows retailers to economically offer larger and larger product variety over time. Second, driven by varying consumer preferences and with the help of recommendation systems, consumers discover and consume products heretofore unavailable. The question is: which effect dominates? To be able to delineate the two effects, it is necessary to introduce definitions of hits and niches that vary with time and that account for ever-changing product variety. We therefore propose a dynamic definition of product popularity that adjusts for increased product variety over time, such as the top 1% and the bottom 1% of total products demanded, to normalize the impact of product variety. This definition is not new: in fact, it is in the spirit of the Pareto principle which suggests that 20% of products account for 80% of demand. Given the current limitations of recommendation systems and the relative infancy of electronic customer relationship management systems, we hypothesize that product variety increases faster than the speed at which people discover products, so in relative terms the Long Term effect is absent.

HYPOTHESIS 1b: *If popularity is measured in relative terms, over time, the demand for hits will increase, while the demand for niches will decrease.*

HYPOTHESIS 2b: *If popularity is measured in relative terms, over time, individual consumers will demand more hit products and fewer niche products.*

4 Data

4.1 Research Setting, Data Collection, and Descriptive Statistics

To examine our research hypotheses, we utilize data available from Netflix, a major US online movie/TV series rental service with annual revenues in excess of $1 billion in 2008. Netflix is known for offering a wide selection of niche movies and it currently offers about 100,000 DVD titles to its 10 million subscribers. The data that we possess consist of the movie ratings submitted by consumers through the Netflix web site from 2000 to 2005. Netflix encourages its users to rate the movies that they have watched both outside and within Netflix to improve its recommendations for them, so users have direct incentives to provide truthful and complete ratings. In addition, Netflix constantly reminds users to rate the movies and it streamlines the rating process. As a result, Shih et al. (2007) suggest that Netflix has the world's largest collection of movie ratings. Netflix made these data available to the public during the Netflix Prize competition, which offered $1 million to the team that could use this data to create the most accurate movie recommendation system. The data set contains approximately 480,000 user IDs, 17,770 movies/TV series[1], and over 100 million

[1] In the rest of the paper, we refer to both the movies and the TV series at Netflix as movies for simplicity.

ratings, which are about a 10% random sample of all ratings submitted by consumers.

We believe that our data provide rich evidence to study the Long Tail effect for several reasons. First, as Anderson (2004; 2006) argues, Netflix operates in the environment that is most conducive to observing the Long Tail effect because Netflix is an Internet company which capitalizes on the favorable economics of DVD distribution to offer much larger product variety to consumers than its closest brick-and-mortar competitors such as Blockbuster. Second, Netflix is known to utilize one of the most advanced movie recommendation systems and therefore Netflix users are very likely to be able to discover niche titles fast. Moreover, consumer-to-consumer communication plays a very important role in the movie industry (De Vany and Walls, 1996), so unlike in the case of narrowly focused Internet retailers, we expect consumers to also be able to discover new products through the word-of-mouth effect. Third, Internet penetration increased dramatically in the 2000-2005 time period and, therefore, by studying Netflix data over time, we can observe temporal changes in movie consumption and study the Long Tail effect. The data that Netflix offers allows doing so both at the movie and at the individual consumer level, which is quite rare. Finally, the Netflix Prize data set is the largest among all the related studies.

Table 1 presents descriptive statistics of the ratings by year as well as the product variety collected from the company's annual reports[2]. The average movie rating, measured on a scale of one to five, has improved from 3.36 in 2000 to 3.67 in 2005, although there was a slight dip in 2002. In addition, the number of rated movies almost exactly quadrupled from 2000 to 2005, indicating that the consumers at Netflix watched an increasingly wider range of movies. Despite this demand diversification, the number of unrated movies increased more than five times from 2002 to 2005. Compared with the fast-growing number of different movies, the total number of ratings and the rater base expanded even faster; both of them increased about 50 times from 2000 to 2005. The substantial increase in demand is likely to be attributed to the quick adoption of DVD players among consumers as well as to Netflix's successful transition to a prepaid subscription service from a pay-per-use pricing model (Shih et al., 2007).

[2]The company reports are available from 2002 when Netflix went public.

Table 1: Descriptive Statistics

Year	Average Movie Rating	Total Number of Rated Movies	Total Number of Ratings	Number of Raters	Reported Product Variety	Number of Unrated Movies
2000	3.36	4,470	924,443	8,227	NA	NA
2001	3.39	6,538	1,769,030	19,801	NA	NA
2002	3.38	8,418	4,342,870	51,732	14,500	6,082
2003	3.40	11,949	9,985,340	117,500	18,000	6,051
2004	3.59	15,506	30,206,570	259,407	35,000	19,494
2005	3.67	17,768	53,250,070	451,435	55,000	37,232
Mean/year	3.47	10,775	16,746,387	151,350	30,625	17,214
Stdev/year	0.14	5,214	20,918,458	173,543	18,553	12,791

Source: Netflix Prize Data and Company Reports

Figure 1 shows that the number of monthly ratings increased exponentially from January 2000 to October 2005[3] and that the number of rated movies increased quadratically during the same period. A comparison of the two functions suggests that the number of ratings has grown considerably faster than product variety has, which further indicates that the demand for niche movies should emerge very quickly. These observations seem to offer favorable conditions for us to expect an increasing demand for niches. A relevant question to ask is whether the active product variety is growing because a lot of new movies are being released or because consumers keep discovering previously released titles. Our data indicates that slightly over 1,000 new movies are released every year, on average, while about 3,000 of movies are newly rated every year. In addition, between 70 and 80 percent of Netflix rentals are reported to originate from the back catalog movies rather than new releases (Flynn, 2006). These observations suggest that the product variety growth is primarily due to discovery of older movies by consumers. The more precise answer to this question is complicated by the fact that many movies are released on DVDs later than in theaters, but this gap continues to decrease over time. Further, most movies are released in several DVD versions at different points in time which makes it hard to exactly delineate ratings of "old" vs. "new" movies.

[3]The data were incomplete in November and December 2005, so we excluded them from our analysis.

Figure 1: Monthly Numbers of Rated Movies and Ratings

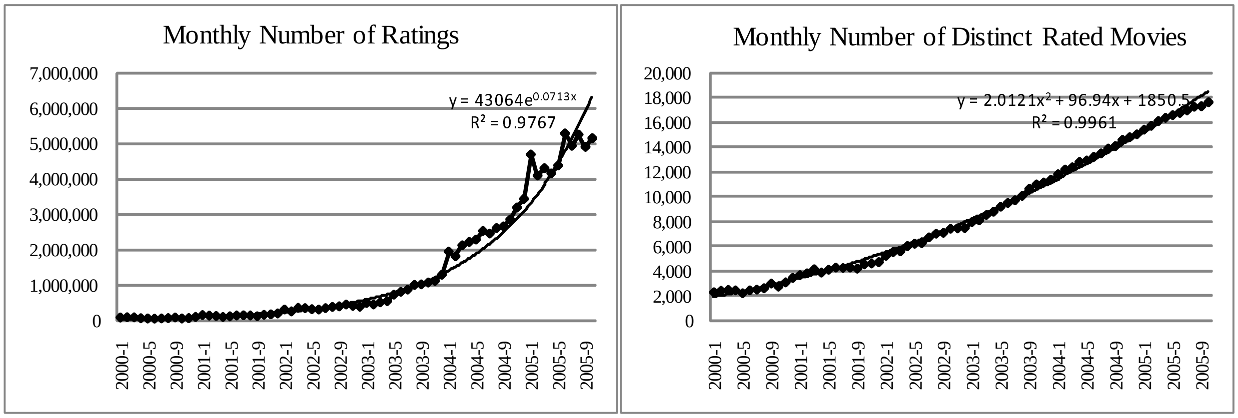

In much of our analysis, we elect to work with monthly (instead of daily or yearly) data and therefore we aggregate all variables at the monthly level. By so doing, we ensure both an adequate sample size in each month for each movie and enough observations over time for statistically significant estimates. This approach also makes sense given that Netflix charges its consumers monthly for services and therefore we expect consumer behavior to be tied to a monthly horizon. As shown in the descriptive statistics, the number of ratings, the number of rated movies, and the average ratings contain time trends that may cause a false conclusion of the relationship of the three variables. To mitigate this problem, we add a time trend of the 70 months from January 2000 to October 2005 to our data. Table 2 reports the correlations for each movie.

