

How Data Happened

A History from
the Age of Reason to the
Age of Algorithms

Chris Wiggins and Matthew L. Jones



W. W. NORTON & COMPANY

Independent Publishers Since 1923

Persuasion, Ads, and Venture Capital

In an information-rich world, the wealth of information means . . . a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients.

—Herbert Simon, 1971¹

Watch time was the priority. . . . Everything else was considered a distraction.

—(ex)-Google engineer Guillaume Chaslot, describing, in 2018, YouTube’s recommendation engine’s sole KPI²

“**T**he sidewalks were a bright medley of color’ on Fifth Avenue at the Easter Parade on Sunday, March 31, 1929, as “modern, prosperous New York was celebrating.” According to the front page of *The New York Times*, “About a dozen young women strolled back and forth between St. Thomas’s and St. Patrick’s while the parade was at its peak, ostentatiously smoking cigarettes. One of the group explained that the cigarettes were ‘torches of freedom,’ lighting the way to the day when women would smoke on the street as casually as men.”³ What *The Times* reported on, credulously, was not in fact a spontaneous uprising of women smokers advancing gender equality via tobacco products. The performance had

been engineered by Edward Bernays, “the father of public relations,” with funding from the American Tobacco Company. Bernays went on to create the field of public relations and advocate for *Propaganda* (as he titled his 1926 book on the subject, not yet exclusively a term of derision) to ensure functioning democracy.⁴ Passing seamlessly between the worlds of politics and marketing, and astutely aware of the overt and covert dynamics of “The Engineering of Consent,” his was a vision a century ahead of its time. In this chapter we’ll trace the rise of efforts over this past century to monetize attention, including the business forces which have recently been brought to bear on this drive. Like fertilizer and gasoline, advertising and venture capital (VC) investing might seem rather banal on their own; yet, as we will see, they give rise to an explosive combination when mixed.

Monetizing Attention

As the epigraph suggests, the economist and AI pioneer Herbert Simon realized early what the rise of computers and information processing would bring: an attention economy. Simon argued that as computers made the storage and transmission of information nearly free, greater scarcity of attention would accrue greater value, and thereby an economy for people’s attention would emerge. “It is not enough to know how much it costs to produce and transmit information; we must also know how much it costs, in terms of scarce attention, to receive it.”⁵ This early digital deluge met head-on with an existing industry: the business of advertising. Not long after Simon’s prophecy, the artists Richard Serra and Carlota Fay Schoolman decried the all-enveloping, corporate-sponsored worldview packaged, sold, and imposed on people by the dominant mechanism of information transmission of the time: television. Serra explained that they strove “to

make explicit” the “capitalist status quo” of broadcast television and its business model: a free service is provided to people in exchange for occasional interruptions funded by those who wish to persuade. In their video piece “Television Delivers People,” they argued

Commercial television delivers 20 million people a minute.

In commercial broadcasting the viewer pays for the privilege of having himself sold.

It is the consumer who is consumed.

You are the product of t.v.

You are delivered to the advertiser who is the customer.

He consumes you.

The viewer is not responsible for programming

You are the end product.⁶

In the 1980s, Neil Postman, later the chair of NYU’s Department of Culture and Communication, offered a warning about how this model could go wrong.⁷ Namely, he illuminated how this advertising model perverts the relationship between the sponsor and the content producer, and specifically in the case of 1980s broadcast television, the relationship between the advertisers and media. Postman argued that the media creates content that is more attention-grabbing than factual, since the actual number of viewers grabbed leads to more money earned for the broadcaster. Indeed, this creates a danger: that content creators more generally are incentivized to create content which is as entertaining

as possible; he entitled his book on the subject *Amusing Ourselves to Death* (1985).⁸ A limit to the negative impact of such a perversion was the finite (though large) stranglehold broadcast television had as the primary source of truth in defining the mass public worldview.

Said otherwise: The scale of the damage was limited by the scale of television viewing public.

The Net

The WorldWideWeb (WWW) project aims to allow links to be made to any information anywhere. . . . If you're interested in using the code, mail me. It's very prototype [sic] . . .

—Sir Tim Berners-Lee, 6 Aug 91 14:56:20 GMT,
*announcing the WWW on alt.hypertext*⁹

How did this dynamic—the nascent “attention economy,” in which “you are the product”—change as our diet of information (and the economy of advertising along with it) moved to the World Wide Web? “The Web” is often given the birthday of August 6, 1991, when Tim Berners Lee posted to a Usenet group about “the WorldWideWeb project.” By 1994, multiple companies had formed simply to try to organize the spiraling complexity of the web, as it rapidly grew without any centralized control, becoming difficult for users to navigate. Online advertising dates from the time of email (the 1970s) and Usenet discussion groups, even before the web. By 1996, several companies were selling online “banner ads” and interruptive “pop-ups.”¹⁰ By the mid-1990s, abundant advertising flourished along with “e-commerce” (now simply called “commerce”) companies, notably including the sites eBay and Amazon, both founded in 1995 as well as a welter of companies that failed in the first “dot-com” bust.

