



Rethinking the Attention Economy

The American Beauty rose can be produced in its splendor and fragrance only by sacrificing the early buds which grow up around it.

—John D. Rockefeller, Jr. on trusts, quoted in Ida Tarbell's
The History of the Standard Oil Company

In early 2000, Google conducted one of its first online experiments. The result was a disaster.

Google's mistake started with a simple question: How many results should it return per search query? The young company had always given users ten results, because that was what previous leader AltaVista had done. Competing search engines like Yahoo! returned twenty, though, and focus groups suggested that users preferred more.¹

Google's researchers decided to try a real world test, splitting off groups of users to receive twenty, twenty-five, or thirty results instead of ten. But when they checked a month later, they found—to their shock—that more results had produced huge *drops* in traffic. Searches in the thirty-result group had fallen by more than 20 percent, and tens of thousands of users had abandoned Google altogether.

The researchers rushed to figure out what had gone wrong. Were users overwhelmed by more results? The data showed no evidence of this. Were users just clicking the “next” button less? Few users clicked on the next button to begin with, so this effect was tiny.

Google eventually traced the traffic drop to a surprising source: It took a fraction of a second longer to return more results. The control group waited 0.4 seconds on average, while those in the twenty-five-result group

waited 0.9 seconds. Over a day or two this added lag meant little. But as the weeks wore on, the effects of that extra half-second multiplied. People visited Google less often, and performed fewer searches when they did visit. Even when the experiment ended, slowed-down users did not come back immediately. Their usage started to increase again, but from the new, lower baseline.²

There are several morals to the story of Google's early experiments, which this book will unpack at length. But the most important lesson is about how to understand online advantage.

Digital survival depends on *stickiness*—firms' ability to attract users, to get them to stay longer, and to make them return again and again. Stickiness is like a constantly compounding internet interest rate, in which a small early edge in growth creates a huge long-term gap. Differences in stickiness do not add up, they *multiply together*.

Google's ascent from upstart to the world's most valuable company came from learning this lesson. Google spent billions to make its site faster—but it also did much, much more. It reinvented itself not as a search engine but as a bundle of the stickiest online activities: email, video, maps, mobile, even office software. And in its pursuit of an ever-larger slice of users' attention Google rebuilt the fundamental infrastructure of the internet: new data centers larger than any before, new fiber optic cables, new ways of running rich applications over the web and a speedy new browser to match, new forms of artificial intelligence running on new types of computer chips, even a new mobile operating system that now runs on two billion active smartphones.

How did online audiences and digital revenue get so concentrated? What does this concentration mean for business, politics, news, even national security? Is online oligopoly inevitable, or is there a way out of the internet trap? These are the questions this book seeks to answer. For us today, just as for Google two decades ago, the first lesson is this: small effects that compound quickly *are not small effects*.

A SCARCITY OF ATTENTION

The World Wide Web is the most astonishingly successful technology of modern times, a technology that now underpins our social, economic,

and political lives. You likely ordered this book online—or even had it wirelessly delivered to your phone or tablet or reader. The web is so popular, powerful, and omnipresent that we forget just how badly it has failed at its original goal.

The World Wide Web was built in an explicit attempt to eliminate hierarchy in communications. It tried to do this by combining other technologies developed with that same goal in mind. The first such technology was hypertext, originated in the 1960s by sociologist Ted Nelson. Hypertext documents were designed to be consumed nonsequentially. Links within the hypertext could point to other relevant passages, definitions, graphics, tables, or even other documents on the same computer. The second technology was the internet, which by the 1980s had become ubiquitous in universities and research labs. The internet had been created as a peer-to-peer network, in which there were no central hubs: each computer could send and receive data with any other computer.

Tim Berners-Lee, the creator of the web, saw that hypertext could piggyback on top of the internet. Instead of being sandboxed within a single computer, hypertext could link documents on computers continents apart. Berners-Lee called the project the World Wide Web to emphasize that “any node can be linked to any other,” and to reflect “the distributed nature of the people and computers that the system could link.”³ In the opening sentence of the www’s project overview, one of the very first pages online, Berners-Lee declared, “There is no top to the Web.” All pages, and all sites, were supposed to be created equal.

If there were no “top” sites on the World Wide Web in 1991, there certainly are now. The fact that some sites and apps are far more popular than others is the most important fact of online life. We network on Facebook, while competing sites get less than 1 percent of Facebook’s traffic. We search the web using Google or Bing, and competing search engines have minuscule market share. We read and write our email using Google or Yahoo! or Microsoft. We use eBay for auctions, and we buy our books (and increasingly everything else) from Amazon. Among hundreds of millions of sites on the web, the four largest internet firms—Google, Facebook, Microsoft, and Yahoo!—capture a third of all web visits.

