

## The Long Tail and Its Implications for Media Audience Measurement

Scott McDonald  
Condé Nast Publications

### LONG TAIL THEORY AND THE MEDIA

Chris Anderson coined the term "The Long Tail" to describe a family of business models emerging as a result of the revolutionary changes caused by digital technologies. The internet provides only one salient example: because of the internet, many businesses are no longer confined to selling to just their local markets, but instead can serve customers across much larger territories. The resulting expansion of markets improves the likelihood of selling low-occurrence niche items that would have trouble finding their customers in conventional, geographically-bounded markets. For internet-based businesses, locality no longer dictates the perimeter of market demand.

But it is not just the internet that accounts for the transformative role of digital technologies: Moore's Law plays an important role as well. Moore's Law describes the exponential growth of computing power, a phenomenon that causes the power of silicon chips to double about every two years. As a result of this rapid acceleration in the computational power of this fundamental building block of digital machines and digital media, businesses built on a digital foundation experience two wonderful benefits—declining cost of digital storage and declining cost of digital distribution.

Anderson argues that as a result of these three important aspects of the digital age—global-scale markets, falling costs of storage, and falling costs of distribution—the digitally-based businesses of today can now sell niche items and niche content profitably. And when consumers gain access to a greater range of choices, they gravitate toward exercising those choices, awarding fewer of their "votes" to the big hits and more of their "votes" to specialized niche choices. Anderson argues that people always wanted more choices, but their desires previously were obscured by distributional bottlenecks imposed by cost or locality. As a result, we erroneously inferred that they only wanted the "hits"—the lowest common denominators of consumer demand.

In this new digital era of 21st century markets, however, we can at last see the true shape of consumer demand, and that shape has a "long tail." Anderson is not speaking metaphorically here: he literally is referring to the shape of a specific family of statistical distributions known as power law distributions—distributions that have a short head (the few choices that are popular with many—the "hits") and a long tail (the many choices that appeal only to small groups of customers). By contrast, the limitations of locality, inventory costs, and distribution costs that characterized 20th century markets led marketers mistakenly to view consumer demand through the prism of the Normal Curve, with its focus only on the products with appeal to the "average" masses (see Figures 1 and 2).

