

Annual Review of Economics

Bestseller Lists and the Economics of Product Discovery

Alan T. Sorensen

Department of Economics, University of Wisconsin, Madison, Wisconsin 53711;
email: sorensen@ssc.wisc.edu



ANNUAL REVIEWS **Further**

Click [here](#) to view this article's
online features:

- Download figures as PPT slides
- Navigate linked references
- Download citations
- Explore related articles
- Search keywords

Annu. Rev. Econ. 2017. 9:87–101

First published as a Review in Advance on
April 7, 2017

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-080614-115708>

Copyright © 2017 by Annual Reviews.
All rights reserved

JEL codes: D22, D83, L15, M37

Keywords

bestseller lists, product discovery, observational learning, long tail

Abstract

Innovations in information technology have increased the prevalence of markets with large numbers of products. Each week, for example, an average of over 800 books are published, and an average of over 1,100 iOS apps are released in Apple's App Store. This review summarizes existing research about how consumers learn about new products in such markets, focusing on mechanisms like bestseller lists and user-generated product reviews. In addition to reviewing research findings about how these mechanisms directly influence sales, this article also discusses how these mechanisms affect broader market outcomes, such as the shape of the success distribution and the likelihood that good products are discovered.

1. INTRODUCTION

Conventional models of demand in differentiated product markets assume that consumers are fully informed: They know about all of the available products in a market and all of those products' relevant characteristics, and they choose the product they like best. But in many markets, this full-information assumption is unrealistic due to the volume of new products that constantly flow into the market. Some notable examples of these markets are listed in **Table 1**. It is difficult to imagine that a consumer shopping for books in 2015 was fully informed about all of the 29 million available titles, or that a consumer shopping for music was fully aware of the 4 million available albums. Of course, markets with large numbers of products are not new—even in the 1950s, the flow of new books into the market was over 10,000 per year. But the prevalence of such markets is growing, largely due to innovations in information technology. Most notably, digital technologies have reduced the costs of producing and distributing digitizable goods such as books and music, leading to a dramatic increase in the rate of new product introductions in those industries,¹ and have also created entirely new product categories, such as smartphone apps, where the production and distribution technologies naturally allow for an enormous volume of products. Even in more conventional markets, the Internet has essentially created an unlimited supply of virtual shelf space, allowing sellers to offer large numbers of products, including niche products.

Markets like the ones listed in **Table 1**, where consumers cannot plausibly be fully informed about all products, raise important economic questions about the process of product discovery. In these types of markets, what mechanisms emerge to facilitate product discovery, and how do these mechanisms affect market outcomes? Are these mechanisms efficient—i.e., do consumers always find the products they like best? If not, how large are the implied welfare losses?

This article offers a brief summary of existing economic research into these questions. Special attention is given to positive-feedback market mechanisms, such as bestseller lists, which can cause a product's success (or lack thereof) to be self-reinforcing. I begin, in Section 2, by reviewing studies that have measured the sales impact of three different mechanisms that affect product discovery: bestseller lists, published reviews and endorsements, and information spillovers. These studies

Table 1 Examples of markets with many products

Market	New products introduced in 2014	Total products available in 2014
Books	300,000	29,000,000
Music albums	75,000	4,000,000
Movies	700	325,000
PC video games	2,700	6,000
iOS apps	400,000	2,000,000
Android apps	500,000	2,400,000
Restaurants in Manhattan	1,200	10,000

All numbers are approximate. Data taken from the following sources: books, Bowker and Amazon.com; music, Nielsen industry reports and Amazon.com; movies, Motion Picture Association of America and Internet Movie Database; video games, Wikipedia and Amazon.com; iOS apps, Apple's App Store and appfigures.com; Android apps, appbrain.com and Google Play store; restaurants, New York City Department of Health Inspections database.

¹For example, Waldfogel (2011) discusses how piracy and production cost reductions have affected the rate of new product introductions in the music industry.

generally find that the information conveyed by these mechanisms has a meaningful effect on sales and that this most likely results from consumers learning about products of which they were previously unaware. In Section 3, I discuss the broader implications of these findings, including the effects of incomplete product awareness on market-level outcomes. The main finding of the research to date is that, in markets like the ones listed in **Table 1**, the distribution of success is much more concentrated—i.e., success is more concentrated among a small number of popular products—than it would be if consumers were fully informed about all available products, and the welfare losses associated with undiscovered products are substantial.

2. MECHANISMS THAT AFFECT PRODUCT DISCOVERY

In general, consumers have many ways of learning about products. This section reviews empirical evidence on three mechanisms that seem especially influential in markets with large numbers of products: bestseller lists, recommendations and reviews, and information spillovers (i.e., learning about product A from news about a related product B).

Before proceeding, it is worth mentioning two other topics that, although relevant, are excluded from the present discussion. First, paid advertising obviously plays an important role in product discovery. Especially for newly introduced products, firms' decisions about how (and how much) to advertise those products directly influence consumers' awareness. Economic theorists have studied the welfare implications of informative advertising (see, e.g., Butters 1977, Grossman & Shapiro 1984), and empiricists have studied the degree to which advertising determines consumers' consideration sets (see, e.g., Goeree 2008). Several interesting and important questions remain about advertising's role in product discovery, but I exclude advertising from the present discussion simply because it is a broad topic that has been thoroughly reviewed elsewhere (see, e.g., Bagwell 2007, Evans 2009).

Second, it is natural to think of search models as relevant to any discussion of product discovery because uninformed consumers can obtain more information through costly search. However, I do not review the literature on search here, in part because (like advertising) it is a large literature that has been reviewed elsewhere, and in part because most studies of search in product markets have focused on price search, which is a narrower form of learning than we are concerned with here.