Table 2: Correlation Table of Monthly Netflix Ratings

Variables	Trend	Number of Ratings	Average Rating
Trend	1.000		
Number of Ratings	0.135	1.000	
Average Rating	0.146	0.140	1.000

Figure 2 (left) illustrates the distribution of the number of ratings after pooling all the observations from 2000 to 2005. As can be seen, demand is approximately logarithmically distributed: the top 5% of movies constitute close to 65% of the total number of ratings; the top 10% of movies contribute to almost 80%; the top 20% of movies generate approximately 90%, implying that demand is highly likely to be concentrated on a few titles, i.e., the hits.

It is important to stipulate here that the ratings data that we possess only reflect the number of movies rated, but not all customers rate all movies that they watch. On the other hand, customers do not have to watch the movie to be able to rate it, so ratings data might actually present a fuller picture of consumer demand for movies than rental data would. Furthermore, previous literature (see Chen et al. 2004) suggested a strong connection between product demand and the number of consumer reviews. One potential problem with utilizing movie ratings as a proxy for demand is that online review data is known to be biased because users tend to review items that they extremely like or dislike, thus causing the histogram to be J-shaped or U-shaped, see (Hu et al., 2007; Dellarocas and Narayan, 2007; Dellarocas and Wood, 2008), and citations therein. To check for the possibility of this bias in our data, we plot the histogram of the rating values on a scale from one to five in Figure 2 (right). We observe that the users gave the rating of four most frequently (31,754,330 ratings), followed by the ratings of three, five, two, and one. The bell-shaped histogram is dramatically different from the data reported by other papers (e.g., Dellarocas and Wood 2008 report that 99% of eBay transactions result in positive reviews) and it seems to suggest that Netflix users are not biased toward giving only extremely high or extremely low ratings. We believe that, while previous studies focused on the propensity to review products on eBay.com and Amazon.com, our data are fundamentally different for two reasons. First, we have pure *ratings* data rather than review data, and giving a rating is much less costly to a user than writing a review. Second, the recommendation system of Netflix directly incentivizes users to reveal their truthful and complete preference for movies to improve their recommendations. We therefore proceed with utilizing ratings as a proxy for demand, although it should be understood that we imply the number of ratings.

Figure 2: Number of Ratings Distribution and Histogram of Rating Values (2000-2005)

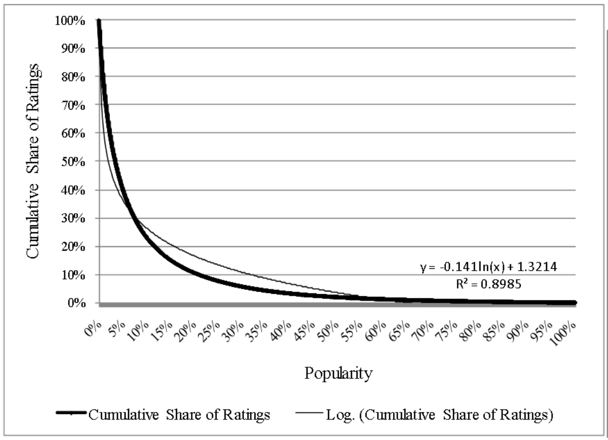

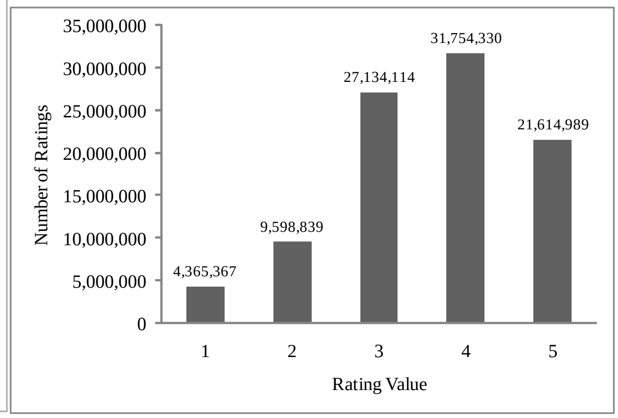

Table 3 demonstrates the summary statistics at the consumer level. The number of movies rated per person every month is highly skewed toward the high percentiles, indicating that a small group of people rate a massive number of movies each month. It is possible that the large number of movies rated, such as 43 for 90th percentile and 227 for 99th percentile, contain a large number of ratings given by users to train the recommendation system because a user can only watch a limited number of movies every month. Ratings submitted during the training process can result in a "contamination" of the data because the ratings of previously watched movies may not reflect the current popularity of a movie. In order to alleviate this issue and provide a robustness check, we purged the data with the monthly number of rated movies more than 30. We choose the cutoff point of 30 because watching 30 movies a month is probably the maximum number of movies that a heavy user is technically allowed to watch within Netflix rental system. Our results remain qualitatively and quantitatively similar so in the paper we report results without dropping any ratings.

Furthermore, it does not appear that heavy users have a tendency to give higher or lower ratings because the frequency of ratings very weakly correlates with the average rating (correlation = 0.007) and with the variance of ratings (correlation = 0.107). Figure 3 further shows that consumer ratings are almost normally distributed except that the right tail is censored at the rating of 5 because of the limit of the rating scale. This nearly normal distribution of consumer ratings provides further evidence that the users at Netflix may not be biased toward rating the movies that they extremely like or extremely dislike. Furthermore, the variance of ratings is skewed, which suggests that the majority of the people tend to be consistent in their ratings.

Table 3: Monthly Consumer Breakdown (N=4,740,731)

	Mean	Std	Skewness	1%	10%	25%	50%	75%	90%	99%
Number of Movies Rated	19	51	16.75	1	1	3	7	16	43	227
Average Rating	3.58	0.72	-0.44	1.3	2.75	3.14	3.6	4	4.5	5
Variance of Ratings	0.70	0.64	1.42	0	0	0.22	0.59	1	1.55	2.88

Source: Netflix Prize Data

Figure 3: Monthly Rating Distributions

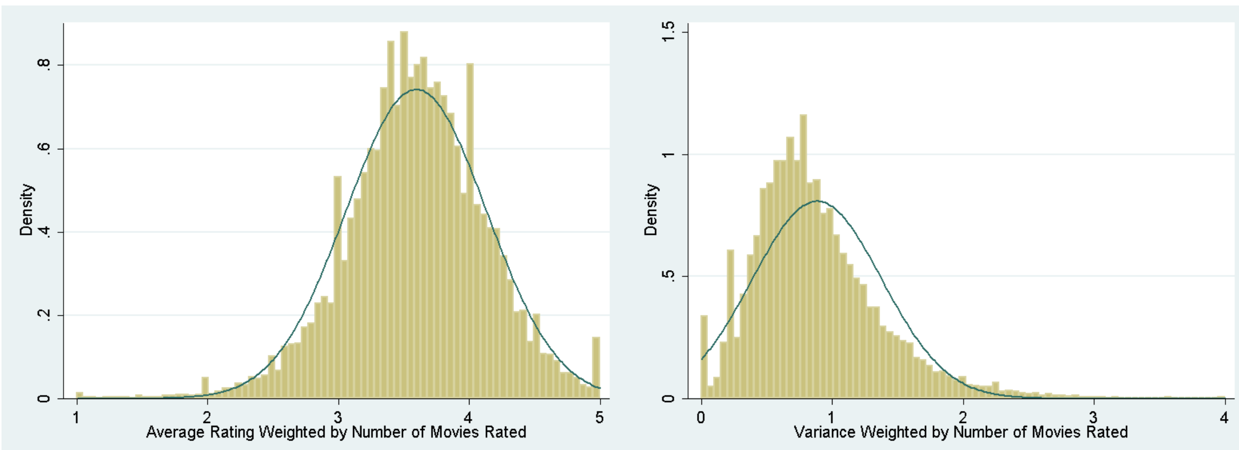

Although the Netflix Prize data provide millions of movie ratings, they contain no movie characteristics, such as genre, critics' ratings and MPAA ratings. The movie characteristics have shown a direct association with movie demand in some studies (Eliashberg et al., 2006) and the lack of this information may cause an omitted variable bias. In order to cope with this potential problem, we designed a web crawler to query the Yahoo Movies website, a database of movie information, and to automatically retrieve movie characteristic information on about 1,500 movies. We then joined the movie characteristics to the Netflix rating data to control for movie heterogeneity.

Table 4 exhibits the descriptive statistics of the combined data from Netflix and Yahoo Movies. Note that one movie may be tagged with multiple genres. Similar to the original Netflix Prize data, the total number of ratings and the total number of rated movies rapidly increased from 2000 to 2005. Admittedly, the Yahoo data is not nearly as comprehensive as the Netflix Prize data since critics tend to cover only hit movies, which limits the product variety on Yahoo. Table 5 presents the correlations for each month at the monthly level, which are all quite small.