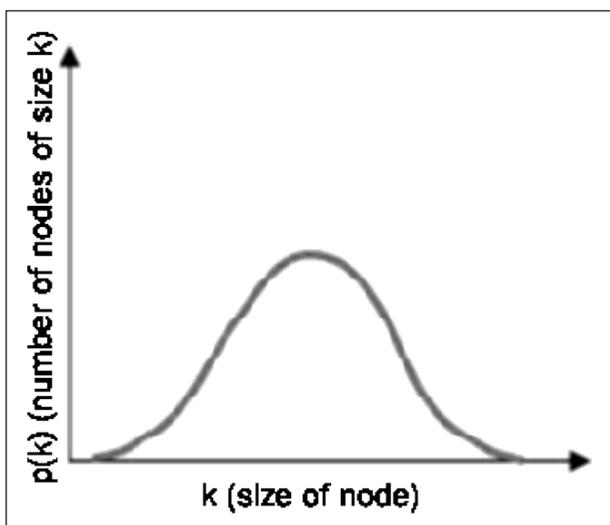


Figure 1: A Normal Distribution

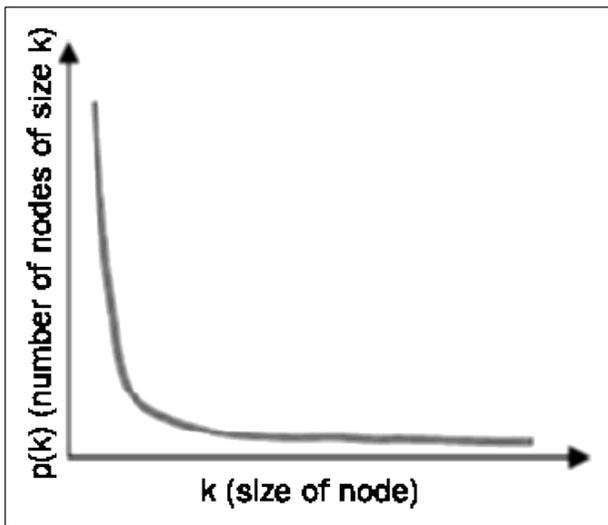


Figure 2: A Power Law Distribution

The long-tail phenomena described by Anderson have occupied statisticians for many years. Indeed, the term “Long Tail” refers to the shape of the frequency distribution described by the power law function—a family of distributions in which the “head” is very high and short (the hits) and the “tail” is very long (the niches). The distribution is sometimes called the Pareto Distribution after the Italian economist, Vilfredo Pareto, who calculated that 20 percent of the population owned 80 percent of the wealth. It is also sometimes called the Zipf Distribution, after the Harvard linguist, George Zipf, who found that the frequency with which a word was used in a language was proportional to 1 divided by the word’s frequency rank among all words, a  $1/x$  mathematical formulation that, when graphed, shows the same properties as other  $y = ax^k$  power law functions. Indeed, once statisticians started focusing on power law functions, they were able to find many examples in the natural and the social worlds that were described well by those distributions.

It is no coincidence that media companies have felt a need to come to terms with the theory of The Long Tail—many of the examples that Anderson offers come from the media.

#### Television

Television broadcast networks had already been seeing their share of audience decline as alternative distribution channels such as cable and satellite had catered to more fragmented niches. However, with the advent of the internet as a mechanism for the distribution of video programming (e.g., YouTube), the proliferation of niche-content video is virtually boundless. Television broadcast networks operate in a conventional scarcity-based economy because the broadcast spectrum is very limited and thus each channel renting that spectrum must attract large enough audiences to recoup costs through advertising revenues. However, digital services such as YouTube allow effectively free distribution of video content, thereby providing a platform for micro-niche video to find its audience.

#### Music

Though popular music has long enjoyed a rich diversity of styles, the distributional bottlenecks (limited number of radio stations, limited number of music companies, limited shelf space at record stores) tended to constrict the range of listening options to “Top 40” hits. Brick-and-mortar record stores typically recruit from a 10-mile radius, thereby restricting their market; in practice, this meant that the CDs stocked at such a store had to be popular enough to repay the rent on the shelf space—a high popularity hurdle when the store’s market is so geographically constricted. Because of this, physical stores almost always confined themselves mostly to stocking recently released CDs with the broadest popular appeal. However, with the growth of the internet and digital distribution systems (e.g., iTunes, Rhapsody, eCast), the aggregate market for niche content has expanded enormously. Companies in these businesses find that a large fraction of their sales come from the “ones-ies and twos-ies” in the Long Tail. Because of the popularity of these companies that cater to Long Tail tastes, the brick-and-mortar record stores have nearly disappeared from the landscape. Even the once-powerful Tower Records is now defunct.

#### DVD

Similarly, DVD rental shops typically attract buyers from a fairly small radius, so it is not feasible for them to stock much esoterica. Every item on the shelves must rent with sufficient frequency to repay the rent on the shelf space, and in most markets, this precludes a population large enough to support the stocking of all but the “hit” DVDs. But by leveraging the national scale afforded by the internet, Netflix has transformed the DVD rental business. Only 30 percent of Netflix rentals come from current releases; 70 percent come from the (niche) back catalog. While the brick-and-mortar Blockbuster Video has not yet followed Tower into oblivion, pressure from Netflix has forced it to cater more aggressively to niche tastes—in part by establishing a digital service akin to Netflix’s. Despite that shift, 70 percent of Blockbuster’s rental volume still comes from current release, and only 30 percent from back catalog.