2.1. Bestseller Lists and Popularity Rankings

Consumers typically spend time and resources learning about products before deciding which one(s) to purchase. If the number of available products in a market is very large, it will be prohibitively costly to learn about all products, so demand depends not only on consumers' preferences but also on which products they choose to learn about. The set of products a consumer learns and knows about can be influenced by the choices of other consumers because in most markets—including those listed in **Table 1**—consumers can readily observe information about other consumers' purchases. This form of observational learning has been recognized and thoroughly explored in the economic theory literature. Early studies by Banerjee (1992) and Bikhchandani et al. (1992) show with simple models that individuals may rationally choose to ignore their own private information to follow a herd, and subsequent studies, like that of Smith & Sørensen (2000), extend the results to more general environments with preference heterogeneity. Bikhchandani et al. (1998) and Mobius & Rosenblat (2014) provide reviews of the theoretical literature on observational learning; my purpose in this article is to review empirical studies that have delivered evidence of the effects of observational learning in real-world markets.

Sorensen (2007) was the first study to directly measure the impact of published sales rankings, examining the impact of the *New York Times* (NYT) bestseller list on sales of hardcover fiction titles. As noted above, there are clear theoretical reasons to expect bestseller lists to directly influence consumer behavior in addition to reflecting that behavior. But measuring the causal impact of a bestseller list is a tricky empirical problem: The set of products that receive the treatment of being listed as bestsellers is clearly not random. Sorensen's study exploits two features of the construction of the NYT list that yield quasirandom variation. First, due to sampling error, the NYT list sometimes mistakenly included or excluded books, especially at the bottom of the list. Second, the NYT list was published with a 3-week lag, so books that attained bestseller status initially did not appear on the list. Together, these features of the NYT list construction led to many instances where books were excluded from the bestseller list even though their sales were as high as or higher than other books that were on the list, enabling a cleaner comparison of the sales trajectories of listed versus nonlisted books. Using data from Nielsen BookScan on a sample of 1,200 hardcover fiction titles, the study uses these sources of quasirandom variation to show that appearing on the NYT list causes a small but statistically significant increase in a book's sales. More importantly, the study explores the potential reasons for this effect and finds that patterns in the data are more consistent with an information-based explanation than a promotion-based explanation. Whereas retailers' promotions persist for the duration of a book's term on the list, the estimates suggest that the impact of appearing on the list is transitory, with the bulk of the effect realized in the first week. Moreover, the impact on sales is most pronounced among relatively unknown authors (new authors in particular), a pattern that favors information over promotion as an explanation for the effect.

Due to the challenges of identifying the effects of bestseller lists from observational data, most other studies have relied on experimental data. Cai et al. (2009) ran a clever experiment in a restaurant in Beijing to distinguish between observational learning effects and saliency effects. Diners at the restaurant—the experimental subjects—were exposed to one of three information conditions: a control condition, in which they were given the ordinary menu (which had 60 hot dishes to choose from); a ranking treatment condition, in which they were shown the top five most popular dishes from the previous week, in order; and a saliency condition, in which they were given a plaque with the names of five sample dishes. They find that sales of the top five dishes increased by 13–20% for diners given the ranking treatment condition, but no such increase occurred for diners given the saliency condition. They also find that the effects were somewhat larger for infrequent diners, which is consistent with the predictions of observational learning models because less informed individuals should be more responsive to the information revealed by other individuals' choices.

Tucker & Zhang (2011) describe another interesting study based on experimental data. They make the point that the effects of observational learning may be different for broad-appeal versus narrow-appeal products. Although it is clear from both theoretical and empirical studies that reporting a product's high popularity can increase its sales, most systems that report popularity do not account for differences in the sizes of products' markets. Some products may gain popularity simply due to their broad appeal—i.e., because they serve a large population of consumers. If a broad-appeal and a narrow-appeal product achieve the same level of popularity, rational consumers should draw stronger positive inferences about the quality of the narrow-appeal product than that of the broad-appeal product. The authors examine this hypothesis using a field experiment run by a website that serves as a directory of wedding services. The site changed its listings from conventional, alphabetically ordered listings to a more modern, bestseller-list style. In one randomly chosen category, vendors were listed in descending order of popularity, with previous clicks prominently shown for each vendor. In a control category, vendors were listed alphabetically,

and in a third category, vendors were sorted by popularity but without showing the number of previous clicks. The authors analyze the clicks in each category separately for vendors located in high-population towns versus low-population towns, which they argue is a proxy for broad appeal versus narrow appeal, respectively. They show that popular vendors' clicks were significantly increased in the treatment group and that this effect was much more dramatic for narrow-appeal vendors. This result—that narrow-appeal or niche products can benefit disproportionately from the display of popularity rankings—is similar in a way to Sorensen's (2007) finding that the sales increase resulting from an appearance on the NYT bestseller list is larger for debut authors, and also to Cai et al.'s (2009) result that displaying the most popular menu items in a restaurant is more influential for infrequent diners than for frequent diners. In each case, the pattern of results is consistent with rational observational learning, where the learning effects are largest when the consumer's own information is weakest (either because the consumer himself has limited knowledge or because the product in question is new and/or relatively unknown). However, in the case described by Tucker & Zhang, the mechanism is slightly different: Narrow-appeal products get larger sales bumps from popularity because consumers recognize that these products had to clear a higher hurdle to become popular.