Table 4: Descriptive Statistics of the Yahoo Movies

Year	Total Number of Ratings	Total Number of Rated Movies	Average Movie Rating	Average Critics Rating
2000	118,039	492	3.40	6.90
2001	247,419	692	3.40	6.94
2002	608,559	865	3.79	6.94
2003	1525,327	1,134	3.41	6.91
2004	4521,530	1,391	3.59	6.88
2005	7835,578	1,546	3.67	6.91
Mean/Year	2476,075	1,020	3.54	6.91
Stdev/Year	3092,796	373	0.15	0.02

Variable	Number of Movies	Percentage of Movies	Variable	Number of Movies	Percentage of Movies
MPAA-R	702	45.41%	GENRES-HORROR	114	7.37%
MPAA-PG13	460	29.75%	GENRES-SEQUEL	109	7.05%
MPAA-PG	195	12.61%	GENRES-MUSICAL	103	6.66%
MPAA-NR	128	8.28%	GENRES-ANIMATION	80	5.18%
MPAA-G	52	3.36%	GENRES-DOCUMENTARY	68	4.40%
MPAA-NC17	5	0.32%	GENRES-WAR	63	4.08%
GENRES-DRAMA	806	52.13%	GENRES-BIOPIC	60	3.88%
GENRES-COMEDY	640	41.40%	GENRES-SPORTS	60	3.88%
GENRES-ACTION	428	27.68%	GENRES-TEEN	49	3.17%
GENRES-ADAPTION	426	27.55%	GENRES-REMAKE	37	2.39%
GENRES-ROMANCE	354	22.90%	GENRES-POLITICS	32	2.07%
GENRES-THRILLER	339	21.93%	GENRES-WESTERN	17	1.10%
GENRES-CRIMES	199	12.87%	GENRES-HOLIDAY	10	0.65%
GENRES-FANTASY	192	12.42%	GENRES-CLASSICS	4	0.26%
GENRES-FOREIGNART	189	12.23%	GENRES-MISCELLANEOUS	2	0.13%
GENRES-FAMILY	122	7.89%			

Source: Netflix Prize Data and Yahoo Movies

Table 5: Correlation Table of the Monthly Yahoo Movies Rating

Variables	Trend	Number of Ratings	Average Rating	Average Critics Rating
Trend	1.000			
Number of Ratings	0.178	1.000		
Average Rating	0.143	0.147	1.000	
Average Critics Rating	-0.004	0.011	0.012	1.000

4.2 Method and Evaluation

We are interested in studying demand for hit vs. niche movies and therefore we define the variable $POPULARITY_{jt}$ to reflect how popular movie j is at time t. In particular, we rank the number of ratings that each movie receives within each month in a descending order and we use this rank to reflect the movie's popularity. Note that a higher (lower) rank indicates a less (more) popular movie. To be able to test our hypotheses, we use two definitions of $POPULARITY_{jt}$: one is in absolute terms and the other is in relative terms. The absolute rankings, e.g., the top 1, the top 100 movies, were used in the previous literature including Anderson (2004). Alternatively, we propose to rank movies in relative terms, e.g., the top 0.1%, the top 0.2%, thus adjusting for current product variety (the total number of movies rated this month).

We define $PRODUCT\_VARIETY_t$ as the total number of different movies that were rated during the time period t. Note that this variable reflects the active product variety as there are many movies which are not rated in a given month. We believe that the active product variety is a more relevant variable than the total variety because 1) products that are not discovered by consumers (or that are discovered but forgotten) should not be taken into account when ranking popularity and 2) it accounts for both the product offering and consumer demand. Clearly, by using active rather than total product variety, we are more likely to find evidence of the Long Tail effect. Furthermore, we define the demand for individual movies with a proxy $DEMAND_{jt} = \frac{\sum_i \mathbf{1}_{ijt}}{\sum_i \sum_j \mathbf{1}_{ijt}}$, which reflects the share of the number of ratings among all the rated movies during a given time period. Similar to other "webmining" data, such as sales ranks, page views, click streams, and online music samplings, shares of ratings have been used in previous research as proxies for demand (Dewan and Ramaprasad, 2008; Chevalier and Goolsbee, 2003; Tucker and Zhang, 2009). Using $POPULARITY_{jt}$, $PRODUCT\_VARIETY_t$ and $DEMAND_{jt}$, we can analyze the distribution of the cumulative demand every year both in absolute terms and in relative terms and we can also compute the yearly Gini coefficients.

Recommendation systems with average ratings have proliferated on the Internet and have had a growing impact on consumers' decision making (Dellarocas, 2003). In this paper, unlike in other studies, we are able to control for product rating using the AVG_RATING_{jt} variable, which reflects the average rating that the movie received in a given month. This control is important because the Long Tail effect might be simply a manifestation of the fact that hit movies have deteriorated in quality over time, which maybe an important driver of the Long Tail effect. Movie ratings in our data are probably reflective of the public's tastes and acceptance of the movies and it is known that ratings can influence the choices of future consumers (Eliashberg et al., 2006). All variable definitions at the movie level are presented in Table 6.

Table 6: Movie-level Analysis Variable Definitions

Variable	Definition
$POPULARITY_{jt}$	Order statistics ranked by the number of ratings of a particular movie j during time period t, either in absolute or relative terms.
$DEMAND_{jt}$	Share of the number of the ratings among all the rated movies during time period t.
$PRODUCT\_VARIETY_t$	Total number of rated movies during time period t.
AVG_RATING_{kt}	Average rating of the movies whose popularity ranks are below cutoff k.
$MPAA_k$	Fraction of different MPAA ratings within the movies whose popularity ranks are below cutoff k, e.g., 20% Rs, 30% PG-13s.
$AVG\_CRITICSRATING_k$	Average critics' rating of the movies whose popularity ranks are below cutoff k.
$GENRES_k$	Fraction of different genres within the movies whose popularity ranks are below cutoff k, .e.g., 10% dramas, 20% comedies.
$TREND$	1, 2, 3, ... a sequence representing the time trend.

In addition to the movie-level variables, we further define variables for our user-level analysis. We define $NICHESEEKING_{it}$ as the average ranking of the movies that the consumer rates in a given month. In essence, $NICHESEEKING_{it}$ is a summary statistic of the movie-level variable $POPULARITY_{it}$. A high value of $NICHESEEKING_{it}$ means that this particular consumer tends to watch more niche movies. Likewise, Brynjolfsson et al. (2007) use the average sales rank of all the products that a consumer purchases. However, we believe that focusing on averages may be too restrictive if the distribution is skewed. For example, a consumer with a skewed set of purchases ranked 1st, 2nd, 3rd, and 10,000th, has an average sales rank of 2,501.5. Another consumer with a set of purchases ranked 2,400th, 2,410th, 2,420th, and 2,430th has an average sales rank of 2,415. Although the average sales rank of the second consumer is smaller than that of the first consumer, it is hard to argue that the second consumer is more inclined to watch hits than the first consumer. Hence, in addition to the mean measurement, we calculate the median, the top 10%, and the bottom

10% of the $POPULARITY_{jt}$ values to obtain more complete information about consumer choices. Furthermore, we divide these metrics by monthly product variety to obtain relative measurements. These relative measurements adjust for both the increasing product variety and the skewness of demand distribution.

In order to control for consumer heterogeneity over time, we define $FREQUENCY_{it}$ as the number of movies that user i rated in month t. In marketing, certain theoretical constructs such as the Dirichlet model suggest a strong link between purchase frequency and brand choice. In particular, it is often found that most consumers of a brand are low-frequency buyers (Chatfield and Goodhardt, 1975; Goodhardt et al., 1984). These light buyers often constitute the majority of the customers who purchase the popular brand (McPhee, 1963) because of the "super-star" effect (Rosen, 1981). In addition, McPhee (1963) explains that consumers who are familiar with the alternatives tend to consume the niche products. Therefore, consumers with high-consumption frequency are likely to consume more niche products than those with low-consumption frequency because the former may be better informed of the variety of products than the latter.

Furthermore, we define $RATING\_PROPENSITY_{it}$ and $RATING\_VARIANCE_{it}$ as the average and the variance of the ratings that user i gives in month t. These two measurements are likely to reflect people's tastes and movie acceptance. For example Clemons et al. (2006) demonstrate the relationship between variance of ratings and demand for products. Further, Hu et al. (2007) recommend controlling for the standard deviation of ratings as well as for two modes to overcome consumer under-reporting bias. Since in our case the distribution of ratings is symmetric, we do not control for the modes. All user-level variables are defined in Table 7. Table 8 presents the correlations of consumer characteristics that we possess. We do not report any strong correlations.