#### Books

The average Borders book store carries about 100,000 titles and recruits buyers from about 10 miles around each store. As such, it has to worry that it is stocking the most popular choices because it is saddled with high inventory costs that oblige it to turn over its stock often enough to pay the rent on its stores. The digitally-based [Amazon.com](http://Amazon.com) has a market not bounded by geography, the benefit of vastly lower inventory costs, and an ability to pass along at least some of its distribution costs (mail) to the consumer. With an inventory of nearly 4 million books, Amazon can afford to cater to every niche taste, and its digital platform allows it to sell to a global market. In breaking the tyranny of locality, Amazon allows its buyers to shop for books down the Long Tail. Indeed, more than 25 percent of Amazon’s sales are not available for purchase in offline retail stores.

As Anderson admits, once one starts looking for Long Tails, one starts to see them everywhere. Google has tapped the Long Tail of small advertisers. eBay has tapped the Long Tail of tag sales. Craigslist has tapped the Long Tail of classified advertising. Blogs have tapped the Long Tail of print journalism. Wikipedia has tapped the Long Tail of human knowledge. Open Source has tapped the Long Tail of programming expertise. Barack Obama has tapped the Long Tail of political donors. There even are defense analysts who consider Al Qaeda's terrorism and "asymmetrical warfare" as representative of the Long Tail of political violence. Anderson's book is replete with examples of Long Tail phenomena in numerous spheres—even those well beyond the digital realm (e.g., the Long Tail of beers, the Long Tail of flour, the Long Tail of KitchenAid mixers). In the balance of this article, we will consider the implications of The Long Tail for media audience measurement.

## MEDIA AUDIENCE MEASUREMENT AS WE HAVE KNOWN IT

Though the specific techniques for measuring media audiences have varied over time and across media, there are three features that have been fairly constant:

### Audience Measurement as "Currency"

The advertising-supported media have functioned as three-way markets among media, consumers, and advertisers: the media have attracted audiences that they in turn have re-sold to advertisers. Media audience measurement companies have been the neutral arbiters of the process, providing the "currency" of exchange between media buyers and sellers. As such they have been responsible for verifying the size and composition of those media audiences, thereby rationalizing media planning and reducing friction in the marketplace. Because of the high cost of collecting data and the reluctance of market participants to accept multiple "currencies" with differing values, there has tended to be only one single, dominant supplier of media audience data within each media sector (e.g., in the United States, Nielsen for TV, Arbitron for radio, MRI for magazines, etc.).

### Probability Sampling

Because of their critical role in enabling economic exchange in these three-way marketplaces, the media audience measurement companies have been very closely scrutinized for accuracy and methodological rigor. In the United States, this has taken the form of detailed process audits, conducted on behalf of media buyers and sellers by the Media Rating Council. These process audits are particularly attentive to sampling methods, with the "gold standard" associated with the random probabilistic methods at the core of inferential statistics. These methods generally assume Gaussian normal distributions (i.e., the bell curve) and prescribe detailed procedures for maintaining the integrity of samples designed to make estimates assuming those population distributions. Media audience measurement methods that do not adhere to the rules laid out in this field of sampling theory for inferential statistics have generally been rejected.

### Only Measure what is Efficient to Measure

Though media have been fragmenting into ever-smaller niches for some time, it has only been economically viable to measure the audiences of the largest media outlets. Hence, individual programs on broadcast networks each get their own Nielsen ratings, but the much smaller cable channels might only get a rating for the entire channel. SRDS reports more than 5,000 magazines published in the United States, but MRI can only afford to make statistically-sound projections for the largest of them, and thus MRI only estimates the readership of about 250 out of the 5,000 magazines published. In effect, the sample-based methods of audience measurement have limited our focus to the "hits" found at the head of The Long Tail. Our reliance on Gaussian inferential statistics has made it economically impractical to estimate the audiences for the niche media in the tail part of The Long Tail.

## MEDIA AUDIENCE MEASUREMENT IN A LONG TAIL FUTURE

Each of these three features is likely to change as a result of changes initiated or accelerated by the Long Tail dynamics described by Anderson.