Given the ubiquity of bestseller lists and popularity rankings and their evident impact on sales, sellers have strong incentives to get their products listed as bestsellers. A recent study by Li et al. (2016) explores these incentives in the market for iOS apps. As shown above in **Table 1**, this market has an enormous volume of products, prompting the authors' motivating question: How does an app get noticed out of the millions of apps available? They report that app developers' most significant costs are marketing costs and that most of these costs are incurred at or near launch. Appearing on the iTunes App Store's Top 300 most-downloaded chart is considered vital for the success of an app, but making it to the list is difficult: In a separate study, Davis et al. (2014) report that of 328,428 apps that were released in the iTunes App Store between September 2010 and August 2011, only 3,431 (1.04%) ever appeared in the Top 300. Interestingly, in this market developers can literally buy their way onto the chart by purchasing downloads—a process facilitated by companies that give incentives to otherwise uninterested consumers to download and install an app, for example, by giving them in-game bonuses or credits that can be converted to gift cards or cash. Using data from a download-buying platform, Li et al. estimate the effectiveness of buying downloads and the implied economic returns, finding that \$100 spent on buying downloads can improve an app's ranking by roughly 7% and that it would cost roughly \$17,000 for a typical app to buy a place in the top 10 for a day.

2.2. Reviews and Recommendations

The literature on observational learning in product markets mostly focuses on indirect learning, where consumers draw inferences about product quality from other consumers' choices. In many markets, consumers can learn from direct communication, in the sense that individuals tell each other directly about their opinions and experiences with a product. User-provided product reviews, which have become a standard feature of online shopping, are examples of this kind of direct communication. Chevalier & Mayzlin (2006) were the first to show direct empirical evidence that such reviews affect consumer behavior. They analyze book reviews posted on the two largest online booksellers, Amazon.com and bn.com (Barnes & Noble), and use a difference-in-differences approach to measure the effects of those reviews. In essence, their strategy was to calculate the changes in a book's sales rank from one period to the next on Amazon.com and bn.com and examine whether differences in those changes were related to differences in the reviews that had been posted during that period. As Chevalier & Mayzlin (2006, p. 345) more colorfully describe

it, “If a cranky consumer posts a negative review of a book on bn.com but not on Amazon.com, would the sales of that book at bn.com fall relative to the sales of that book at Amazon.com?” They find that the answer is yes: Positive reviews lead to increases in sales, and negative reviews lead to decreases in sales, with the impact of negative reviews being somewhat larger in magnitude.

Subsequent studies have examined the influence of reviews in other settings, mostly finding that reviews affect consumers in the ways we would naturally expect. For example, Luca (2011) finds that a one-star increase in a restaurant’s rating on Yelp.com increases its revenue by 5–9% and that this effect is only present for independent restaurants—changes in Yelp ratings have no significant effect on chain restaurants. A similar result is reported by Newberry & Zhou (2016), who find that reviews are less important for sellers and products that are already well-known. Using data from Alibaba, China’s largest online business-to-consumer marketplace, they show that sales are sensitive to posted ratings from previous consumers and that this sensitivity is 12 to 16 times greater for local sellers—i.e., sellers that do not have a national offline presence and reputation—than for national sellers.

Celebrity endorsements can be especially powerful product recommendations. In Sorensen’s (2007) study of the NYT bestseller list, he notes that the positive impact of appearing on the bestseller list appeared to be much smaller than the impact of a recommendation by celebrity talk-show host Oprah Winfrey. Garthwaite (2014) examines the effect of Oprah’s book club recommendations in more depth, using a larger sample of her recommended books, and finds the effect of these recommendations to be enormous: When Oprah endorsed a book, sales of that book increased by an average of 6,000% in the following week. This sales boost diminished over time, but remained positive for more than 6 months. Some of this effect could have been less about learning than about coordination: Because people generally enjoy discussing the books they read, the Oprah announcements could have enabled readers to coordinate on reading a book at the same time. However, given the magnitude of the effect, it also seems likely that many purchases were made by consumers who otherwise would not have known about the book Oprah recommended.

Internet technologies have amplified the importance of product reviews and recommendations both by making it less costly to disseminate information and by making it easier to collect information. Prior to the Internet, product ratings were published in magazines like *Consumer Reports*, covering a narrow range of products and reaching a limited audience. Today, by contrast, user-provided reviews and ratings can be cheaply organized and published, so it is rare to find products that do not have ratings and reviews posted online. Waldfogel (2015) considers how this affects product discovery in the movie industry. He points out that the Internet has significantly increased both the number of published reviews per movie and the number of movies that have any reviews at all. He postulates that the second effect, the increase in coverage of movie reviews, makes it easier for consumers to discover a wider range of movies.

User-provided product reviews are easy to collect but difficult to monitor. To the extent that reviews influence consumer behavior, firms have incentives to submit fake reviews—and a recent study by Mayzlin et al. (2014) provides evidence of such manipulation in the hotel industry. They examine the difference in review distributions between Expedia and TripAdvisor, which have different reviewer verification policies, and across different competitive conditions. They find that hotels with next-door neighbors have more negative reviews on TripAdvisor (where it is easier to submit fake reviews), and also that independent hotels with small owners have more positive reviews.

Online sellers’ ability to gather large quantities of user-generated product reviews—either explicitly through users’ direct ratings of products or implicitly through their choices of which products to purchase—makes it possible for them to construct automated big data tools to help

consumers find products. In fact, automated recommenders or recommendation engines may be the most important modern mechanisms for facilitating product discovery, but their influence has received little attention in the economics literature. Most online shopping sites (e.g., Amazon.com) display some form of “if you like product A, we think you’ll like product B” recommendations, and some (e.g., Netflix) encourage users to submit their rankings of different products as a way of improving the site’s recommendations for new products. These recommendation systems generally take one of two approaches. Systems based on content filtering generate recommendations based on profiles of the users’ and/or products’ observed characteristics: for example, recommending an action movie to a teenage male or a book about parenting to someone who has young children. Systems based on collaborative filtering rely only on past user behavior, generating recommendations based on similarities between users or products implied by users’ transaction data: for example, recommending movie A because the user purchased movie B, and other users who purchased B also purchased A. The machine learning algorithms underlying these systems are powerful, fast, and surprisingly scalable, especially in the case of collaborative filtering. A fascinating question that has not yet been addressed by academic studies (but that one suspects has been explored internally by many of the large online sellers) is how much these recommendation systems increase consumers’ ability to find products they like—i.e., how much they increase the efficiency of the matching of consumers to products.