Table 7: User-level Analysis Variable Definitions

Variable	Definition
$NICHESEEKING_{it}$	Popularity of the movies that user i rated in month t, measured as the mean, the median, the top 10%, and the bottom 10% of the $POPULARITY_{jt}$, both in absolute and relative terms.
$FREQUENCY_{it}$	Number of movies rated by user i in month t.
$RATING\_PROPENSITY_{it}$	Average rating given by user i in month t.
$RATING\_VARIANCE_{it}$	Variance of the ratings given by user i in month t.

Table 8: Correlation Table of Consumer-level Variables

Variables	*TREND*	$FREQUENCY_{it}$	$RATING\_$ $PROPENSITY_{it}$	$RATING\_$ $VARIANCE_{it}$
TREND	1.000			
$FREQUENCY_{it}$	0.032	1.000		
$RATING\_PROPENSITY_{it}$	0.110	0.007	1.000	
$RATING\_VARIANCE_{it}$	-0.021	0.107	-0.211	1.000

5 Estimation and Results

5.1 Movie-level Analysis

We begin our analysis by presenting the distribution of cumulative demand at various cutoff points by year as well as by calculating the yearly Gini coefficients. Table 9 reports this initial check of the dynamics of the demand shift over time. As can be seen from the upper part of the table, cumulative demand tended to decline throughout our study period at all the absolute cutoff points, suggesting that both hits and niches have had decreasing demand over time. In particular, demand for the top 10 titles has declined from 4.39% to 2.43% over the course of five years. However, a different picture emerges in the bottom part of Table 10, where cutoff points are relative to active product variety. We see that the top quintiles, i.e., the hits such as the top 1%, the top 5%, and the top 10%, exhibit increasing cumulative demand over time. However, the bottom quintiles exhibit no obvious patterns. Furthermore, the Gini coefficients exhibit no obvious trend either.

While this initial check suggests that there might be at least partial support for Hypotheses 1a and 2b, it has limited statistical power. In order to perform a statistical hypothesis test, we aggregate demand data at the monthly rather than the annual level. Since demand for movies increases exponentially in our data, we employ the following logarithmically transformed time series model

$$\log(\sum_{j \in k} DEMAND_{jt}) = \beta_0 + \beta_1 TREND + \epsilon_{kt} \quad (1)$$

$$k \in \text{cutoff points according to } POPULARITY_{jt}$$

to examine the dynamics of demand over the 70 months from January 2000 to October 2005. Table

Table 9: Distribution of the Cumulative Demand by Year

Cumulative of Movies	Cumulative of Demand 2000	2001	2002	2003	2004	2005
Bottom 10	0.01%	0.01%	0.00%	0.00%	0.00%	0.00%
Bottom 50	0.10%	0.03%	0.01%	0.01%	0.00%	0.00%
Bottom 100	0.11%	0.06%	0.02%	0.01%	0.00%	0.01%
Bottom 1,000	0.24%	0.12%	0.11%	0.08%	0.03%	0.03%
Top 500	70.84%	66.04%	60.36%	51.74%	47.11%	48.25%
Top 100	26.64%	27.67%	24.25%	19.63%	16.08%	16.96%
Top 50	15.66%	17.15%	14.57%	11.64%	9.92%	9.95%
Top 10	4.39%	4.51%	3.66%	3.05%	2.28%	2.43%

Cumulative of Movies	Cumulative of Demand 2000	2001	2002	2003	2004	2005
Bottom 1%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Bottom 5%	0.02%	0.02%	0.02%	0.03%	0.02%	0.03%
Bottom 10%	0.06%	0.06%	0.08%	0.12%	0.08%	0.09%
Bottom 20%	0.20%	0.20%	0.35%	0.41%	0.29%	0.30%
Top 20%	86.60%	88.80%	87.10%	86.30%	88.50%	90.08%
Top 10%	67.50%	73.10%	72.90%	72.60%	75.40%	78.50%
Top 5%	46.50%	54.80%	55.90%	56.10%	58.40%	62.60%
Top 1%	14.20%	20.70%	21.5%	22.20%	22.40%	25.50%
Gini Coefficients	80.67%	83.15%	81.93%	81.55%	83.39%	84.91%

10 presents the results of Model 1. We observe that the demand for the top 10, the top 100, and the top 1,000 movies decreases, while the demand for the top 0.1%, the top 1%, and the top 10% of movies increases over time. All coefficients are highly significant and this simple regression model has significant exploratory power. These results imply that the demand for hits tends to be less and less concentrated in absolute terms, but more and more concentrated in relative terms, which agrees with our Hypotheses 1a and 1b. In particular, the demand for the top 0.1% of movies increases over four times as fast as the demand for the top 10% of movies (coefficients 0.0087 and 0.0019, correspondingly), indicating that the demand for the "hits of the hits" significantly outpaces the demand for less popular movies when popularity is measured in relative terms. Further, demand for niche movies decreases over time, whether popularity is measured in absolute or relative terms.

In particular, the bottom 0.1% movies lose demand slightly faster than the bottom 1% and the bottom 10% of movies (coefficients -0.0174, -0.0173, and -0.0112, correspondingly), indicating that the demand for the "niches of the niches" deteriorates faster than for more popular movies. These results render support to Hypothesis 1b but not 1a: no matter how popularity is measured, there is no evidence that demand for niche movies increases.

Table 10: Regression of Hits and Niches Over Time

	Top 10	Top 100	Top 1,000	Top 0.1%	Top 1%	Top 10%
Intercept	-1.2040	-0.4289	-0.0167	-1.5054	-0.8455	-0.2049
	(0.0176)***	(0.0116)***	(0.0038)***	(0.0228)***	(0.0170)***	(0.0062)***
TREND	-0.0023	-0.0040	-0.0025	0.0087	0.0055	0.0019
	(0.0004)***	(0.0003)***	(0.0001)***	(0.0006)***	(0.0004)***	(0.0002)***
R^2	0.29	0.74	0.91	0.78	0.72	0.69

	Bottom 10	Bottom 100	Bottom 1,000	Bottom 0.1%	Bottom 1%	Bottom 10%
Intercept	-3.2564	-2.6342	-1.5280	-3.2564	-2.9608	-2.3238
	(0.0241)***	(0.0237)***	(0.0201)***	(0.0241)***	(0.0243)***	(0.0182)***
TREND	-0.0174	-0.0310	-0.0305	-0.0174	-0.0173	-0.0112
	(0.0006)***	(0.0006)***	(0.0005)***	(0.0006)***	(0.0006)***	(0.0004)***
R^2	0.93	0.98	0.98	0.93	0.93	0.90

1) *p-value<0.05. **p-value<0.01, ***p-value<0.001

2) The rows below the estimates are standard errors.

To conclude this part of the analysis, we conduct a hypothesis test of the famous 80/20 Pareto Rule on the entire data set from 2000 to 2005. We reject the null hypothesis that 20% of the movies contribute to 80% of the demand with 95% confidence ($t = 10.04$) and we estimate that the top 20% of movies constitute between 86.4% and 84.4% of total demand, significantly more than 80%. In fact, the percentage of demand share for the top 20% of movies has increased from 86.6% in 2000 to 90.08% in 2005. This result, once again, goes against the notion of the Long Tail effect that the top 20% of movies will constitute a smaller and smaller share of demand.

So far, our econometric model is limited to capturing the time trend. To provide a robustness test which controls for the movie characteristics data, we utilize the combined data set from Netflix and Yahoo Movies and employ Model 2. Once again, the goal is to study demand for the hits and the niches over time at both absolute and relative cutoff points.

$$\log(\sum_{j \in k} DEMAND_{jt}) = \beta_0 + \beta_1 TREND + \beta_2 \log(AVG\_RATING_{kt}) + \beta_3 \log(AVG\_CRITICSRATING_k) + \beta_4 MPAA_k + \beta_5 GENRES_k + \epsilon_{kt} \qquad k : \text{cutoff points.} \tag{2}$$

Tables 11 and 12 present regression results for the hits and the niches, respectively[4]. As is clear from the tables, even after controlling for movie characteristics, the results of our analysis are very consistent with the results of Model 1. Namely, *TREND* coefficients for the top 10 and the top 100 movies are -0.0125 and -0.0045, suggesting that the demand for hits decreases over time in absolute terms. However, the same coefficients for the top 1% and the top 10% of movies are 0.0067 and 0.0032, implying that the demand for hits tends to increase over time in relative terms. This evidence is consistent with our Hypotheses 1a and 1b. For the niches, we see that demand decreases in both absolute and relative terms, which is consistent with the findings in Model 1 and supportive of Hypotheses 1b but not 1a. We note that, although movie characteristics add to the explanatory power of the model, there are no consistent results to report here, which is not surprising given the small size of the sample. In particular, we find no evidence that the demand for the hits or the niches is associated with the user or critic ratings. Furthermore, since the high R^2s in these regressions indicate potential multicollinearity problems, we examine the correlation table[5] of the variables in each regression and find no high correlations.