Against this backdrop, the former particle physicist turned media scholar Michael Goldhaber updated Simon's observation about the scarcity of attention in a new internet studies journal called *First Monday* (formed in 1996):

Information, however, would be an impossible basis for an economy, for one simple reason: economies are governed by what is scarce, and information, especially on the Net, is not only abundant, but overflowing.¹¹

Information is not scarce. And so, the thing that is flowing or that is constrained is attention.

Goldhaber was writing from a particular time of lack of crisis; he writes that “there's so much leisure time and yet we all feel busy” because “all of our creature comforts are taken into account.” (Obviously, he was writing from a First World perspective of affluence and comfort.) “And so, we have lots of time to surf the web.”¹² A similar observation was made a decade later by Jonah Peretti, founder of BuzzFeed; BuzzFeed's target market he called the “Bored at Work Network,” meaning the millions of people at jobs which provided ‘Net access but which left them bored and thus with a surfeit of attention: time to surf.¹³ Buzzfeed met this glut, initially, with cute kittens and racy celebrity pictures; it would come to be a “unicorn”—a start-up company valued at over a billion dollars.

Goldhaber realized that while attention was already valuable, the web made it possible for anyone to consume anyone else's (online) attention. Bernays's 1929 Easter Day contrivance required tremendous coordination and the existence of an attention-getting parade which could be exploited; with the web, anyone can publish and thereby potentially consume the attention of anyone else online. Goldhaber argued that this change would increase the importance of the brand

of the individual relative to that of the individual's employer, or any corporation. In the same way that star journalists might leave a newspaper to establish their own blog, or later their own newsletter (or even, more recently, their own Substack), the web allows individuals to develop their own command of online attention without the backing of a well-financed or powerful incumbent.¹⁴

It's useful to reflect on how different the relationships were between economy and attention before the web. In 1997, on the subject of electronic books, Goldhaber wrote, "At present, it's impractical to distribute books directly over the Internet, though it's easy to foresee that this won't be true for long, and physical books will be seen as cumbersome and quaint." Indeed, over the course of the COVID-19 pandemic, scholars' dependence on their books and libraries gave way to e-copies and scans thereof, assisting in the writing of the book you hold in your hands (or the e-book you hold in the palm of your hand)—and yet printed book sales exploded. Similarly, the relationship between norms of paywalls has clearly evolved in the intervening decades. Goldhaber advised, "If you have a website, don't charge for it, because that will just reduce attention. If you can't figure out how to afford it without charging, you may be doing something wrong."

However, twenty years later, we write this book while listening to music via a (paid) digital subscription service and note that *The New York Times* has a growing business based on a digital paywall. "Information wants to be free," many people believed in the 1990s, but the information-producing companies now depend on information being, in fact, expensive and requiring massive infrastructure to store and process. Just as information on the web can be organized in many ways, its infrastructure can be paid for in many ways, including, but not limited to, ads. Some vested interests, however, have long claimed that there's really one

sustainable way to pay for information on the internet: ads based on surveillance of the activities of users. The falsity of the claim doesn't negate its historical importance.

The attention economy, like others, is not a purely invisible hand; the government imposes, for example, copyright restrictions which limit the copying and distribution of some information, retaining its scarcity and thereby its value. The role of the state in setting the balance of powers will be central in our next chapter.¹⁵

If Information Wants to Be Free, Who Is Paying? And Who Will Build It?

Information wants to be free because it has become so cheap to distribute, copy, and recombine—too cheap to meter. It wants to be expensive because it can be immeasurably valuable to the recipient. That tension will not go away.

—Stewart Brand¹⁶

Goldhaber delivered his prophecies at a time when BackRub/Google at Stanford was a mere NSF-backed graduate student project to turn the labor of website authors into a ranking algorithm for web pages. Very soon, though, this project became the venture capital-backed start-up Google, proverbially born in a garage in Palo Alto. The original paper describing the PageRank algorithm—the technical advance differentiating Google from scores of other contemporary innovators organizing the web—included no mention of monetizing the algorithm's needed infrastructure via advertising. One could have imagined many other income models such as subscriptions, affiliate fees, or sponsored links. Ads won out.¹⁷

The early 2000s saw the emergence of a technological norm termed “Web 2.0”: the idea that all users could become publishers by providing user-generated content (UGC) to the sites who would host the content in exchange for, and, increasingly, by monetizing, these users’ labor and creativity. Dating from 1999, the phrase “Web 2.0” gained more prominence after the “Web 2.0 Conference” started in 2004 by O’Reilly Media.¹⁸ Tim O’Reilly, the company’s founder, began his career writing technical books, but after the “dot-com crash” of 2000, he diversified his business model to include conferences as well as technical books; later, he founded a venture capital firm of his own. A variety of sites proliferated to host UGC, further encouraging the explosion of such content. These sites required a combination of design and algorithmic optimization to navigate as well as resources to pay for server space and bandwidth (especially as video became prominent). The growing use of user-created content—the continuation of the promise of the democratization of creation on the internet—came ironically with the creation of new intermediaries that organized, but also increasingly sought to profit from, their users’ creativity.