Concentration in online revenue is even more dramatic. The ten largest digital firms have always dominated digital advertising, taking

three-quarters of ad dollars since at least the mid-1990s. But the shift to mobile and video has intensified concentration at the very top. As of mid-2016, Google and Facebook together combined for more than 73 percent of digital advertising in the United States, a remarkable duopoly over a \$60 billion-a-year industry.⁴

At its core, this is a book about the *attention economy*. It focuses on the interplay between money and attention, and how each is exchanged for the other. It shows how digital firms spend money—often mountains of money—to attract and keep audiences. The book shows why investments in things like server farms and fiber and backend software matter so much, and how sites with more content and better designs are able to tilt online competition in their favor.

The book also details how audiences are traded for revenue through ads and subscriptions. This exchange is highly unequal: for several reasons, larger outlets earn far more per audience member than smaller outlets. This reciprocal relationship between money and attention creates feedback loops, in which sites and apps that have early success can invest to become more successful still.

In itself, the idea of attention economics is not new. The notion is often traced back to Herbert Simon, political scientist and Nobel prize-winning economist. Writing about the problems of information overload in the 1960s, Simon wrote that

in an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.⁵

Simon's passage has been repeated so often that it verges on cliché, with Google claiming to find the quote on more than 4.5 million web pages.

Many scholars have claimed to pick up where Simon left off. In the past two decades many books and articles have been written about “attention scarcity,” though this phrase has animated radically different ideas about digital media. Some discussions have been cloaked in business jargon or cyber-utopian rhetoric. One early treatment even declared that “attention

transactions” would replace traditional money.⁶ More recent works, such as James Webster’s excellent *The Marketplace of Attention*, have relied not on an idealized vision of the web but on data on how digital media is actually consumed.

This book differs from previous work on attention economics in numerous ways. Too often attention economics has been invoked to argue that “old” economics does not apply, that the digital world is throwing off the shackles of the analog one. This book uses new data sources, together with models borrowed from economic domains such as international trade and mathematical finance, to show that the digital world is not so different after all. Large media firms still dominate for reasons economists will find both novel and familiar. Diverse preferences can serve to *concentrate* audiences rather than spread them out.

Even more fundamentally, this book challenges claims that “money cannot reliably buy attention.”⁷ This belief that digital attention cannot be bought is woven throughout internet scholarship and public debates—but it is flatly, undeniably wrong. Showing exactly how audiences are bought and sold is a central theme of this book.

There is one key point, though, in which this book is in agreement with previous work: understanding attention economics is crucial for understanding digital media’s social impact. As James Webster writes, “Media need an audience before they can achieve a purpose.”⁸ Much is at stake beyond just commercial interests. Thinking about attention economics leads us to consider the *political economy* of attention—the ways in which attention economics impacts politics, and vice versa. On one hand, political choices and public policies help determine who wins and loses online. On the other, the digital attention economy increasingly shapes public life, including what content is produced, where audiences go, and ultimately which news and democratic information citizens see.

FORCES OF CONCENTRATION

Building a better version of attention economics starts with a key problem: *our understanding of the internet has been lopsided*. The forces that disperse digital attention are widely understood, while the forces of concentration are not.

This sort of asymmetric understanding in internet scholarship looks a lot like the historically uneven advance of knowledge in other fields. Consider the history of geology. When geology was organized as a scientific field in the late 1700s, it was mostly the study of erosion. While geologists quickly learned how wind and water wore mountains down, figuring out what built them up in the first place was harder. It was not until the surprise discovery of plate tectonics in the 1960s that geology had good models of geologic uplift.⁹ In a sense, geology went for 150 years as only half of a discipline.

This sort of uneven understanding has occurred in the social sciences, too. Writing in 1997, Paul Krugman argued that economic geography—the study of where production and consumption happen in space—was like geology before the discovery of plate tectonics: “economists understood why economic activity spreads out, not why it becomes concentrated.”¹⁰ Economists had long understood how high rents push economic activity away from an urban center, a process modeled well by Econ 101 supply-and-demand models.¹¹ But such models had a critical problem: they could not explain why cities would ever form in the first place. Good models of city formation emerged only once economists understood how imperfect competition can create urban concentration—for example, how stockbrokers in lower Manhattan in New York City can outperform those in Manhattan, Kansas.

Those who write about digital media today face the same sort of problem that previously afflicted geology and spatial economics. The distribution of online audiences, like the distribution of people in the real world, is the result of a tug-of-war between forces of concentration and forces of dispersion. But while the forces of dispersion have been widely understood and celebrated, the forces that concentrate audiences have been systematically ignored.