[4]Due to a smaller sample size of the combined data set, the top 0.1%, the top 1,000, the bottom 0.1% and the bottom 1,000 of the movies are not available in certain months. Therefore, we drop these cutoff points.

[5]We omit to report them in the paper because of the page limitation.

Table 11: Regression of Hits Over Time Controlling for Movie Characteristics

	Top 10		Top 100		Top 1%		Top 10%	
(Intercept)	-0.2112	(1.625)	0.9319	(1.1805)	-1.6326	(1.1511)	1.1570	(1.3265)
TREND	-0.0125	(0.0025)***	-0.0045	(0.0012)***	0.0067	(0.0022)**	0.0032	(0.0014)*
AVG_RATING	-0.0564	(0.8887)	-0.2291	(0.3471)	-1.1448	(0.8123)	-0.2506	(0.4908)
AVG_CRITICSRATING	-0.3145	(0.3208)	-0.3154	(0.288)	0.0110	(0.1731)	-0.0955	(0.4128)
NC17	NA	NA	NA	NA	NA	NA	NA	NA
NR	-0.0081	(0.9899)	0.0760	(0.7423)	2.0797	(0.872)*	-1.9477	(0.9997)·
PG	-0.1730	(0.6137)	0.0033	(0.4497)	0.3988	(0.3874)	-0.0288	(0.496)
PG13	-0.1782	(0.7271)	0.3352	(0.6488)	1.3259	(0.4337)**	-0.5183	(0.7168)
R	-0.0618	(0.6764)	-0.4911	(0.6738)	0.6289	(0.4324)	-1.1511	(0.6988)
COMEDY	0.4124	(0.3002)	-0.0777	(0.2159)	0.4539	(0.2497)·	-0.3945	(0.2888)
DRAMA	-0.4005	(0.2687)	0.1138	(0.3051)	-0.1223	(0.2039)	-0.3823	(0.3095)
MUSICAL	0.1683	(0.6432)	-0.4754	(0.4324)	-0.1159	(0.8235)	0.7918	(0.4473)·
ROMANCE	-0.1585	(0.2387)	-0.2622	(0.3077)	0.4138	(0.1905)*	0.1661	(0.3516)
ADAPTION	-0.0074	(0.2442)	0.1018	(0.211)	0.1163	(0.2753)	0.4643	(0.2738)·
CRIME	0.2056	(0.3758)	-0.0941	(0.2564)	-0.0534	(0.2581)	-0.0346	(0.3235)
TEEN	-1.0394	(0.7395)	0.3700	(0.7304)	-0.6559	(0.3364)·	0.6621	(0.7024)
ACTION	-0.0549	(0.2592)	0.1113	(0.3553)	-0.0734	(0.1943)	-0.5702	(0.2935)
REMAKE	0.0891	(0.6555)	-0.0598	(0.6308)	-0.1714	(0.8997)	-0.7901	(0.675)
HORROR	-0.6001	(0.4106)	-0.0381	(0.4686)	-0.5429	(0.4044)	0.6323	(0.4308)
THRILLER	-0.2265	(0.2443)	0.2012	(0.2889)	0.2365	(0.2114)	-0.7378	(0.4296)·
FAMILY	-0.3910	(0.7283)	-0.2961	(0.5232)	0.4709	(0.5009)	-1.5292	(0.4923)**
FANTASY	0.2360	(0.3898)	-0.6669	(0.3036)*	0.4020	(0.3863)	-0.0407	(0.5242)
ANIMATION	0.5634	(0.5272)	0.6598	(0.5482)	0.9530	(0.5259)·	0.8018	(0.4322)·
FOREIGNART	-0.1924	(0.4566)	-0.0635	(0.3561)	-0.0199	(0.4648)	0.782	(0.4056)·
BIOPIC	1.224	(0.6166)·	1.0227	(0.5376)·	1.3600	(0.385)**	-0.8277	(0.5837)
WESTERN	0.4241	(0.4344)	-2.4168	(0.9217)*	0.3939	(0.3732)	-3.1213	(1.3105)*
SEQUEL	-0.7725	(0.3531)*	-0.0031	(0.3412)	-1.4728	(0.3678)***	0.4473	(0.4762)
DOCUMENTARY	0.3343	(0.6254)	0.1212	(0.5647)	0.6374	(0.7691)	0.7207	(0.6162)
HOLIDAY	NA	NA	NA	NA	NA	NA	NA	NA
WAR	0.0959	(0.6129)	-0.4474	(0.5955)	-0.7623	(0.5320)	0.0292	(0.5784)
POLITICS	-0.0787	(0.9745)	0.0433	(1.0524)	-0.3161	(0.6154)	1.8766	(1.1944)
SPORTS	-0.6627	(1.0154)	-0.1196	(0.5578)	2.9889	(1.8628)	1.0517	(0.8126)
CLASSICS	-0.0524	(1.1264)	5.2394	(2.5497)	-0.8332	(1.3703)	7.5167	(2.4063)**
MISCELLANEOUS	NA	NA	NA	NA	NA	NA	NA	NA
R^2	0.87		0.97		0.93		0.95	

1) · p-value<0.1, *p-value<0.05. **p-value<0.01, ***p-value<0.001

2) The columns in parentheses are standard errors.

3) Categorical variable G is not included in the regression to avoid multicollinearity and provide base-line scenario.

4) *NC17, HOLIDAY, MISCELLANEOUS* are small portion of the movies, so they are dropped in some regressions by the statistics package R.

Table 12: Regression of Niches Over Time Controlling for Movie Characteristics

	Bottom 10		Bottom 100		Bottom 1%		Bottom 10%	
(Intercept)	-5.5706	(1.4132)***	-10.2595	(3.3938)**	-8.5711	(0.6288)***	-6.3935	(2.4689)*
TREND	-0.0761	(0.0025)***	-0.0626	(0.0037)***	-0.0483	(0.002)***	-0.0348	(0.0025)***
AVG_RATING	0.1321	(0.2281)	0.4685	(0.4391)	0.0700	(0.1832)	0.5399	(0.3915)
AVG_CRITICSRATING	0.0047	(0.3418)	0.9259	(1.0538)	-0.2198	(0.2197)	-0.1958	(0.6978)
NC17	1.1090	(2.8282)	10.5379	(6.1251)·	1.8434	(2.1605)	-0.6503	(5.6898)
NR	-0.5430	(0.9089)	5.1206	(2.5358)·	1.4033	(0.5991)*	-1.6417	(2.0826)
PG	-0.1868	(0.8478)	3.1803	(2.0526)	1.3746	(0.6251)*	0.6438	(1.4855)
PG13	-1.1532	(0.9923)	4.8037	(2.0598) *	1.2972	(0.5697) *	0.4963	(1.5737)
R	-0.5430	(0.9539)	4.5993	(2.1163)*	1.0991	(0.5695)·	0.9569	(1.6598)
COMEDY	0.0809	(0.3674)	-0.1557	(0.8403)	-0.2305	(0.2819)	0.2564	(0.7464)
DRAMA	-0.4223	(0.3436)	0.2800	(0.8442)	0.1970	(0.3202)	0.1873	(0.7273)
MUSICAL	0.0336	(0.5062)	-2.3684	(1.3791)·	0.3728	(0.4005)	-0.9844	(1.0259)
ROMANCE	0.4458	(0.3525)	1.9933	(1.0637)·	0.4282	(0.2783)	0.8982	(0.6542)
ADAPTION	-0.1524	(0.2879)	-1.5423	(0.8447)·	-0.2054	(0.2231)	0.2105	(0.6525)
CRIME	0.0678	(0.4562)	-0.3655	(1.325)	0.1142	(0.356)	-0.6137	(0.8921)
TEEN	0.6102	(0.7373)	-2.5442	(2.1831)	-0.2967	(0.6663)	0.1453	(1.8509)
ACTION	-0.2826	(0.3164)	0.808	(0.9907)	-0.1199	(0.3018)	-0.5783	(0.7174)
REMAKE	0.9343	(1.1546)	2.9576	(2.3504)	0.7553	(0.9053)	1.7089	(2.0801)
HORROR	-0.3486	(0.5275)	0.7117	(1.8792)	-0.3753	(0.4507)	-1.1074	(1.3466)
THRILLER	0.3597	(0.4547)	-2.5496	(1.2024)*	0.3210	(0.2723)	-0.4651	(0.9413)
FAMILY	-0.5721	(0.6725)	-2.2637	(1.703)	1.4683	(0.6359)*	-0.3804	(1.076)
FANTASY	-0.1365	(0.3568)	1.6073	(1.0053)	0.1317	(0.2876)	1.1371	(0.8915)
ANIMATION	-0.3968	(0.6606)	2.9043	(1.9104)	-1.0386	(0.5726)·	-1.9543	(1.4957)
FOREIGNART	0.1067	(0.4396)	-0.9004	(1.2902)	0.2958	(0.2936)	0.5178	(0.8469)
BIOPIC	-0.176	(0.7532)	1.0666	(1.8999)	-0.1679	(0.4928)	1.3764	(1.3385)
WESTERN	0.327	(2.0035)	-3.4022	(6.3635)	-1.7692	(1.6385)	-3.9491	(5.5983)
SEQUEL	-0.1200	(0.5144)	1.1387	(1.1508)	0.5035	(0.4477)	0.8430	(0.8951)
DOCUMENTARY	-0.1049	(0.9472)	1.6888	(1.8753)	0.4108	(0.6719)	2.9324	(2.2056)
HOLIDAY	0.7530	(1.0892)	-1.1441	(2.5587)	-0.264	(0.7889)	3.1889	(1.7753)·
WAR	0.8700	(0.8567)	-0.6802	(2.47)	0.2846	(0.7257)	-0.2296	(2.0855)
POLITICS	0.2928	(0.6438)	-7.2928	(2.6826)*	0.6005	(0.5822)	-2.6774	(1.808)
SPORTS	-0.7414	(0.7984)	-3.1143	(1.7759)·	-1.5846	(0.7797)*	-1.4592	(1.9339)
CLASSICS	2.4196	(3.2742)	0.8266	(7.0691)	1.3783	(2.5106)	11.7492	(5.9701)·
MISCELLANEOUS	NA	NA	-12.2229	(10.1459)	NA	NA	-0.1515	(13.1357)
R^2	0.98		0.99		0.97		0.96	