Organizing the deluge of information on the internet had no one solution. On sites such as Reddit, communities are organized into dynamic subgroups (or “subreddits”) of different topics, with posts algorithmically sorted based on user voting. Without these design constraints and community labor, however, there was an opportunity to sort a structureless feed of posts algorithmically. While some sites, like Pinboard, the “antisocial bookmarking site,” chose a subscription model for revenue, the dominant norm by the end of the 2000s was advertising. Ad-supported UGC hosting sites took up the algorithmic challenge of choosing which of billions of pieces of content to present. These algorithms, as with all machine learning, are optimization algorithms,

requiring technologists to commit to a subjective design choice: What is the function to be optimized? Increasingly, designers chose time spent on a site—and therefore among ads—as the function to be optimized. This world, with ad-supported, algorithmically optimized UGC, saturates our daily life. You are soaking in it. While some readers no doubt grew up with it, its pervasiveness should not be conflated with necessity. This is not the way things were in the first decade of the World Wide Web. The victory of ads over other ways to pay for internet services is today seen as natural, maybe even inevitable. Powerful constituencies have worked hard to make us think this way.

To understand the growth of internet advertising, we need to see that major advertisers saw the web as far less accountable than television or print media. It’s highly counterintuitive to us today but they believed they understood the success of legacy media advertisements far better than they did online advertising. In the 1980s and 1990s, advertising firms competed based on their computational models of how best to reach audiences, which “served as a testing ground for the coming age of ubiquitous digital media.”¹⁹ An article in 1998 in *Advertising Age* noted, “Lack of accurate measurement and difficulty tracking return on investment are cited as the biggest barriers to buying online media in a survey conducted by the Association of National Advertisers earlier this year.”²⁰ Legacy advertisers helped create a very different web by demanding metrics and leveraging technologies to allow tracking, as historian Joseph Turow has documented. “Advertisers will be excited to spend more money online,” one major agency head who handled Proctor & Gamble and Dell remarked, “when the medium proves it is accountable.”²¹ With the coming of the internet, advertisers pushed for more accountability, requiring more data on audiences, using techniques that could track their attention.

Surveillance wasn't a foreign quality grafted onto capitalism by the new internet firms; it was something that emerged as a dance between the demands of the major legacy advertisers and advertising firms that pushed technologists to meet those metrics to provide accountability about the effectiveness of on-line ads, increasingly by developing granular profiles of users.

Tracking users required web browsers to include technologies that enable surveillance—most notoriously cookies and hidden tracking pixels. The loss of privacy quickly became obvious to many technologists responsible for the standards undergirding the web, some of whom sought to alter the standards for web browsers to afford users more protection by default. The nascent internet advertising industry fought back aggressively, including lobbying browser manufacturers. They ultimately won the battle to let a thousand million cookies bloom (bake?) on your computer. “It seems to me it’s an extreme reaction,” one executive argued, “from a bunch of people who are saying . . . ‘We’re going to convince you there is a privacy problem on the Web.’”²²

Naturally, there *was* a privacy problem on the web—that was the entire point. Successfully depicted as radicals against commerce, the technologists and privacy activists lost this battle. Industry pushback emphasized the idea that the web required advertising—and advertising needed to be able to track users. Once again, the burden of protecting privacy was left to individual users. In a 1998 report the Federal Trade Commission explained, “choice easily can be exercised by simply clicking a box on the computer screen that indicates a user’s decision with respect to the use and/or dissemination of the information being collected.”²³

Most infamously, DoubleClick, founded in 1995, connected selling ads to collecting data on millions of users. DoubleClick president and CEO Kevin O’Conner explained,

“The great paradox with targeting ads is that the more you are micro-targeting, the more reach you have to have”—the more you want to focus ads, the more you need to collect large amounts of data on each user.²⁴ While privacy advocates and government regulators alike challenged DoubleClick around 2000, the firm weathered the limited restrictions on its practices strikingly well.²⁵ The company, to be sure, made opting out possible, but, as Matthew Crain notes, the small number of “opt-outs were a drop in an ocean of surveillance” by default.²⁶ By the end of the 1990s, advertising on the web had become the norm, as had the lack of user privacy when browsing. Over the next decade, defaulting to the ad model would become the common norm among a flood of “disruptors” as well, with Facebook and many other now-successful companies choosing the ad model as their post-venture capital lifeblood.