2.3. Information Spillovers

Hendricks & Sorensen (2009) attempt to empirically quantify consumers’ lack of information in the music industry, employing a novel strategy based on measuring the effects of new album releases on sales of previous albums by the same artist. Their study analyzes music album sales between 1993 and 2002, a time when consumers still learned about albums primarily through radio play and purchased them mainly at brick-and-mortar stores. Promotional activity and radio airplay surrounding a new album release typically boosted the sales of the artist’s past albums—an effect that the authors labeled the backward spillover. The authors measured the strength of the backward spillovers in the data and documented a few key patterns. First, the spillover effects tended to begin a few weeks prior to the new album’s release, which was consistent with the typical timing of airplay and promotional activity for new albums. Second, the effects were surprisingly persistent: Sales of the artist’s past albums remained inflated for several months after the release of the artist’s new album. Third, the effects were largest when the newly released album was a hit and especially large when the new release was a hit and the previous album was not. Finally, they show that the backward spillovers were meaningfully smaller in the artist’s home market (the city in which she began her career), even though sales were higher on average in the home market. Taken together, these patterns are strongly suggestive of an information-based interpretation of the spillover effects: namely, that the new album releases were causing some consumers to discover the artist and purchase the artist’s past albums.

Kumar et al. (2014) conduct a similar exercise in the movie industry. They examine the information shocks that result from the timing of films’ releases to different distribution channels: Typically, films are first released in theaters, then to DVD, and then to paid cable channels like HBO. The essential hypothesis of the study is that consumers may be unaware or ill-informed about many films at the time they are released to DVD because, whereas virtually all movies get released to DVD and paid cable, relatively few movies are heavily advertised during their theater runs. In this context, the backward spillover comes from paid-cable broadcasts: When a movie is shown on HBO, for example, the consumer might learn about it and like it enough to buy the DVD or tell her friends about it and increase the likelihood that they will purchase the DVD. The

authors use detailed sales data on 314 movies and show that DVD sales indeed increase when a movie is released to paid cable and that the effect is stronger for lesser-known movies.

Garthwaite’s (2014) study of the Oprah endorsement effect also finds evidence of significant information spillovers. After Oprah announced a pick for her book club, sales of other titles by the same author increased by 53%, and the effect persisted for several months. Taken together, the spillover effects described in these studies suggest that many products “lose” sales due to a lack of consumer information; in the next section, I discuss this consequence of incomplete learning in more depth.

3. CONSEQUENCES OF INCOMPLETE PRODUCT DISCOVERY

To the extent that bestseller lists and product reviews influence consumer learning, they can also influence overall market outcomes. To date, most research has focused on how learning mechanisms affect the efficiency of the match between consumers and products and shape the distribution of success—for example, by concentrating it more heavily among the few most-popular products. To fix ideas and clarify the issues I aim to discuss, I first briefly outline a simple framework for thinking about the inefficiencies associated with incomplete product discovery. I then describe how previous studies, including several of the ones mentioned in the previous section, have measured the broader market impacts of incomplete learning.

3.1. A Simple Framework

Consider a market with N available products, in which consumer i ’s utility from purchasing product n is $u_{in} = \theta_n + \epsilon_{in}$. Product quality that is common across consumers is captured by θ_n , whereas consumer i ’s idiosyncratic tastes are captured by ϵ_{in} , which we can also call the match quality. I begin with this utility specification because it is standard in empirical models of discrete-choice demand; the ϵ_{in} term is typically assumed to be a type I extreme-value random variable, so that the choice probabilities have the convenient logit form. In that case, if S_i represents the set of available products known to consumer i , the probability that i chooses n is

$$p_{in}(S_i) = \begin{cases} \frac{e^{\theta_n}}{\sum_{k \in S_i} e^{\theta_k}} & \text{if } n \in S_i \\ 0 & \text{if } n \notin S_i. \end{cases} \quad 1.$$

In this context, questions about product discovery are questions about which products are in S_i . In markets where N is large, we do not expect S_i to include all available products; if information in the market is such that consumers always learn about the products they like most, then it does not matter (from a welfare perspective) whether they know about all products. If, on the other hand, consumers do not always know about the products they would like most, then there are welfare losses stemming from the inefficient matching of consumers to products—and the magnitude of these losses depends on how incomplete is the awareness set S_i .

Even if consumers are not aware of all products, it is natural to expect higher-quality products to be more widely known. One way to parameterize this is to write the probability of product n ’s inclusion in S_i as

$$q_n = \frac{e^{a+b\theta_n}}{1 + e^{a+b\theta_n}}. \quad 2.$$

The parameter a determines the baseline awareness probability; as a goes to $-\infty$, the probability that a consumer knows about a low-quality product goes to zero; as a goes to ∞ , the awareness

probability goes to 1 for all products regardless of quality. The parameter b represents the degree to which awareness is a function of product quality: The larger is b , the more awareness is driven by product quality.

To calibrate the implied welfare losses from incomplete product awareness within this simple setup, I ran simulations of a market with 1,000 products with qualities (θ 's) drawn from a distribution calibrated to match the distribution of market shares described by Chevalier & Goolsbee's (2003) study of book sales on Amazon.com. If the a and b parameters are chosen so that the highest-quality product has a 95% chance of being known and the median-quality product has a 10% chance of being known, then roughly 75% of consumers are unaware of their highest-utility product. The welfare loss from this inefficiency is equivalent to the loss that would result from removing over half of the products from the market. These numbers are striking, and although they should be taken with a grain of salt—when each consumer gets many independent draws of the logit error, it is not surprising that idiosyncratic matches are important—they nevertheless highlight that welfare losses from incomplete product awareness can be large.