1) · p-value<0.1, *p-value<0.05. **p-value<0.01, ***p-value<0.001

2) The columns in parentheses are standard errors.

3) Categorical variable G is not included in the regression to avoid multicollinearity and provide base-line scenario.

4) *NC17, HOLIDAY, MISCELLANEOUS* constitute a small portion of the movies, so they are dropped in some regressions by the statistics package R.

5.2 Consumer-level Analysis

We now turn our attention to examining how individual consumers change their ratings over time in order to gain insights into the movie-level analysis. In particular, we examine how the propensity of each consumer to discover niche movies evolves over time, while controlling for observed user heterogeneity, such as rating frequency and variance of ratings. The data that we have lack other potentially significant consumer characteristics, such as demographics. In order to cope with this issue, we introduce a time-invariant preference for each consumer's movies through the panel data analysis and we further assume that the preference correlates with the observed characteristics of the consumer. This correlation is likely to be caused by the recommendation systems, which can influence an individual's preference based on his/her observed characteristics. The Hausman test (Hausman, 1978) further provides strong evidence of this correlation. Therefore, we employ the following fixed-effect regression to predict consumer propensity to rate movies:

$$\begin{aligned}\log(\mathit{NICHESEEKING}_{it}) &= \beta_0 + \beta_1 \mathit{TREND} + \beta_2 \mathit{FREQUENCY}_{it} + \beta_3 \mathit{RATING\_PROPENSITY}_{it} \\ &\quad + \beta_4 \mathit{RATING\_VARIANCE}_{it} + \beta_i + \epsilon_{it}, \qquad t \in 70 \text{ months.} \end{aligned} \tag{3}$$

The top of Table 13 presents results for absolute movie rankings while the bottom of the table presents the same results for relative movie rankings using Model 3. As is evident from the top portion of the table, all time-trend coefficients are positive and highly significant, suggesting that the absolute popularity rankings of the movies watched by the average consumer consistently increase. In other words, consumers tend to watch more and more niche movies over time when movies are ranked in absolute terms. In particular, the *TREND* coefficient for the bottom 10th percentile of the movies that a person watches (i.e., the obscure titles) is 0.0149, which is 15 times higher than the same coefficient for the top 10th percentile of the movies (i.e., the popular titles). This comparison suggests that consumers discover niche products much faster than they move away from the hits (again, if popularity is measured in absolute terms).

However, the picture changes completely when the popularity of the movies is measured in relative terms. The bottom part of Table 13 shows that time-trend coefficients are consistently negative, suggesting that, relative to the product variety that is available at that point of time, consumers tend to watch more and more popular movies. In particular, the *TREND* coefficient for the top 10th percentile of movies is -0.0256, which is about twice as high as for the bottom 10th percentile of movies, suggesting that consumer demand shifts toward more popular hits faster than it shifts away from less popular niches. Taken together, the results of Model 3 suggest that the growth rate of product variety is substantially higher than the speed at which consumers discover niche products. This finding is consistent with the results of Fleder and Hosanagar (2008) that recommendation systems guide similar consumers to the same products, which does not effectively help consumers discover products at the tail of the distribution.

Table 13: Fixed-Effect Model Results

	mean	median	top 10%	bottom 10%
(Intercept)	5.7328	5.5667	4.8737	6.2022
	(0.0051)***	(0.0057)***	(0.0074)***	(0.0057)***
TREND	0.1488	0.1250	0.0010	0.0149
	(0.0001)***	(0.0001)***	(0.0001)***	(0.0000)***
$FREQUENCY_{it}$	0.0008	0.0009	0.0029	0.0009
	(0.0000)***	(0.0000)***	(0.0000)***	(0.0000)***
$RATING\_PROPENSITY_{it}$	-0.1137	-0.0048	-0.2612	-0.0728
	(0.0010)***	(0.0012)***	(0.0014)***	(0.0011)***
$RATING\_VARIANCE_{it}$	0.2568	-0.0048	-0.6596	0.5166
	(0.0010)***	(0.0012)***	(0.0015)***	(0.0011)***
Overall R^2	0.0391	0.0132	0.0302	0.0663

	relative mean	relative median	relative top 10%	relative bottom 10%
(Intercept)	-2.3428	-2.4343	-3.1272	-1.7987
	(0.0052)***	(0.0057)***	(0.0074)***	(0.0057)***
TREND	-0.0118	-0.0142	-0.0256	-0.0118
	(0.0001)***	(0.0001)***	(0.0001)***	(0.0001)***
$FREQUENCY_{it}$	0.0006	0.0009	0.0028	0.0008
	(0.0000)***	(0.0000)***	(0.0000)***	(0.0000)***
$RATING\_PROPENSITY_{it}$	-0.1192	-0.1557	-0.2611	-0.0726
	(0.0010)***	(0.0011)***	(0.0014)***	(0.0011)***
$RATING\_VARIANCE_{it}$	0.2434	-0.0072	-0.6619	0.5142
	(0.0010)***	(0.0012)***	(0.0015)***	(0.0012)***
Overall R^2	0.0278	0.0114	0.0428	0.0586

*p-value<0.05, **p-value<0.01, ***p-value<0.001

Furthermore, the consistently positive and highly significant coefficients of $FREQUENCY_{it}$ indicate that heavier users tend to watch more niche movies. In particular, the coefficients for both absolute and relative medians are 0.0009. In other words, if an average consumer watches five more movies per month or one more movie per week, the median of his/her propensity for niche movies is likely to increase by 0.45% on average, holding other factors constant. Thus, it appears that heavy users are the ones that drive toward the Long Tail effect. Nevertheless, these heavy users constitute a relatively small segment of the entire population: as we demonstrated earlier, heavy users with a monthly frequency over the mean constitute less than 25% of all users. Although this small group of people tends to watch more niche movies, it does not seem to shift the entire demand from hits to niches. A comparison of coefficients for the top 10th percentile (0.0028 and 0.0029) and coefficients for the bottom 10th percentile (0.0008 and 0.0009) suggests that heavier users shift away from the hits three times as fast as they discover niches. That is, even heavy users are not as fast in discovering niches as they are in "forgetting" about hits.