Advertising on Google originally focused on search terms and context, not surveillance of users. Its entire business model soon underwent a dramatic transformation. In 2005, Google purchased a company founded to provide infrastructure to monetize the network of advertisements and users, a start-up which itself had already been acquired by a private equity company—DoubleClick. And on March 11, 2009, Google announced that surveillance advertising was to be its future: interest-based ads “will associate categories of interest—say sports, gardening, cars, pets—with your browser, based on the types of sites you visit and the pages you view. We may then use those interest categories to show you more relevant text and display ads.”²⁷

By the middle of the 2000s, it was clear that having the technology to allow anyone to buy and sell ads on other platforms was going to be extremely profitable for whomever could dominate as intermediaries. By this time, the ad model was accepted by the relevant sides in the content producing-consuming exchange: the founders, the investors, and the

users (albeit begrudgingly). It became normal for the founders of the companies; e.g., Google's Larry Page and Sergey Brin, and Facebook's Mark Zuckerberg, had to acquiesce that they were going to make money primarily by selling ads. Before long, they became the dominant platforms for web advertising. Investors had to find this first plausible and then normative. And users had to consent, even if unconsciously. If users collectively reject an ad-supported website because it has ads on it, that's the end of the website; if users instead shrug and accept that it's normal for sites to show ads despite the deleterious and interruptive effect on user experience, then it's a plausible business model. Similarly, if content producers—all those who post and tweet and upload their hopes, memes, and fears—agree to do so for free, then UGC continues to thrive as a product-technology-business model. Each of us makes these choices, and the continuation of such business models requires our public norms and markets to evolve, keeping up with private companies' changes in technical architecture.

DATA AND ADS

Let's double-click on the technological challenge of ad targeting we mentioned above: How would you get this done? What is necessary to create an ad exchange like DoubleClick? It requires software to display ads, but also an investment in machine learning to determine to whom to show ads and how to price them. By the late 1990s, large advertisers like Procter & Gamble worried that internet advertising firms were promising much but weren't yet delivering. Advertisers worried that they knew more about the audiences of traditional media than they did about web users. "To prove its worth," communications scholar Matthew Crain explains, "the internet advertising industry needed to improve its capacity to target specific groups of consumers and demonstrate that online ads could move the needle of consumer

behavior."²⁸ To demonstrate they could provide more return on investment, these firms claimed they needed more data on users—as well as better machine learning.

Sites experiment by showing different ads to different people and learning which ads are going to drive outcomes desired by marketers—usually a click or purchase. Undertaking these experiments and putting them into practice required abundant data, performant algorithms, and robust software engineering. One person who wrote about this at the moment was Jeff Hammerbacher, briefly (2006–2008) of Facebook. Hammerbacher attended Harvard with Facebook founder Mark Zuckerberg, originally in his class of 2004. After graduating Hammerbacher went to Bear Stearns became deeply bored with finance, and then moved to Facebook for two years, helping his then-boss Adam D'Angelo. D'Angelo is credited with creating Facebook's "growth team," which he later described as "a team of engineers who changed the product in various ways to make it more viral and get more users to sign up."²⁹ As discussed in chapter 10, Hammerbacher described this time in 2011 in an interview, "best minds of my generation are thinking about how to make people click on ads. That sucks."³⁰ Facebook and Google would go on to become two of the largest companies today, collectively dominating digital advertising.

"In advertising I'm wasting half my money. The problem is I don't know which half." Often attributed to the advertiser John Wanamaker, this quip's cultural impact can be measured by the street named after him in New York City, later the headquarters of AOL and Nielsen, and the site of Facebook's NYC office. In the digital domain, however, marketers fought hard to know which half they were wasting: since the late 1990s, clicks have been easy to record and users to track, and software may be deployed to use the same science extolled by R. A. Fisher in 1925: the randomized controlled trial, to learn which of several "treatments" (here, advertisements) opti-

mizes engagement. In fact, only a bit more algorithmic effort allows marketers not only to learn what is best (a descriptive analysis) but preferentially to serve the ad with greater engagement (a prescriptive analysis). Simple methods for doing so are now commonplace in web programming, exploiting mathematical methods dating as far back as 1933.³¹

What is needed to “make people click on ads”? Above all, a tremendous abundance of varied and fresh content, provided now largely by user-generated content. Choosing which of these billions of sources to share is the main output of these algorithms in digital products like Facebook’s Newsfeed or TikTok or YouTube’s recommended videos. This also requires historical data on engagement—the response of similar prior users to prior content (clicks, shares, likes, etc.). Recall what an advance in storing and making information usable has been achieved since 1945, when technologists were attempting to move off punch cards. Personalizing this optimization requires information about users; while “coarse” information like the device type or the geographic location can be useful to optimization, demographic information as had been understood in advertising for decades is not only useful for optimization but saleable to marketers. That is, marketers possessing a craft and domain intuition hope to spend on a particular “customer segment,” such as “NASCAR moms” or “soccer dads,” to use examples from a political marketer in *The New York Times*, and so it’s useful for advertising platforms to describe users in this language and to model predictively whether users fall into these groups.