Most writing to date has assumed that digital media will produce “centrifugal diffusion”¹² in audiences. Again and again, observers have said that the web is a “narrowcasting” or “pointcasting” medium, that it tilts the playing field toward small content producers, that it will, to paraphrase the U.S. Supreme Court in *Reno v. ACLU* (1997), make everyone a pamphleteer. We have been told that “the monolithic empires of mass media are dissolving into an array of cottage industries,”¹³ that the internet has

empowered an “army of Davids,”¹⁴ that the internet will mean “the end of big.”¹⁵ Argument has focused not on whether audiences were diffusing, but rather how far and how fast they were spreading out.

Traffic data, though, show that audiences keep stubbornly refusing to decentralize. This book is, in part, an outgrowth of my previous book *The Myth of Digital Democracy*, which showed that online audiences follow concentrated power law patterns. Some commentators dismissed this evidence of concentration as premature, saying that it ignored the rapidly evolving nature of the internet. Micah Sifry argued, for example, that “it is dangerous to make conclusive statements about such a young and dynamic space.”¹⁶ The book’s evidence of audience concentration online, Matt Bai suggested, “only reflects a particular moment in time.”¹⁷

This wait-and-see argument is a natural consequence of our lopsided understanding of online audiences. If the only forces that shape internet traffic are dispersive, then, indeed, we should just wait, and all will be leveled. Like Vladimir and Estragon in *Waiting for Godot*, many still hold out hope that the internet they’ve been waiting for will arrive. Eventually.

All of these commentators are right that some old forces of concentration do not apply in digital media. There is no scarcity of spectrum on the web. There is no need for digital-only media firms to maintain printing presses or a fleet of delivery vans. Motivated citizens can share video without a broadcast license. Single journalists can start their own digital publications with only a laptop and a Starbucks latte.

But the unrelenting focus on new, small-scale sites ignores the elephants in the room—the large digital platforms where users spend the most time, and that soak up nearly all online profits. In celebrating the profusion of citizen blogs, we must also understand why the vast majority are abandoned and unread. The centripetal forces deserve just as much attention as the centrifugal ones do.

So the body of this book begins by documenting the other half of the story. Chapter 2 details how large internet firms can take advantage of a host of economies of scale, even beyond simple network effects. As a group, large sites load faster. They are prettier and more usable. They have more content updated more frequently. They rank higher in search results. They have established brands, and visitors are more practiced in navigating them. They are more efficient platforms for advertising. There

is substantial evidence that each of these factors, individually, serves to concentrate traffic.

Chapter 3 goes further, showing that large firms and websites are far better at personalizing their content for users. Digital media defined by “The Daily Me” does not favor small-scale content producers, as most have assumed. Instead, it favors those with resources: money, staff, data, computing horsepower, intellectual property, an established audience.

BETTER MODELS

Detailing the forces that concentrate internet audiences is a start. But what has been especially needed are not new facts but *new models*, theories that explain the broad patterns of digital traffic and revenue from Facebook and Google all the way down to tiny personal blogs. What we need are simplified stories that can explain both the enormous concentration at the top of the web and the (very) long tail of smaller sites.

Chapters 4 and 5 take on the task of model building, using different but complementary approaches. Chapter 4 builds a formal economic model of online content production. This deductive model is based on three key assumptions. First, larger sites are assumed to have economies of scale, both in the production of content and in their ability to turn traffic into revenue. Second, users are assumed to have at least modest preferences for diversity. Third, users are assumed to face search costs or switching costs in seeking out new content.

Individually, these assumptions are uncontroversial—but combining them leads to some surprising results. Portal sites and aggregators are a market response to these stylized facts. Big sites can dominate *even if* smaller sites produce better content that perfectly matches users’ preferences. The model suggests, too, that search engines or social networks can have countervailing effects: pushing readers further afield, but making audiences of the largest search engines and portals even more valuable. Facebook’s push to host other sites’ news articles, or the economic logic of content farms and so-called fake news, are captured well by this simple model.

Chapter 5 takes a different tack, using data on fluctuations in web traffic. The wait-and-see crowd is right about one thing: the web is a dynamic medium, and there have been big gaps in our understanding of how web

traffic evolves over time. Sites gain and lose traffic every day. New sites constantly emerge, and old sites decline into irrelevance.

For perennial Internet optimists this churn is a cornerstone of their faith. Look at Facebook, the argument goes: barely a decade old, and now the most visited property on the Internet. Or look at the Huffington Post, founded in 2005, which became a top-ten online news outlet even before it merged with venerable AOL. Large web firms have been especially aggressive in pushing this line. Google has repeatedly told regulators that it does not need to be regulated, since “competition is only a click away.”