While media audience measurement will continue to play a role as a neutral third-party arbiter in the exchange between media buyers and sellers, it will be harder for measurement firms to maintain their "natural monopolies" within their discrete market sector. Indeed, digitalization increasingly blurs the distinctions among the sectors, and consumers increasingly tend to consume their media with precious little concern for the platform of delivery. Thus, TV programs flow to the internet, print content is sent to telephones, and radio listeners break the bounds of locality by listening to favorite programs from distant radio stations via satellite or internet. Advertisers are also breaking out of silos, planning cross-platform campaigns that frequently mix traditional media with PR activities, sponsorships, events, product placements, and other forms of promotion. Even the most ambitious and agile audience measurement company is likely to be hard pressed to measure and capture such diverse activities at the level of precision currently associated with "media currency." As their coverage becomes more partial, their role as neutral arbiter in the market will become less central—less like a "currency."

The industry's near religious faith in random sampling and probabilistic inferential statistics will also come under increasing pressure in the Long Tail future. Back when the high cost of inventory or distribution artificially limited the range of consumer media choices, one could afford to build expensive random probability samples to estimate the audiences for the major media because those audiences were, in fact, big enough to justify the cost. However as "true" consumer demand reveals itself and shifts down the Long Tail, it will not be economically viable to apply conventional random sampling methods to the measurement of increasingly minute audiences. Indeed, digital technologies are also undermining the overall apparatus of the random probability sample survey by empowering consumers to evade the survey taker (e.g., caller ID, computerized do-not-contact lists, spam filters). As a result, the costs of gathering samples based on the random probability sample paradigm continue to climb, while the quality of the data thereby generated continues to decline. With ever-falling response rates, many in the industry have questioned whether the basic assumptions of Gaussian statistics are now being met by the existing practices of the media audience measurement enterprises. Even if the measurement of Long Tail niche media could be accomplished (expensively) through a vast expansion of sample sizes, the underlying decline in respondent cooperation, the concomitant rising cost of data collection, and the resulting decline in data quality call into question the efficacy of the entire paradigm.

These adverse trends likely will lead, in time, to a loss of confidence in the random probability sample as the sole source of "truth" in media audience ratings. But what can replace this venerable and scientifically-grounded paradigm? Surely the large and heterogeneous media economy will continue to need valid and reliable data to abet market decision making, even in a Long Tail future.

Anderson does not address the question himself, though he hints that the “votes” of interested elites may provide an alternative metric, if not for audience size, then at least as a guide to which media outlets are most influential. He cites, for example, Technorati’s rankings of media by the number of incoming web links as a non-sample-based metric for the measurement of popularity in the blogosphere. Indeed, the frequency distribution for Technorati’s rankings serves as a splendid example of a Long Tail, with a few big players (e.g., [nytimes.com](#)) showing themselves as highly popular hits, but with the distribution of links quickly falling to a nearly infinitely Long Tail of blogsites. In a similar vein, he points to websites like [digg.com](#) that put current news stories up for a vote among site visitors, as a means of finding out what stories spark the greatest interest among news-oriented elites. Neither Technorati nor Digg rely upon survey sampling methodologies, nor do they purport to be projectable to a large, general population. They interest Anderson because they represent economically sustainable mechanisms for capturing Long Tail phenomena, and not just the “hits” that are captured by the Gaussian-based random probability methods of current media audience measurement. But are such methods a viable alternative for the measurement of advertising-based media? Probably not. Indeed, they do not begin to address the fundamental questions that advertisers ask such measurement systems to answer: How many people saw my advertisement? How often? What kinds of people? As difficult as it may be to measure reach and frequency in the Long Tail future, these still will remain basic questions—without which we will have difficulty measuring other advertising outcomes or assigning value to media exposures.

So what does this portend for media audience measurement? One possible direction, at least in the near term, may be the adoption of what I will call “hybrid systems.” Hybrid systems will continue to use random probability sampling to measure the bigger media events in the head of The Long Tail (e.g., the biggest TV programs, the largest magazines), but will use (non-sample-based) measures for the niche media events in the Long Tail. Such methodological stratification has precedent in media audience measurement. Consider, for example, that media planners use (sample-based) MRI audience data for the largest magazines, but (non-sample-based) ABC circulation data for the magazines that are too tiny to be measured by MRI. Similarly, for many years, the U.S. television market was characterized by a stratification between the methodologies used for national program ratings (continuous measurement through people meters) and the methodologies used for smaller local markets (diary measurement in February, May, and November): though both systems used probability sampling, the cost and caliber of measurement were made to fit the economics of the respective national and local markets.