An alternative and more realistic formulation makes the awareness probability a function of popularity:

$$q_{nt} = \frac{e^{a+bs_{nt}}}{1 + e^{a+bs_{nt}}}, \quad 3.$$

where s_{nt} is product n 's market share among the first t consumers who purchase. This framework is based on the same idea as the influential Bass (1969) model, where the probability that a product is purchased is an increasing function of its cumulative past sales. As discussed by Bass, there are a variety of market forces, including word of mouth, that would make consumers more likely to buy popular products. In most models, higher-quality products are more likely to become popular, so in a sense, this setup is similar to one based on quality alone. The important difference is that outcomes in this setup are path dependent: The set of products that gain high market shares depends on the idiosyncratic preferences of the early consumers in the market. If consumer arrival is random, then products' eventual market shares are also random, and this model admits the possibility that high-quality products remain unnoticed. In simulations of markets with this setup, we find outcomes for consumers to be similar to the setup in which awareness is directly linked to product quality; the main difference is that products' market shares were more volatile.

3.2. Sales Distributions and Lost Sales

When product awareness is driven by observational learning processes, early success can be self-reinforcing, so that sales become even more concentrated among the most popular products. Moreover, the path dependence of the observational learning process makes product-specific outcomes stochastic, and the highest-quality products are not guaranteed to succeed. Salganik et al. (2006) illustrate these effects with a clever experiment in which thousands of subjects were recruited to participate in artificial online music markets. Participants arrived sequentially and were presented with a list of 48 songs, which they could listen to, rate, and then download (for free) if they so chose. In real time, each participant was randomly assigned to one of nine worlds. In the treatment worlds, of which there were eight, participants were shown information about the downloads of previous participants in their world: Song listings included the total number of previous downloads, and the song list was sorted by download rank. In the control world, the songs were shown in a random order, with no information about previous participants' listens or downloads. Importantly, the eight treatment worlds operated independently of one another, so that the researchers could observe eight separate realizations of the stochastic learning process for

each song. The authors show that outcomes in the treatment versus control worlds were different in two important ways. First, the distribution of success (as measured by number of downloads) was more unequal in the treatment worlds: Displaying the relative popularity of songs tended to increase downloads of the most popular songs at the expense of the less popular songs. Second, success was more unpredictable in the treatment worlds: For good songs in particular—a good song being one with relatively high downloads in the control world—there was a wide range of outcomes across the eight treatment worlds. As Salganik et al. (2006, p. 855) put it, in the treatment worlds, “the best songs never do very badly, and the worst songs never do extremely well, but almost any other result is possible.”

Hendricks et al. (2012) argue that, from a social welfare perspective, this unpredictability of outcomes is the most important aspect of observational learning. In markets with large numbers of products, where observational learning mechanisms appear to be most influential, the key question is whether observational learning leads consumers to ignore superior products or waste time and resources learning about inferior products. In other words, what is the likelihood of bad herds in these markets? The authors develop a variant of the herding model described by Smith & Sørensen (2000) in which a large number of consumers with heterogeneous preferences arrive sequentially and decide whether to buy a product. Consumers do not know the quality of the product, which is either high or low, or their idiosyncratic preferences for it. Each consumer observes the aggregate purchases of previous consumers and a private informative signal about her utility for the product and uses this information to decide whether to investigate the product and learn more about her utility—an action the authors refer to as search—before making a purchase decision. Consumers only purchase products that they search and like; however, search is costly, so consumers do not search products they believe they are unlikely to buy. The authors derive an interesting testable implication from this model, which is that the long-run outcomes for high-quality products should have a two-point support: Either the probabilities of search and purchase will converge to zero (a bad herd) or the market will converge to complete learning, in which case the probabilities of search and purchase will converge to positive constants. This prediction is tested using the data from the Salganik et al. (2006) experiment. The distribution of songs’ long-run listening rates (that is, the fraction of experimental subjects who chose to sample a song) across the eight treatment worlds does indeed exhibit bimodality. For example, the highest-quality song (as measured by its download rate in the control world) had long-run listening rates between 18–30% in five of the eight treatment worlds but, in two of the worlds, had listening rates of 2.5% and 3.5%, respectively. This is consistent with the idea that, although observational learning can facilitate convergence to the right market outcome, it can also cause the market to miss good products.

A common characteristic of markets with many products is that sales tend to be highly skewed: A large share of industry profit is claimed by a small number of very successful products. Although some of this skewness surely reflects differences in product quality, it is natural to ask how much of the skewness can be attributed to consumers’ lack of information. Hendricks & Sørensen (2009) use their estimates of the backward spillovers in the music industry, in combination with a simple model of consumer demand, to measure the extent to which consumers’ lack of product awareness skews the distribution of sales. Their model is similar to the Bass (1969) model alluded to above, positing that the probability that a consumer purchases an album is the product of two probabilities: the probability that she knows about the album and the probability that she likes the album enough to buy it (conditional on knowing about it). The strength of the backward spillovers in the data—and, more importantly, how the magnitudes of the spillovers depend on the relative success of the artist’s past versus new albums—enables the estimation of an artist discovery function characterizing the probability that a typical consumer would know about a given artist. Having estimated that function, the authors then simulate what the distribution of