But these views hide a basic error. The easiest way to see the mistake is to consider another venue defined by both ceaseless concentration and constant change: the stock market.

Though thousands of stocks trade on public exchanges, most of the market’s value is concentrated in just a few dozen companies: big blue-chip firms like Apple, Microsoft, Google, Exxon, General Electric, and Bank of America. The smaller a company is, on average, the more volatile its stock price. Blue-chip companies are generally safer investments than smaller companies. Even when a blue-chip firm faces an unexpected negative shock—like, say, it spills hundreds of millions of gallons of oil into the Gulf of Mexico—its stock price drops only modestly. Moreover, the *structure* of the stock market is much more stable than the stocks themselves. We don’t know which stock will be the fiftieth largest tomorrow, but we know how much it will be worth relative to the other stocks in the market.

Chapter 5 shows that web traffic follows remarkably similar patterns. Big sites have more stable audiences—day to day, month to month, year to year. Smaller sites are much more volatile. While individual sites constantly rise and fall, the overall structure of web traffic is largely constant. We cannot predict which site will be the one-hundredth most visited site tomorrow, but we know what share of traffic it will receive.

It is intuitive to think that the traffic of Facebook or Google varies less, in percentage terms, than the traffic of CNN.com or the *New York Times*. It makes sense, too, that traffic to NYTimes.com is more stable than traffic to a small blog. The consequences of web traffic being structured this way are not obvious, but they are profound. As we will see, the “size equals stability” pattern—*by itself*—is enough to produce the power law patterns that dominate online life. This audience agglomeration is baked into the

math. And while the notion of a dynamic, constantly changing web is often invoked to claim that the web empowers small sites, smaller outlets turn over *much* faster than big sites.

NEWS AND THE PUBLIC SPHERE

The models of the attention economy that this book proposes are quite general. Much of the evidence in the first few parts of the book comes from the commercial sphere of the web and digital media, where the dynamics of the attention economy are particularly stark, and where the online dynamics are not so different from familiar offline patterns.

But one of the biggest contributions of these models is illuminating areas of the web that go beyond purely commercial content. The most important implications of the attention economy concern the public sphere—the mix of news and information, citizen discussion and collective action, that is at the core of democratic politics. And so the last few chapters of this book focus on the online public sphere, and—especially—on online news.

This turn toward news comes for several reasons. Not only is the health of news of crucial democratic importance, but it is also the area where the digital attention economy has wrought the biggest changes. Local newspapers in the United States have always produced most of the nation's journalism, and employed most American reporters. But as online news consumption has reached a tipping point, we have witnessed a historic collapse in print readership and revenue.

Even beyond the civic importance of news, and the rapidity of change in the news business, there are theoretical reasons for looking at news content. News content provides some of the most challenging tests of our models of attention economics, and it is the area where the book's analysis most strongly challenges previous work.

So chapters 6 and 7 look at news through the lens of attention economics. Chapter 6 begins by looking at local news on the web. Many have hoped that new online news sources would take up the slack as newspapers and local broadcasters struggled. Using data from comScore, based on a panel of more than 250,000 web users in the one-hundred largest local media markets, this chapter provides the most comprehensive look at online local news to date.

Newspaper publishers have repeatedly claimed that local newspapers have a revenue problem, not a readership problem—that “lots of people came, but lots of advertising didn’t.”¹⁸ In fact, local news sites get only about one-sixth of news traffic, or just one-half of 1 percent of traffic overall. Within local news markets, newspaper and television news sites soak up nearly all of the attention. Not a single web-native local news outlet—*not one* in any of the one hundred largest media markets—comes close to the traffic of a typical newspaper site or local TV site. Newspapers may be weakened, but these ailing monarchs are still larger than their upstart competitors.

Chapter 7 goes further, showing that our models provide actionable intelligence to strengthen local journalism. Publishers, technology leaders, academics, and policymakers have proposed a broad and contradictory set of “solutions” to the local news crisis. To date, nearly all of them are rooted in a misdiagnosis of the problem. Like it or not, preserving local journalism is mostly about helping newspapers make the transition to the digital age. Newspapers cannot monetize audience they do not have.

Any proposal to save local journalism must start from the dynamic nature of digital audiences. Local newspapers, and especially *smaller* local newspapers, have long broken all of the rules for building sticky sites. As a group they are slow to load, cluttered, and—let’s be honest—often ugly. And while newspapers increasingly pay attention to digital traffic, they often do not understand what online metrics really mean.