Hybrid systems are, to some extent, already being used for website audience measurement. While third-party audience measurement companies like comScore and Nielsen Online provide sample-based probabilistic estimates of overall monthly audiences for the largest websites, the granular details about daily traffic on specific parts of those websites come directly from the sites’ (non-sample-based) logfile reports. And smaller websites even fall below the radar of the monthly statistics produced by the third-party measurement companies; for these websites, the only traffic data available are those derived from the (non-sample-based) log-files. The result is that most companies are forced to integrate sample-based and non-sample-based data streams in order to make sense of their own website data. For example, at my own company, Condé Nast, sample-based data are only available for our largest sites (e.g., [style.com](#), [epicurious.com](#), [wired.com](#)), but not for many of the smaller websites associated with our magazine brands (e.g., *Architectural Digest*, *Gourmet*, *GQ*, etc.). And even on the larger sites for which third-party data are available, the data resolution is so poor that one can only see overall traffic statistics for the overall site for the month. To see where traffic went on those sites, we must revert to the more granular (non-sample-based) logfile data reports generated by our web servers. Thus, in current practice, participants in the web media markets effect their own hybrids, mixing data from sample-based and non-sample-based sources to fit the requirements of the situation. Such provisional solutions are likely to become more common with other media in the Long Tail future.

A similar hybrid system is beginning to take shape in the television sector. Program ratings still come from Nielsen panels built according to the rules of conventional probability sampling. But experiments are already underway to capture detailed program viewing data from the video servers of cable companies. Under the “old rules” of the sample-based paradigm, this would never be permissible because cable users only represent a fraction of the total universe of television viewers. Even if one had a complete census of the program viewing of all of the subscribers to every cable system, this would still exclude information about those who get their TV signal from other sources (e.g., satellite systems, internet, over-the-air broadcast). What’s more, the experiments are only being done in specific markets of the largest cable systems; in other words, they do not purport to be a random sample even of the entire universe of cable viewers. Data on viewing that come directly from cable systems do not include the other information traditionally associated with Nielsen’s sample-based systems—information on how many viewers are watching the set, their ages, and genders. However, the tradeoff is clear: the data from a cable server, like the data from a web server, can provide most granular detail on TV set usage, even if it lacks demographic details. Though cable households only represent a part of the universe of television viewers, the industry is poised to countenance a hybridization of data streams on TV viewing—with the sample-based probabilistic panels providing the overview of the market, but with the non-probabilistic cable data providing a granular view of a key sector of the market. We have yet to work out a process for integrating these data streams and interpreting their meaning, but that does not seem to be slowing the efforts in this direction. As TV programs are increasingly sourced from the internet, look for this hybridization process to accelerate.

Over the longer term, digital technology may render obsolete all such efforts to measure the audiences for advertising-supported media. After all, advertisers ultimately are more interested in the audiences for their advertisements than in the audiences for the TV programs, magazines, radio shows, or websites containing their advertisements. On the internet, display advertisements are usually inserted digitally by third-party advertising servers, which, in turn, provide a (non-sample-based) accounting to advertisers of the number of advertisements served in each website location. In principle, such third-party digital insertions could also be used for TV advertisements in a digital TV future. Or if one dares to imagine a world in which magazines are downloaded to some kind of digital reader, one can just as easily imagine third-party digital insertion of advertisements in that medium. To some extent, this will relieve us of the need for third-party media audience measurement, though not entirely. Though such digital insertion systems can tell us where the advertisements ran, there likely will still remain questions about how many people were present, their demographic and behavioral characteristics, and their level of attention. In other words, we still are probably going to need to integrate different data streams to be able to get the full picture needed for effective media planning—and for the establishment of the advertiser value of the advertising exposure afforded by the media. While direct-response advertising may avoid this need through Google-like pay-per-response systems, brand advertising—especially for products with longer consideration and purchase cycles—will still need to calculate the probable value of a media exposure even if that exposure happens in a niche far down The Long Tail.