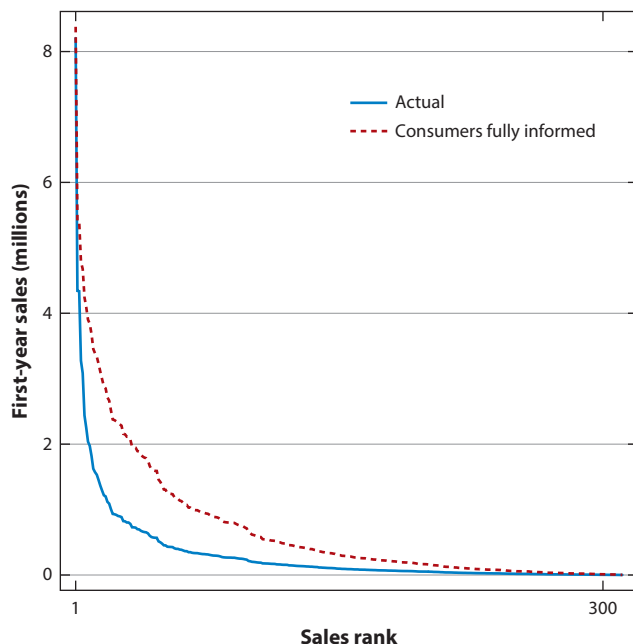


Figure 1

Sales by rank for top 300 music albums: actual sales (*solid line*) versus full-information (*dashed line*). Figure reproduced with permission from Hendricks & Sorensen (2009).

sales would have looked like if all artists were always discovered. The results from this simulation, shown graphically in **Figure 1**, imply that incomplete information has a substantial impact on the distribution of sales in the music industry. In particular, the estimates indicate that, whereas almost all consumers learn about an artist with a major hit, only 32% of consumers learn about an artist whose album reaches the median level of sales. In the data, sales of the top artist exceeded sales of the median artist by a factor of 90; the authors estimated that, if consumers had been fully informed about the median artist, then her sales would have been lower than the top artist's by only a factor of 30.

In the movie industry, the backward spillover effects found by Kumar et al. (2014) tend to reduce the overall skewness of sales. In the first few weeks that a cohort of movies becomes available on DVD, the distribution of sales across movies almost exactly matches the (very skewed) distribution from theater box office sales; however, after the movies are released to paid cable, this distribution becomes meaningfully less skewed as sales of lesser-known films increase proportionally more than sales of well-known films. The authors estimate a learning model very similar to the artist discovery model of Hendricks & Sorensen (2009), and find results that parallel what Hendricks & Sorensen find for the music industry. Consumer awareness is very high for blockbuster movies (over 99% for the biggest movie in their sample), but the awareness probability for movies in the bottom quartile of the sales distribution is only 57%. The authors estimate that the median movie's DVD sales would increase by 28% if consumers were fully informed.

Li et al. (2016) explore similar questions about market outcomes in their study of download buying in app markets. They outline a modified Bass model in which an app's diffusion does not begin until it becomes visible, which happens when it reaches a top sellers list, and developers can purchase downloads in order to gain this visibility. The authors describe two testable implications

of the model: (a) Diffusion patterns should be double-humped because the download buying period causes an accelerated early diffusion that combines both bought and organic (nonbought) downloads but then settles into a diffusion based only on organic downloads, and (b) rankings should become more persistent over time because bought downloads make early rankings noisy indicators of eventual success. Both of these patterns are shown to be true for apps listed in the iTunes App Store. The authors then use data on iTunes App Store rankings to estimate their model and simulate counterfactual diffusion paths that would result if all apps were visible from the start. They find that the speed of apps' diffusions, as measured by their half-lives (the time it takes to attain half of their eventual downloads), would accelerate by roughly 95% without the visibility problem. The implication is that, although the iTunes Top 300 list mitigates the visibility problem, there is still an enormous gap between the observed diffusion rates and the ones we would see if all apps were fully known to all consumers.

In some cases, firms have chosen pricing schemes to encourage experimentation and product discovery. Newberry (2016) studies product discovery in the (now defunct) Amie Street online music market, in which songs were priced as a function of their popularity. All songs posted on the site were initially free to download; after that, their prices increased with the number of downloads. Users shopping for songs on the site would see the price, which served as a signal of previous downloads, and also the number of times the song had been listened to. Intuitively, the site's pricing structure should have encouraged experimentation, and Newberry finds this to be the case. Using detailed data for a large number of songs, he estimates a model of users' listening and purchase choices, and then simulates the model to predict outcomes under alternative pricing schemes. His results indicate that fixed prices would slightly increase the rate at which good songs fail (bad herds) and would significantly reduce consumer surplus. He also finds that provision of information—i.e., the ability to sample songs prior to purchasing them—reduces the rate at which good songs fail and increases consumer surplus. He notes that the measured effects are not large because his estimates indicate considerable heterogeneity in consumers' preferences and a relatively low cost of experimentation. However, the qualitative implication of his study is an important one: Firms' pricing schemes can have a meaningful impact on the product discovery process.

3.3. Learning and the Viability of Niche Products

Anderson (2007) popularized the hypothesis that Internet technologies should lead to increased sales of niche products—products that occupy the thin right tail of the sales distribution when sales are plotted against sales ranks (as in **Figure 1**). Brynjolfsson et al. (2003) estimate that, in the year 2000, online sales of niche books—i.e., books that are not typically offered in brick-and-mortar bookstores—generated nearly \$1 billion in consumer surplus. In a later study (Brynjolfsson et al. 2010), the same authors compare the shape of the sales distribution on Amazon.com in 2008 to that of the same distribution in 2000. Between 2000 and 2008, online search tools for finding niche products—such as popularity ranks, recommendation engines, and pervasive product reviews—presumably improved. The authors show that the long tail grew longer over that period: Niche books accounted for a larger fraction of total sales in 2008 than in 2000, and the estimated consumer surplus generated by niche books increased by a factor of 5.