Consistently negative and highly significant coefficients of $RATING\_PROPENSITY_{it}$ suggest that consumers who, on average, give higher ratings and may therefore be more satisfied tend to watch more popular movies. For example, the coefficients for both absolute and relative means are around -0.11, which suggests that increasing the average rating by one unit is associated with a 10% increase in the average movie popularity. In other words, the more popular movies generally satisfy people better than the obscure titles. Of course, it is possible that consumers who watch popular movies are systematically different from consumers who watch niche movies in that the former tend to rate all movies higher than the latter. Since we are unable to observe characteristics of individual consumers, out finding is subject to this limitation.

Finally, we note that $RATING\_VARIANCE_{it}$ is negatively associated with the median and the top 10th percentiles, while this variable is positively associated with the mean and the bottom 10th percentiles (in either absolute or relative terms). We interpret these mixed signs to imply that consumers with highly disperse rating tend to watch extreme hits and niches. In other words, the extreme hits and the extreme niches receive more polarized ratings from those consumers. Further, the popularity of the movies that those consumers watch tends to skew toward the niches. In other words, consumers with highly disperse ratings watch a larger quantity of hits than niches, but the niches that they watch are generally extremely obscure.

6 Conclusion and Discussion

The original definition of the Long Tail effect by Anderson (2004) states that demand for hit products will decrease while niches will constitute a larger and larger proportion of demand on the Internet. In this paper we argue that one has to be careful about defining hits and niches in the Internet era. In a brick-and-mortar world, where product variety is relatively stable and all products are consumed at some rate, hits and niches are typically defined in absolute terms (e.g., the top 10, the bottom 100 movies). However, product variety has been skyrocketing in the Internet age and therefore more and more products can be left unnoticed by consumers, or are being discovered very slowly, even though the customer base is also expanding. To evaluate the consumer propensity to discover niches and to separate this effect from the entirely different effect of increasing product variety on the Internet, we suggest that product popularity should be measured in relative terms, thus dynamically adjusting for the "active" product variety at that point of time. By doing this, we bring the distribution of demand to a common scale and we analyze how it changes over time.

We use a novel data set that contains over 100 million online movie ratings from 2000 to 2005 and we supplement this data by web-crawling movie characteristics data from Yahoo Movies. In this large data set we find that, when the popularity of a movie is defined dynamically, i.e., in relative terms, there is essentially no evidence of the Long Tail effect: the demand for hits increases over time, while the demand for niches decreases. Additionally, even in the absolute ranking definition, the Long Tail effect is only partially present: the demand for niches decreases over time although the demand for hits decreases too.

In order to gain insights into these findings, we further examine changes in the demand distribution at the consumer level. Once again, we find that consumers over time indeed watch more niche movies in absolute terms, but we also discover that the rate at which consumers shift demand from the hits to the niches is considerably lower than the growth rate of product variety. Therefore, if we normalize product ratings for currently active product variety and measure popularity in relative terms, we find that consumers tend to watch more and more hits over time.

Figure 4 visually illustrates the comparison between the absolute and relative ratings of the movies that consumers watch. In Figure 4 (left, bottom line), we observe the median movie that the average consumer watched over time, which has a linear upward trend, indicating that consumers increasingly discover niches. In particular, at the beginning of our study, the median movie rated by an average consumer was ranked slightly above 150, while at the end of the study the median movie ranking had increased about seven times to over 1,150. However, Figure 4 (left, top line) also indicates that product variety increased even faster, which creates an impression of the lengthening tail of demand distribution: there are more and more obscure movies over time. However, once we bring distribution to the common scale through dividing by current product variety, the long tail disappears. Not surprisingly, Figure 4 (right) shows that, when we look at the median rating in relative terms, the average consumer gravitates more and more toward hits. In fact, at the beginning of our study the average consumer watched, on average, movies in the 11th percentile of product variety while at the end of the study the average consumer watched movies in the 5th percentile. Hence, we conclude that although consumers do venture into niches, new movies appear more quickly than people can actually discover them.

Figure 4: Average Median and Relative Average Median Popularity

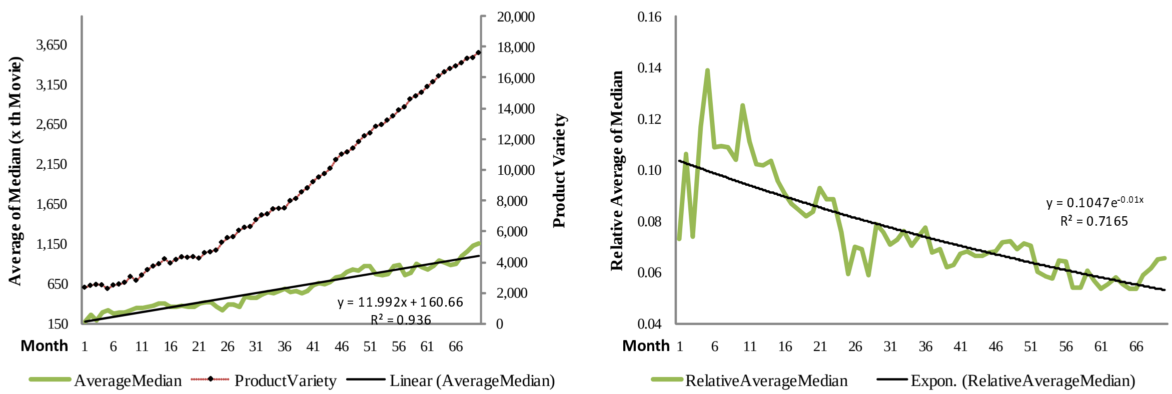

We make a number of additional observations based on our user-level analysis. We find evidence that the consumers who give high average ratings tend to watch more popular movies. Hence, we do not find any evidence that niche products satisfy consumer tastes better and better over time, which is suggested by Anderson (2004; 2006). Furthermore, we find that the consumers who watch the niches tend to be heavy users, constituting only a small part of the entire user base. Light users, however, tend to focus on the popular items and since most users are in this category, hits continue to drive the market.

Our findings have a number of managerial implications as they shed new light on the controversy surrounding the Long Tail effect. First, the promise of the Long Tail effect became a basis for many new business models and business ideas (see, e.g., http://en.wikipedia.org/wiki/The_Long_Tail). Our findings suggest that caution needs to be used when assessing the potential benefits of focusing a business on supplying niche products. While it may be true that niche products are much more profitable for companies (e.g., Anderson 2006 rightfully suggests that niche movies cost a fraction of hit movies to make), this argument does not account for the fact that for each niche product that consumers demand, there might be several that are never discovered, thus potentially adding to the costs but not to the revenues. Further, a large number of products might take a while to be discovered. This finding seems to suggest that much more attention needs to be paid to recommendation systems, review forums and other means of aiding product discovery. Although Netflix employs what is widely considered to be a sophisticated recommendation system, even this system does not allow numerous consumers to discover titles as fast as they appear. This raises an important issue of carefully forecasting how long will it take for a given title, once it is added to the inventory, to begin accumulating demand. More improvements to the recommendation systems, such through the Netflix Prize and the algorithm proposed by Park and Tuzhilin (2008) should be implemented.

Insights from our consumer-level analysis suggest that consumers are generally much more satisfied by hit products than by niches. This is an important consideration: while Netflix currently achieves extremely high ratings for customer satisfaction (see, e.g., http://en.wikipedia.org/wiki/The_Long_Tail), we do not find any evidence to suggest that customers watching obscure titles find them more satisfactory than other movies. We can speculate that many consumers over time will learn that niches are called niches for a reason and might stop watching them altogether. Our other observation that heavy users tend to watch more obscure movies suggests that the presence of the Long Tail effect might be moderated by the frequency of service. In the case of Netflix, it is physically impossible to rent more than a few DVDs per month (due to the time that the mailing process takes). However, Netflix and other companies (such as Amazon) have started allowing customers to watch DVDs on their computers at home right away, which may increase the number of heavy users who venture into niches. In this case, one will have to re-examine the existence of the Long Tail effect.

Of course, the production and inventory costs of media and entertainment products are steadily decreasing, particularly for some purely digital products. Nevertheless, for more traditional physical products, irrational expansion into niche products can increase operational difficulties, such as maintaining the level of service (Fisher et al., 1994; Randall and Ulrich, 2001). Some companies therefore can benefit by focusing on top-selling products, with the caveat that there are many more of these hits on the Internet.