Advertising and Reality

What does advertising and, in particular, its recent machine learning optimized form mean for the way that we all con-

struct reality from the world we perceive? What does it mean that our primary source of truth, delivered to us in the palms of our hands, is funded by, and optimized for, the surveillance ad model? We don’t have to accept a facile view of the power of advertising or of propaganda to recognize the need to grasp the systems of cultural mediation and selection at work.

To understand this better, first let’s discuss what optimization means for information and persuasion. This topic brings together design, data, and incentives; ultimately, it also involves discussion of who is doing the optimizing and why, which will bring us to understand how the business goals drive design choices, which in turn shape our perception of reality.

Like the common task frameworks introduced for optimizing handwriting recognition, or predicting movie reviews at Netflix, the reduction of advertising to simple metrics has allowed data scientists and product developers to turn the armamentarium of machine learning toward optimizing agreed-upon advertising metrics. An example of this dynamic is illustrated by the common digital advertising metric CPM or “cost per thousand impressions.” Marketers, who buy pixels, choose to place their advertisements on different platforms based in part on low CPMs; advertisers, who sell these pixels, instead seek to maximize the number of total impressions. (To be clear, not all digital advertising deals are so priced and negotiated, but this framing is a useful simplification of programmatic advertising and ad exchange networks.)

Optimizing impressions means getting users to view as much content as possible, and to return to use the product as much as possible. Like any other common task framework, this changes what is valued—not the happiness of the user, or the veracity of the information, simply the total number of

impressions. Similarly, if ad impressions are delivered with a regular cadence as the user uses the product, this means optimizing a “North Star” metric for use, often called “engagement.” For information platforms, then, more engagement means more money, irrespective of the nature of the information shown to the user.

Particularly with the shift to mobile from desktop computing, the rise of information platforms has required visual design choices which facilitate maximum information flow in a small space—often the palm of one’s hand. This includes stripping away contextual cues, such as the details of the source of the information when served in an aggregating feed of content as well as, importantly, mixing types of content: news, entertainment, posts from friends, posts from strangers, and, of course, advertisements. Facebook’s News Feed, which debuted in 2006, exemplifies such aggregation. Content could be sourced from billions of user-generated posts from around the web. This is the culmination of the state Postman warned about: “news,” in the sense of good faith attempts to provide factual, useful information to everyone in a community, no longer is separated from sponsored content, persuasive content, false or satirical content, and entertainment.³²

These design decisions can be optimized, in the same way Fisher advocated in 1925, using randomized controlled trials to drive a North Star metric. Industry terms these “A/B tests.” The real gain in engagement comes from algorithmically choosing which content to show which users. The design of the feed breaks the programming constraints of media such as radio or television, by presenting any form of engaging content adjacent to any other and to sponsored content. An additional constraint faced by traditional media was the “one to many” constraint, meaning that everyone saw the same content. Information platforms optimize

unfettered by this constraint, providing optimally engaging, distinct, bespoke realities to each individual.

Increasingly, technologists and non-technologists have noted patterns over the years in *what* types of content are the most engaging. As Postman warned, the most engaging content is not necessarily the most factual or useful. Alluding to Marshall McLuhan’s famous maxim, “the medium is the message,” law professor James Grimmelmann investigated this dynamic in his essay “The Platform Is the Message” (2008). Platforms “are carefully and constantly watching to see which content beats out its rivals in drawing attention,” Grimmelmann explains. “These platforms prioritize and promote the content most likely to grab users by the lapels.” While Udney Yule had only observational data to inform policy, websites can—and constantly do—perform interventions and record the resulting effects, rather than merely observing activities. When you use a digital product, you are experiencing an experiment involving you all the time. In the case of user-generated content (UGC), the choices of which content to supply to a feed are immense, well into the billions. When you’re a professional YouTuber (and they are legion), you’re carefully watching your numbers. Imagine that you find that your videos on a new topic draw a lot of attention. It is an incentive to go out there and produce more on that topic. The result is to amplify nearly any taste of any user segment. As Grimmelmann observes, “Recommendation engines may only supply car crashes in the sense of suggesting that, since you looked at this one, here’s another one. You may be interested in watching that.”³³ And—he notes—content will now be automatically produced to cater to newly created interests. Even if they don’t work to vend goods or sway our politics in a simple manner, platforms organized around engagement alter our informational world in unpredictable ways.