In principle, there are two reasons to expect Internet technologies to increase the importance of the long tail: On the supply side, the low costs of virtual shelf space make niche products easier to stock, and on the demand side, consumers' ability to browse and search online makes niche products easier to find. Brynjolfsson et al. (2011) show that the latter effect is important. Using data from a clothing retailer that sells both online and via mail catalogs, the authors show

that Internet sales were less concentrated than offline sales (Gini coefficients of 0.49 and 0.53, respectively) even though the two channels offered exactly the same products. A natural question is whether this difference was merely due to differences in the preferences of consumers who choose to shop online versus offline; online shoppers may choose to shop online precisely because they are looking for harder-to-find products. However, the authors show that the difference in sales concentration remains even after controlling for consumer characteristics. They also examine server logs to measure the type and extent of customers' searches and find that sales of niche products are higher among customers who conduct nondirected searches and use the website's recommendation system, which is consistent with the idea that these searches expose the consumer to information about products they otherwise would not have considered.

The long tail hypothesis is intuitively appealing, and evidence from the above-mentioned studies confirm its validity. However, other authors have pointed out that Internet technologies may also amplify the success of the most-popular products. Rosen (1981) argues that a market's tendency to generate superstars—or, more formally, to generate convex returns to product quality—can be explained simply by imperfect substitutability: Low-quality products are poor substitutes for high-quality products. [As Rosen (1981, p. 846) puts it, “hearing a succession of mediocre singers does not add up to a single outstanding performance.”] Moreover, this general market tendency is magnified when consumption value does not suffer from any congestion effects—that is, when one consumer's purchase of a product does not diminish other consumers' utility of consumption. These features are certainly present in mass media markets, and considered in combination with the prominence of bestseller lists, they suggest that Internet technologies may push demand more toward superstars than toward niche products. Weeds (2012) offers a theoretical framework in which a decline in the fixed costs of production can simultaneously generate increased returns to both superstars and niche products. Elberse & Oberholzer-Gee (2007) examine home video sales between 2000 and 2005 and find a clear superstar effect: Sales were more concentrated among the few most-popular titles in 2005 than they were in 2000. Tan & Netessine (2009) report similar patterns in the demand for movies on Netflix, arguing that the demand for hits had increased over time, whereas the demand for niche movies had fallen. Tan & Netessine (2009, p. 1) conclude that “new movie titles appear much faster than consumers discover them.”

4. CONCLUSION

Collectively, the research studies reviewed in this article provide ample evidence that consumers are often not fully informed about the products available in a market and that, consequently, they are responsive to information provided by bestseller lists, reviews, and recommendations. By themselves, such results are not surprising: In markets with many thousands of products, common sense says that consumers will not be fully aware of all the options. However, the magnitude of the information gap is important for understanding the implied efficiency gains that would result from mechanisms aimed at closing that gap.

From an economic perspective, the most interesting questions concern how consumers' lack of awareness affects market outcomes. If consumer awareness is driven by observational learning processes, as most of the above-cited studies indicate, then the success of a new product is path dependent and unpredictable. In particular, such environments allow the possibility of bad herds in which good products are never discovered. Existing studies suggest that bad herds do happen, and that the lost sales and consumer surplus resulting from consumers' lack of information are large. Policies and technologies that facilitate consumer learning or otherwise improve the matching of consumers to products could therefore generate substantial welfare gains.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

Thanks to Adam Smith (a graduate student at University of Wisconsin, not the deceased author of *The Wealth of Nations*) for help with this manuscript.

LITERATURE CITED

- Anderson C. 2007. *The Long Tail: How Endless Choice Is Creating Unlimited Demand*. New York: Random House
- Bagwell K. 2007. The economic analysis of advertising. In *Handbook of Industrial Organization*, Vol. 3, ed. M Armstrong, R Porter, pp. 1701–844. Amsterdam: Elsevier
- Banerjee AV. 1992. A simple model of herd behavior. *Q. J. Econ.* 107(3):797–817
- Bass FM. 1969. A new product growth for model consumer durables. *Manag. Sci.* 15(5):215–27
- Bikhchandani S, Hirshleifer D, Welch I. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* 100(5):992–1026
- Bikhchandani S, Hirshleifer D, Welch I. 1998. Learning from the behavior of others: conformity, fads, and informational cascades. *J. Econ. Perspect.* 12(3):151–70
- Brynjolfsson E, Hu JY, Simester D. 2011. Goodbye Pareto principle, hello long tail: the effect of search costs on the concentration of product sales. *Manag. Sci.* 57(8):1373–86
- Brynjolfsson E, Hu JY, Smith MD. 2003. Consumer surplus in the digital economy: estimating the value of increased product variety at online booksellers. *Manag. Sci.* 49(11):1580–96
- Brynjolfsson E, Hu JY, Smith MD. 2010. *The longer tail: the changing shape of Amazon's sales distribution curve*. SSRN Work. Pap. 1679991
- Butters GR. 1977. Equilibrium distributions of sales and advertising prices. *Rev. Econ. Stud.* 44(3):465–91
- Cai H, Chen Y, Fang H. 2009. Observational learning: evidence from a randomized natural field experiment. *Am. Econ. Rev.* 99(3):864–82
- Chevalier J, Goolsbee A. 2003. Measuring prices and price competition online: Amazon.com and BarnesandNoble.com. *Quant. Mark. Econ.* 1(2):203–22
- Chevalier JA, Mayzlin D. 2006. The effect of word of mouth on sales: online book reviews. *J. Mark. Res.* 43(3):345–54
- Davis JP, Muzzyrya Y, Yin P-L. 2014. *Experimentation strategies and entrepreneurial innovation: inherited market differences in the iPhone ecosystem*. Work. Pap., INSEAD, Fontainebleau, Fr.
- Elberse A, Oberholzer-Gee F. 2007. Superstars and underdogs: an examination of the long-tail phenomenon in video sales. *Mark. Sci. Inst. Rep.* 4:49–72
- Evans DS. 2009. The online advertising industry: economics, evolution, and privacy. *J. Econ. Perspect.* 23(3):37–60
- Garthwaite CL. 2014. Demand spillovers, combative advertising, and celebrity endorsements. *Am. Econ. J. Appl. Econ.* 6(2):76–104
- Goeree MS. 2008. Limited information and advertising in the US personal computer industry. *Econometrica* 76(5):1017–74
- Grossman GM, Shapiro C. 1984. Informative advertising with differentiated products. *Rev. Econ. Stud.* 51(1):63–81
- Hendricks K, Sorensen A. 2009. Information and the skewness of music sales. *J. Polit. Econ.* 117(2):324–69
- Hendricks K, Sorensen A, Wiseman T. 2012. Observational learning and demand for search goods. *Am. Econ. J. Microecon.* 4(1):1–31
- Kumar A, Smith MD, Telang R. 2014. Information discovery and the long tail of motion picture content. *MIS Q.* 38(4):1057–78