It is important to remember the limitations of our findings. First, our study does not directly compare brick-and-mortar and Internet companies and therefore we are unable to comment on this aspect of the Long Tail effect. Rather, our findings need to be interpreted as a temporal study in the Internet environment only. Further, our findings are restricted to the ratings data, which are different from rental data. It is not entirely clear to us that ratings data are better or worse than rental data: consumers often rate movies that they watched elsewhere, providing a richer picture of demand for movies, but on the other hand many consumers probably do not rate the movies that they watched. An interesting future venue for research would be to compare ratings data with time-stamped rental data, and further to compare data for both online and brick-and-mortar stores. In addition, the data set with the time-stamped transaction level information may be potentially used to study consumer purchase patterns and their life-time value. Further research opportunities also include linking recommendation system metrics, such as product ratings, with operations management and marketing strategies (see Netessine et al. 2006 for some initial work in this direction). Finally, empirically distilling the effect of recommendation systems on demand concentration to verify results of Fleder and Hosanagar (2008) is a fruitful direction.

Acknowledgment

The authors gratefully acknowledge financial support from the Mack Center of Technological Innovation at the Wharton School. We would like to thank Chris Anderson, editor-in-chief of *Wired Magazine* and author of *The Long Tail: Why the Future of Business Is Selling Less of More,* for his encouraging comments and constructive advice. We are also indebted to Chris Dellarocas, Josh Eliashberg, Pete Fader, Dan Fleder, Karan Girotra, Andres Hervas-Drane, Kartik Hosanagar, Dylan Small, Christian Terwiesch, Matthew White, and Juanjuan Zhang for their helpful feedback.

We are also grateful to Netflix for providing free public access to its data.

References

Anderson, C. 2004. The long tail. *Wired Magazine* **12**(10) 170–177.

Anderson, C. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. New York: Hyperion.

Baumol, William, Edward Ide. 1956. Variety in retailing. *Management Science* **3**(1) 93–101.

Bockstedt, Jesse, Kim Huat Goh. 2008. Unbundling information goods: An empirical analysis of consumer created custom cds. INFORMS Annual Conference Presentation.

Brynjolfsson, E., Y. Hu, M. D. Smith. 2006. From niches to riches: The anatomy of the long tail. *Sloan Management Review* **47**(4) 67–71.

Brynjolfsson, E., Y. J. Hu, M. D. Smith. 2003. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science* **49**(11) 1580–1596.

Brynjolfsson, E., Y. U. J. Hu, D. Simester. 2007. Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. Working Paper.

Cachon, G. P., C. Terwiesch, Y. Xu. 2008. On the effects of consumer search and firm entry in a multiproduct competitive market. *Marketing Science* **27**(3) 461–473.

Chatfield, Christopher, Gerald Goodhardt. 1975. Results concerning brand choice. *Journal of Marketing Research* **12**(1) 110–113.

Chen, Pei-Yu, Shin-yi Wu, Jungsun Yoon. 2004. The impact of online recommndations and consumer feedback on sales. International Conference on Information Systems (ICIS).

Chevalier, Judith, Austan Goolsbee. 2003. Measuring prices and price competition online: Amazon and Barnes and Noble. NBER Working Paper.

Clemons, E. K., G. G. Gao, L. M. Hitt. 2006. When online reviews meet hyperdifferentiation: A study of the craft beer industry. *Journal of Management Information Systems* **23**(2) 149–171.

De Vany, A., W.D. Walls. 1996. Bose-Einstein dynamics and adaptive contracting in the motion picture industry. *The Economic Journal* **106**(439) 1493–1514.

Dellarocas, C. 2003. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science* **49**(10) 1407–1424.

Dellarocas, C., C.A. Wood. 2008. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Science* **54**(10) 460–476.

Dellarocas, Chrysanthos, Ritu Narayan. 2007. Tall heads vs. long tails: Do consumer reviews increase the informational inequality between hit and niche products? University of Maryland Working Paper.

Dewan, Sanjeev, Jui Ramaprasad. 2008. Consumer blogging and music sampling: Long tail effects. UC Irvine Working Paper.

Elberse, A. 2008. Should you invest in the long tail? *Harvard Business Review* (July-August) 88–96.

Elberse, A., F. Oberholzer-Gee. 2008. Superstars and underdogs: An examination of the long tail phenomenon in video sales. Working Paper.

Eliashberg, J., A. Elberse, M. Leenders. 2006. The motion picture industry: Critical issues in practice, current research, and new research directions. *Marketing Science* **25**(6) 638–661.

Farquhar, Peter, Vithala Rao. 1976. A balance model for evaluating subsets of multiattributed items. *Management Science* **22**(5) 528–539.

Fisher, Marshall, Janice Hammond, Walter Obermeyer, Ananth Raman. 1994. Making supply meet demand in an uncertain world. *Harvard Business Review* (May-June) 83–93.

Fleder, D., K. Hosanagar. 2008. Blockbuster culture's next rise and fall: The impact of recommender systems on sales diversity. *Management Science* (Forthcoming).

Flynn, Laurie. 2006. Like this? you'll hate that (not all web recommendations are welcome). *The New York Times* January 23rd.

Frank, R., P. Cook. 1995. *The Winner-Take-All Society: Why the Few at the Top Get So Much More Than the Rest of US*. New York: Penguin.

Ghose, Anindya, Bin Gu. 2006. Search costs, demand structure adn long tail in electronic markets: Theory and evidence. Working Paper.

Gomes, L. 2006. It may be a long time before the long tail is wagging the web. *The Wall Street Journal* July 26th.

Goodhardt, G. J., A. S. C. Ehrenberg, C. Chatfield. 1984. The dirichlet: A comprehensive model of buying behaviour. *Journal of the Royal Statistical Society. Series A (General)* **147**(5) 621–655.

Gu, Bin, H. Michelle Chen, Prabhudev Konana. 2008. Measuring product competition in online retailers from revealed preferences of online recommendation. University of Texas at Austin Working Paper.

Hausman, JA. 1978. Specification tests in econometrics. *Econometrica: Journal of the Econometric Society* **46**(6) 1251–1271.

Hervas-Drane, Andres. 2009. Word of mouth and taste matching: A theory of the long tail. Working Paper.

Hu, N., P.A. Pavlou, J. Zhang. 2007. Why do online product reviews have a j-shaped distribution? overcoming biases in online word-of-mouth communication. Working Paper.

Kekre, S., K. Srinivasan. 1990. Broader product line: A necessity to achieve success? *Management Science* **36**(10) 1216–1231.

McAlister, Leigh. 1982. A dynamic satiation model of variety seeking. *Journal of Consumer Research* **9**(2) 141–150.

McPhee, William N. 1963. *Formal Theories of Mass Behavior*. Glencoe, NY: Free Press.

Netessine, Serguei, Serguei Savin, Wenqiang Xiao. 2006. Revenue management through dynamic cross selling in e-commerce retailing. *Operations Research* **54**(5) 893–913.

Oestreicher-Singer, Gal, Arun Sundararajan. 2009. Recommendation networks and the long tail of electronic commerce. Working Paper.

Orlowski, Andrew. 2008. Chopping the long tail down to size. *The Register* 7 Nov.

Padmanabhan, Balaji, Alexander Tuzhilin. 2003. On the use of optimization for data mining: theoretical interactions and eCRM opportunities. *Management Science* **47**(10) 1327–1343.

Park, Yoon-Joo, Alexander Tuzhilin. 2008. The long tail of recommender systems and how to leverage it. ACM Conference on Recommender Systems.

Pessemier, Edgar. 1979. Stochastic properties of changing preferences. *The American Economic Review* **68**(2) 380–385.

Randall, T., S. Netessine, N. Rudi. 2006. An empirical examination of the decision to invest in fulfillment capabilities: A study of internet retailers. *Management Science* **52**(4) 567–580.

Randall, Taylor, Karl Ulrich. 2001. Product variety, supply chain structure, and firm performance: An empirical examination of the U.S. bicycle industry. *Management Science* **47**(12) 1588–1604.

Rosen, Sherwin. 1981. The economics of superstars. *The American Economic Review* **71**(5) 845–858.

Schmittlein, David C., Lee G. Cooper, Donald G. Morrison. 1993. Truth in concentration in the land of (80/20) laws. *Marketing Science* **12**(2) 167–183.

Shih, Willy, Stephen Kaufman, David Spinola. 2007. Netflix. *Harvard Business School Case* (November) 9–607–138.

Smith, Michael, Erik Brynjolfsson, Mohammad Rahman. 2008. Profit in the long tail. INFORMS Annual Conference Presentation.

Tucker, Catherine, Juanjuan Zhang. 2009. How does popularity information affect choices? A field experiment. MIT Sloan Working Paper.

Varian, Hal. 2006. Why old media and Tom Cruise should worry about cheaper technology. *The New York Times* October 19th.