CAN'T ONE HEAL THIS AI WITH MORE AI?

Why can we not use supervised learning, which predicts an outcome in the absence of treatment, or reinforcement learning, which chooses the best treatment to maximize an outcome, to fix the “troublesome content” (e.g., fake news, mis/disinformation, hate speech, and abuse) problem? Challenges include the incentives of the companies to eliminate such engaging content (again, Grimmelmann: “Complaining about it doesn’t help, either. Hate clicks are still clicks”), but also the lack of clear labels. Even with tens of thousands of moderators, as in the case of Facebook, or even with abundant crowdsourced “ghost labor,” different readers will often disagree as to the nature of satire, parody, and good- or bad-faith arguments online. Facebook’s “Oversight Board,” for example, is often called upon to adjudicate decisions over particularly troubling content. As an “AI” problem, we are limited by the lack of clear judgment of even *actual* intelligences.

An additional conundrum in content moderation: even when platform companies do take a stand, banning or simply down-weighting the frequency at which troubling content is shown, such actions draw accusations of bias or censorship which in turn can draw attention to—i.e., amplify—the content they sought not to distribute. Design choices which label content as problematic have drawbacks as well, including both the attention drawn to such labels as well as “backfire effect,” a phenomenon sometimes observed in experiments in which correcting information results in increased faith in the contradicted claim.

DEMOCRATIZED PERSUASION, FROM MARKETING TO POLITICS

The composite “persuasion architecture,” as Zeynep Tufekci terms it, includes the machine learning algorithm as well as the product that instantiates the algorithms. These architec-

tures have become so optimized in statistical performance as well as ease of use that anybody can do it. As one example, Facebook’s “lookalike audiences” allows marketers to ask Facebook to find other people who look like the people who performed a certain action. Marketers can thus target content to new users who resemble demographically or behaviorally prior users who performed an action, like clicking on a particularly provocative link. This can be done without any particular savvy or intuition about market research or user psychology; just a small amount of money for the individual purchaser, though vast in aggregate for Facebook or Google.

Bernays realized, even in the 1920s, that the laws of persuasion apply equally in marketing or in politics. Before Bernays was honing the message for tobacco, he was shaping public perception in politics. In 1924 he arranged for a group of popular celebrities to appear with the “practically inarticulate” president Calvin Coolidge to improve his image. Today, aspiring politicians benefit from the same statistical methods as marketers, leveraging digital tools for market surveys to craft campaigns and messages as well as to target the optimal message to the optimal user. Bernays saw no line between advertising and politics; indeed, he saw this engineering of consent as a good thing for democracy. “Bernays saw this as an unavoidable part of any democracy,” according to Zeynep Tufekci. “He believed, like Dewey, Plato and Lippmann had, that the powerful had a structural advantage over the masses. . . . He urged well-meaning, technologically and empirically enabled politicians to become ‘philosopher-kings’ through techniques of manipulation and consent engineering.”³⁴

While these techniques can—and must be—used for good, Bernays was clear, they “can be subverted; demagogues can utilize the techniques for antidemocratic purposes with as much success as those who employ them for socially desir-

able ends.” A leader seeking good objectives, Bernays argued, must “apply his energies to mastering the operational know-how of consent engineering, and to out-maneuvering his opponents in the public interest.”³⁵ Even as Bernays’s bracingly honest use of the term “propaganda” fell out of favor during the Cold War, the effectiveness of “engineering consent” grew dramatically with the increasing use of data and algorithms to inform and optimize persuasive messaging. It was not long before the techniques perfected for digital advertising online came to inform political messaging.

In 2007, Rayid Ghani, then an employee at Accenture Technology Labs, described an “Individualized Promotion Planning system.” He celebrated how data would allow radically new forms of individualized targeting. “In addition to using newspapers, in-store displays, and end caps to highlight their products and run promotions, retailers can influence individuals in a vastly different way using individual consumer models.”³⁶ Business goals can be met by understanding customers at a granular level. This technology allows every company to target every customer as a separate person, not just a representative of a statistical category. Just such targeting is at the heart of marketing and political campaigns.³⁷ Ghani served as Obama’s chief scientist in the 2008 election, in part using these data-empowered views of the electorate as customer segments.

In an op-ed in *The New York Times* soon after the reelection of Obama in 2012, Ethan Roeder concluded with a celebration of the individual at the heart of the Obama campaign’s data strategy. “Campaigns are . . . moving toward treating each voter as a separate person.”³⁸ This rosy picture of persuasion is interpreted more darkly after the concerns of the past few years in which this technology has been characterized as too persuasive. Unbounded granularity, deep contextual information about individuals, and personalized

persuasion architectures which optimized for engagement had created, many feared, too great a success for democracy to function healthily.