- Li X, Bresnahan T, Yin P-L. 2016. *Paying incumbents and customers to enter an industry: buying downloads*. Work. Pap., Stanford Univ., Stanford, CA
- Luca M. 2011. *Reviews, reputation, and revenue: the case of Yelp.com*. Work. Pap. 12-016, Harvard Bus. School, Cambridge, MA
- Mayzlin D, Dover Y, Chevalier J. 2014. Promotional reviews: an empirical investigation of online review manipulation. *Am. Econ. Rev.* 104(8):2421–55
- Mobius M, Rosenblat T. 2014. Social learning in economics. *Annu. Rev. Econ.* 6(1):827–47
- Newberry P. 2016. An empirical study of observational learning. *RAND J. Econ.* 47(2):394–432
- Newberry P, Zhou X. 2016. *Heterogeneous effects of online reputation for local and national retailers*. Work. Pap., Pennsylvania State Univ., State College
- Rosen S. 1981. The economics of superstars. *Am. Econ. Rev.* 71(5):845–58
- Salganik MJ, Dodds PS, Watts DJ. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311:854–56
- Smith L, Sørensen P. 2000. Pathological outcomes of observational learning. *Econometrica* 68(2):371–98
- Sorensen AT. 2007. Bestseller lists and product variety. *J. Ind. Econ.* 40(4):715–38
- Tan TF, Netessine S. 2009. *Is Tom Cruise threatened? Using Netflix prize data to examine the long tail of electronic commerce*. Work. Pap., Univ. Pennsylvania, Philadelphia
- Tucker C, Zhang J. 2011. How does popularity information affect choices? A field experiment. *Manag. Sci.* 57(5):828–42
- Waldfoegel J. 2011. *Bye, bye, Miss American Pie? The supply of new recorded music since Napster*. NBER Work. Pap. 16882
- Waldfoegel J. 2015. *Cinematic explosion: new products, unpredictability, and realized quality in the digital era*. Work. Pap., Univ. Minnesota, Minneapolis
- Weeds H. 2012. Superstars and the long tail: the impact of technology on market structure in media industries. *Inf. Econ. Policy* 24(1):60–68

Contents

Tony Atkinson on Poverty, Inequality, and Public Policy: The Work and Life of a Great Economist <i>Anthony Barnes Atkinson and Nicholas Stern</i>	1
Quantitative Spatial Economics <i>Stephen J. Redding and Esteban Rossi-Hansberg</i>	21
Trade and the Environment: New Methods, Measurements, and Results <i>Jevan Cherniwchan, Brian R. Copeland, and M. Scott Taylor</i>	59
Bestseller Lists and the Economics of Product Discovery <i>Alan T. Sorensen</i>	87
Set Identification, Moment Restrictions, and Inference <i>Christian Bontemps and Thierry Magnac</i>	103
Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure <i>Leonardo Bursztyn and Robert Jensen</i>	131
Quantile Regression: 40 Years On <i>Roger Koenker</i>	155
Globalization and Labor Market Dynamics <i>John McLaren</i>	177
High-Skilled Migration and Agglomeration <i>Sari Pekkala Kerr, William Kerr, Çağlar Özden, and Christopher Parsons</i>	201
Agricultural Insurance and Economic Development <i>Shawn A. Cole and Wentao Xiong</i>	235
Conflict and Development <i>Debraj Ray and Joan Esteban</i>	263
Quantitative Trade Models: Developments and Challenges <i>Timothy J. Kehoe, Pau S. Pujolàs, and Jack Rossbach</i>	295
The Economics of Nonmarital Childbearing and the Marriage Premium for Children <i>Melissa S. Kearney and Phillip B. Levine</i>	327

The Formation of Consumer Brand Preferences <i>Bart J. Bronnenberg and Jean-Pierre Dubé</i>	353
Health, Health Insurance, and Retirement: A Survey <i>Eric French and John Bailey Jones</i>	383
Large-Scale Global and Simultaneous Inference: Estimation and Testing in Very High Dimensions <i>T. Tony Cai and Wenguang Sun</i>	411
How Do Patents Affect Research Investments? <i>Heidi L. Williams</i>	441
Nonlinear Panel Data Methods for Dynamic Heterogeneous Agent Models <i>Manuel Arellano and Stéphane Bonhomme</i>	471
Mobile Money <i>Tavneet Suri</i>	497
Nonparametric Welfare Analysis <i>Jerry A. Hausman and Whitney K. Newey</i>	521
The History and Economics of Safe Assets <i>Gary Gorton</i>	547
Global Liquidity: A Selective Review <i>Benjamin H. Cohen, Dietrich Domanski, Ingo Fender, and Hyun Song Shin</i>	587

Indexes

Cumulative Index of Contributing Authors, Volumes 5–9	613
Cumulative Index of Article Titles, Volumes 5–9	616

Errata

An online log of corrections to *Annual Review of Economics* articles may be found at <http://www.annualreviews.org/errata/economics>