Compounded audience is the most powerful force on the internet. The success of local news in the digital age depends on this compounding process, on measuring stickiness and optimizing for it. In contrast to repeated, contentless calls for “experimentation” and “innovation” in news delivery, this book provides real-world metrics by which to measure success.

THERE IS NO SUCH THING AS A FREE AUDIENCE

Building a more rigorous version of attention economics makes us rethink many “obvious” things about the web. It even challenges the single most important assumption about the digital age: the belief that the internet makes distributing content nearly free.

In the opening pages of *The Wealth of Networks*, Yochai Benkler argues that the internet has abolished the “industrial” economics that pertained to older communication technologies:

The core distinguishing feature of communications, information, and cultural production since the mid-nineteenth century was that effective communication . . . required ever-larger investments of physical capital. Large-circulation mechanical presses, the telegraph system, powerful radio and later television transmitters, cable and satellite, and the main-frame computer became necessary to make information and communicate it on scales that went beyond the very local. . . . Information and cultural production took on, over the course of this period, a more industrial model than the economics of information itself would have required. The rise of the networked, computer-mediated communications environment has changed this basic fact.¹⁹

Many other scholars have followed suit. Clay Shirky writes in *Here Comes Everybody* that the challenges of “mass amateurization” center on two main questions: “What happens when the costs of reproduction and distribution go away? What happens when there is nothing unique about publishing anymore, because users can do it for themselves”?²⁰ Jay Rosen argues that the internet improves journalism by “driving towards zero the costs of getting it to people, and by vastly reducing the capital requirements for quality production”²¹ Or as one book reported, “In one key area the Internet is reducing the cost structure of media firms and content producers: it lowers the cost of distribution.”

Actually, that last quote was from me.²² I was wrong, and this book aims to explain why.

It is true that the cost per byte of moving, storing, or processing information is now cheap, and getting cheaper all the time. But we now move, store, and process so much data that the total spent is massive. *Who pays* those costs has shifted, with important consequences. But the society-wide costs of web servers and site development are far larger, in constant dollars, than the cost of telegraph lines or mechanical presses or television transmitters. Google’s data centers—or Amazon’s or Facebook’s—are exactly the sort of multi-*multi*-billion dollar capital expenditure that was supposed to be obsolete.

The more profound error made by those proclaiming free digital distribution, though, was an overly narrow definition of what counts as a distribution cost. Distribution costs, as Rosen's quote suggests, need to include all of the costs of getting content in front of citizens. But this involves not just data costs, or the costs of servers. Rather, the distribution cost of digital content is *the total cost of building up a digital audience over months and years*.

Distribution costs include the need to constantly post new content, since the volume of new content is a large factor in stickiness. Distribution costs include site design, and all the site features that increase reader engagement. Distribution costs include the staff and effort needed to personalize content, or even to just A/B test news headlines. Distribution costs include the expertise and infrastructure that go into search engine optimization. Distribution costs include the raw buying of traffic through ads and paid-for links. Distribution costs include the costs to build mobile apps and make websites mobile-friendly. Simply because the costs of online distribution are *different* than they are in other media does not mean that they are *small*.

Building an online audience is like pumping air into a balloon with a slow leak. One has to keep pumping, to keep up a constant level of investment, or previous efforts will quickly be lost. These indirect costs of distribution are not optional. For a news site or a blog, maintaining an above-average level of stickiness is a matter of life and death.

Chapter 8 thus concludes by offering an *evolutionary model* of digital audiences. Darwinian competition for attention does not produce the egalitarian internet many have assumed, because the traits needed for stickiness are not distributed equally. A tiny edge in stickiness allows winners to grow faster than the niches they occupy, monopolizing money and attention. Expensive distribution means that internet openness is neither intrinsic nor inevitable. And for all the cuddly talk of the internet as an "ecosystem," digital niches are just as brutal and fragile as those in the natural world.

Many had hoped that the web would make news and political debate less centralized, expand and diversify the number of journalists and news outlets, and make capital less important in gathering an audience. The number of outlets may have expanded, but the public sphere remains highly concentrated. The number of journalists has plummeted and "fake

news” has multiplied, but digital media are just as dependent on a few corporate gatekeepers as ever. Building a consistent news audience remains hugely expensive. The attention economy has doomed most of our civic hopes for the web. This book is both eulogy and postmortem—an explanation of why things have ended up this way, and why this is not a passing phase.

The unwillingness of many to abandon those failed hopes now threatens to make the situation even worse. For all of the internet’s faults, there are concrete steps that we can take to strengthen the online public sphere. But first we need to understand the gulf between our imagined, fictionalized internet and the less-inspiring reality.