In a world of “hits,” it was economically rational only to measure the biggest media events—the top networks and shows, the biggest magazines, etc. However, as the micro-niche events in the Long Tail come to represent a larger share of consumer attention and a bigger fraction of the media economy, it will no longer be rational to consign them to the status of “unmeasured media.” However, this will pose a serious challenge to a media industry that has tended to abhor data integration, modeling, fusion, and hybrid systems. The fragmentation and transformation described by Anderson in *The Long Tail*, however, are likely to make such hybrids a necessity.

Scott McDonald is the senior vice president for market research at Condé Nast Publications. He oversees market research, consumer research, advertising research, editorial research, and development research for all of the Condé Nast magazines. This includes such well-known titles as *Vogue*, *Architectural Digest*, *Glamour*, *Self*, *GQ*, *Vanity Fair*, *Gourmet*, *Bon Appetit*, *Condé Nast Traveler*, *Allure*, *House &*

*Garden, Wired, Lucky, Bride's, Modern Bride, W, Details, Jane, Cookie, Domino, Teen Vogue, Men's Vogue, and The New Yorker*. He also oversees research for Condé Nast, including such sites as [Style.com](http://Style.com), [Epicurious.com](http://Epicurious.com), [Concierge.com](http://Concierge.com), [Men'sStyle.com](http://Men'sStyle.com), [Brides.com](http://Brides.com), [Flip.com](http://Flip.com), and the companion sites to the Condé Nast magazines.

Prior to joining Condé Nast, Dr. McDonald was the director of research for Time Warner Inc. He was instrumental in the redesign of *TIME Magazine*, the growth of the *People* franchise, and the launches of both *Entertainment Weekly* and *Martha Stewart Living*.

Dr. McDonald has been at the forefront of industry efforts to develop standards of audience measurement for interactive media. He organized the first industry conferences on these topics for the Advertising Research Foundation (ARF), and he served as the first research chair for the Internet Advertising Bureau. He is co-author of the FAST Standards for Internet Audience Measurement, standards that have been adopted widely in the United States and Europe. In 2004 he was appointed to the Congressional Task Force investigating allegations of racial bias in the implementation of Nielsen's local people meter television audience measurement system.

Dr. McDonald has served as the chairman of the board of directors of the ARF. He also serves on the Program Committee of the Worldwide Readership Research Symposium, as well as on the board of the Media Research Council. He is a trustee of the Marketing Science Institute. He also is a member of the European Society for Marketing and Opinion Research, of the Market Research Council, of the American Association of Public Opinion Research, and of the American Sociological Association. Over the years he has delivered numerous scholarly papers before each of these professional associations.

Dr. McDonald comes to the media research business from an academic background in both demography and survey statistics. He holds an A.B. in sociology from the University of California at Berkeley and a Ph.D. in sociology from Harvard University. Since 1997, he has taught as an adjunct professor at the Graduate School of Business at Columbia University, and he has previously taught in the sociology departments at both Harvard and New York University. He has also lectured and published on a broad range of topics in the social sciences. He has been a recipient of the National Science Foundation Fellowship and of the Charles Abrams Fellowship of Harvard and MIT.

[scott\\_mcdonald@condenast.com](mailto:scott_mcdonald@condenast.com)

## REFERENCE

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

---

### © Copyright Advertising Research Foundation 2008

Advertising Research Foundation

432 Park Avenue South, 6th Floor, New York, NY 10016

Tel: +1 (212) 751-5656, Fax: +1 (212) 319-5265

All rights reserved including database rights. This electronic file is for the personal use of authorised users based at the subscribing company's office location. It may not be reproduced, posted on intranets, extranets or the internet, e-mailed, archived or shared electronically either within the purchaser's organisation or externally without express written permission from Warc.



[www.warc.com](http://www.warc.com)