This vision of an engineered public fills some with hope, and others with concern that informed consent has been infringed by those in power. Such tools could easily be put to ill use, and we now know they have. Tufekci warned in 2018, “To microtarget individuals with ads, today’s platforms massively surveil their users; then they use engagement-juicing algorithms to keep people onsite as long as possible. By now it’s clear that this system lends itself to authoritarian, manipulative, and discriminatory uses,” of which she gave numerous examples.³⁹ But you can’t microtarget somebody unless you differentiate them based on who they are or what they’ve done. And that requires you to have copious data and machine learning. While many feel that our free-market norms are not threatened by powerful companies delivering deeply personalized advertising, our concerns over power are heightened when these abilities empower the state. The algorithm, though, works for states and corporations alike.⁴⁰ Matthew Salganik warns us, “These capabilities are changing faster than our norms, rules, and laws.”⁴¹ We should add: they are likely changing faster than our analytical tools for understanding the relationship between our social and economic realities and the conceptual worlds, mediated in platforms, through which we experience and act in them.

Kenya’s 2017 election, Brexit, and the 2016 US election greatly amplified—and popularized—concerns about algorithmic manipulation. A flashpoint of concern was the firm Cambridge Analytica; then-CEO Alexander Nix shared a worldview in 2017 that echoed and updated that of Bernays:

There’s no question that the marketing and advertising world is ahead of the political marketing and

political communications world. And there are some things that I would definitely [say] I'm very proud of that we're doing which are innovative. And there are some things which is best practice digital advertising, best practice communications which we're taking from the commercial world and are bringing into politics.⁴²

Researchers have reached no consensus on the ultimate effects of this attempted manipulation.

Worrying about the effects of the adtech and persuasion architectures does not require us to believe the claims of advertisers and technical snake-oil salespeople about the effectiveness of their ads. Tim Hwang and Cory Doctorow have brilliantly stressed the deep limits and deceptions around targeted advertising.* While adtech, either for commerce or politics, surely doesn't work in the ways those hawking it suggest, it has dramatically transformed our media landscape and consolidated a landscape of digital advertisers into a near duopoly (Facebook and Google), with unpredictable effects. Facebook and Google don't need the ads to work as promised—they need advertisers to believe they work. Perhaps it's a shell game, but it's one that dominates our informational landscape ever more, for better and worse.

* Cory Doctorow writes, "Surveillance capitalists are like stage mentalists who claim that their extraordinary insights into human behavior let them guess the word that you wrote down and folded up in your pocket but who really use shills, hidden cameras, sleight of hand, and brute-force memorization to amaze you." "How to Destroy 'Surveillance Capitalism,'" *OneZero* (blog), August 26, 2020, <https://onezero.medium.com/how-to-destroy-surveillance-capitalism-8135e6744d59>; see Tim Hwang, *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet* (Farrar, Straus & Giroux, 2020).

We next discuss what happens when the advertising model—a norm for revenue—meets the venture capital model—a process for accelerating market innovation which, we will see, moves faster than norms and laws adapt.

Moving Fast: Venture Capital

Venture capital is not even a home run business. It's a grand slam business.

—Bill Gurley, general partner at
the venture capital firm Benchmark⁴³

The Ford Automotive Company introduced mass-produced cars in 1916, but it took years to develop norms around integrating cars into society and decades for regulation about consumer protection, such as seat belt laws. This timescale for innovation, in which technology, markets, norms, and laws have decades to equilibrate, seems quaint compared to the way software and information technology rapidly upends our norms today. Venture capital (VC) has helped to expedite this disruption in that massive growth can precede revenue and, in the case of consumer-facing companies, norms (and, even more so, regulation) that otherwise slow adoption of a new product.

Has venture capital been with us forever? Although investments have been, many people point to World War II as a time when, among other things, the venture model was born.⁴⁴ As an example, in 1946, George Doriot, later a Harvard Business School professor, was the quartermaster general in charge of logistics during World War II. After the war Doriot created ARDC, a publicly traded company for investing in long-term research developments, including many developments in the nascent computing industry. In subsequent

decades microprocessors and personal computing were to a great extent funded initially through venture capital; given the positive changes of, for example, personal computing on society and the economy, VC has been widely celebrated. Absent from this celebratory timeline is decades of military funding of computation, threaded through earlier chapters, as well as massive support for innovative small businesses through SBIR programs whereby the government effectively served as venture capital.⁴⁵ E-commerce, we've pointed out, grew in the wake of eBay and Amazon, VC-backed companies that did very well. The author and angel investor Jerry Neuman points out a sixteen-fold rise just over the five years 1978 to 1983 in venture capital, from \$218 million to \$2.6 billion.⁴⁶ This is funded by limited partners—companies, states, sovereign wealth funds, and especially huge pension funds—which, in turn, invest in the VC firms. The record is mixed as to whether these limited partners are investing wisely: as Tom Nicholas points out in *VC: An American History*, in aggregate, venture capital doesn't return much better than other forms of investing.⁴⁷ Nevertheless individual venture capitalists often believe that *they* are the venture capitalists who are going to do it right. Venture capital has been a big part of the data driven attention economy in the last two decades, including in advertising technology companies as well as consumer-facing news start-ups such as Vice and BuzzFeed.

Venture capital de-risks: investors provide abundant resources that enable start-ups to execute a search for a repeatable, scalable business model. To be clear, risk remains inherent in the model: VCs expect the vast majority of their portfolio companies to fail, so long as the few that succeed return such a large financial windfall as to make up for the lost investment in the rest of the portfolio. In many celebrated cases, such as Fairchild Semiconductor's contribution

to the rise of computing, the risk is a technical risk, such as, "can an integrated circuit be created cheaply and reliably from silicon?" A separate risk which VC can obviate is market risk: finding product-market fit means creating a product people will pay for at a price which allows a company to sustain. However, VC investments allow a company to grow to find new users even before deciding on a revenue model. This was the path taken, for example, by Facebook and Google, which grew a huge user base—people willing to integrate the product into their lives and habits—before settling on a revenue model and imposing it on their users.

More recently, venture capitalists have advocated for a scale of investing—so called blitzscaling—which allows companies to attempt to go from pre-revenue to market dominance, effectively by purchasing an entire market. Rather than competing with other companies in the marketplace, this strategy is to provide sufficient investment simply to buy the market. As an example, Uber could afford to undercut entire existing taxi and limousine markets, municipality by municipality, unencumbered by regulations which provided driver and passenger safety protections. Similarly, the co-working company WeWork attempted to flood cities with cheap and flexible office spaces, at a scale which would out-compete competitors reliant on a profitable scale of revenue from consumers. According to *Fast Company*, the venture capitalist Masayoshi Son blitzscaled, investing in the founding CEO Adam Neumann, without extensive diligence:

There, in the back seat, Son took out an iPad and wrote out the terms for a \$4.4 billion investment in the company. He drew two horizontal lines at the bottom, signed his name across one, and then handed the iPad to the then 37-year-old Neumann to scribble his name on the other. Neumann would

keep a photo of the agreement on his phone. “When Masa chose to invest in me for the first time, he only met me for 28 minutes. Okay?”⁴⁸

Monopolies can be bought, but can also be grown through reinforcing effects and various cycles. Often data-empowered companies provide a particular model for this. In the language of the venture capitalist Kai-Fu Lee, “more data leads to better products, which in turn attract more users, who generate more data that further improves the product. That combination of data and cash also attracts the top AI talent to the top companies, widening the gap between industry leaders and laggards.”⁴⁹ This model, along with investment, has helped make services such as Google, Facebook, and others so much better than their competitors—owing to a tremendous abundance of data used to train machine learning algorithms—that they clearly dominate the present markets in search and social media, respectively.

Consequences of Attention Economy and Venture Capital

Above we’ve argued that concerns about the attention economy are at least fifty years old, that venture capital can be traced back almost seventy-five years, and public relations has been with us for at least a century. The combination of optimized computational influence and the rapid scale provided by contemporary venture capital, however, is a particularly potent mix that we are still learning to integrate into our political and personal realities. To illustrate: the very first CTO of Facebook already created a sub-team in 2006 called “growth”—essential for understanding how Facebook works. They are a team of some of the best engineers, excelling at ensuring the growth of new services to

more and more active users. Doing so requires identifying key performance indicators (KPIs) and leveraging the data and computational infrastructure needed to optimize these KPIs. Facebook VP Andrew Bosworth wrote in a June 2016 memo: “The ugly truth is that we believe in connecting people so deeply that anything that allows us to connect more people more often is **de facto** good. It is perhaps the only area where the metrics do tell the true story as far as we are concerned.”⁵⁰ Other engineers, including those who are formerly at information platform companies, have pointed out problems there. The engineer Guillaume Chaslot, formerly of YouTube, writes, “watch time was a priority. Everything else is considered a distraction.”⁵¹ Growth at all cost is good for the CEO, but not always good for the arena. In *The Nature of Economies*, Jane Jacobs warns of the growth of companies that, like species in an ecosystem, can grow in ways that damage the ecosystem itself. What powers exist to check such growth when the damage is to the marketplace and the societal ecosystem itself?⁵²