

This is a section of [doi:10.7551/mitpress/13562.001.0001](https://doi.org/10.7551/mitpress/13562.001.0001)

Selling the American People

Advertising, Optimization, and the Origins of Adtech

By: Lee McGuigan

Citation:

Selling the American People: Advertising, Optimization, and the Origins of Adtech

By: Lee McGuigan

DOI: 10.7551/mitpress/13562.001.0001

ISBN (electronic): 9780262374248

Publisher: The MIT Press

Published: 2023



The MIT Press

1 ADTECH FLOWS: CLAIMS, LOGISTICS, AND OPTIMIZATION IN DIGITAL ADVERTISING

Adtech is a major circulatory system of the digital economy. Marketers, media companies, and the intermediaries between them use an agglomeration of tools and infrastructures to route money and personal data around the world and around the clock. Information about internet users, predictions about their worth and future actions, and bids to serve advertisements ping-pong from computer to computer hundreds of billions of times a day. These flows of information and commerce animate huge swaths of the for-profit internet, helping some of the most highly capitalized companies in history collect almost all their revenue. In 2021 advertisers spent an estimated \$211 billion on digital media in the United States and \$491 billion worldwide. Analysts project those figures will increase by 50 percent and 60 percent, respectively, by 2025. Google, Meta, and Amazon dominate this trade, collecting almost two-thirds of US digital advertising revenue.¹ Thousands of lesser companies orbit these oligopolists, rigging tackle on advertising markets and fishing for scraps of money and data. Whether they are unknown, well known, or just notorious, adtech companies are fixtures of modern life. As Tim Hwang explains, “The need to create a liquid market in human attention influences the architecture of the social spaces of the web.”²

Defining adtech is tricky. The term is used to characterize various means of automating and optimizing digital advertising processes. We could say that adtech is a collection of classification and decision systems running on proprietary data sets, algorithms, software platforms, and artificial intelligence products. The core operations in adtech cohere around two interrelated classes of activities: one is epistemological, involving predictions, conjectures, and

knowledge claims (about value, identity, affinity, authenticity, and probability); the other is logistical, formatting and trafficking information (including HTTP requests, metadata about ads and ad inventory, server logs, sensor readings, identity tokens, user profiles, invoices, and receipts) and facilitating transactions at the inhuman speed and scale required for a business that trades a nearly unlimited volume of commodities that are each, in some sense, unique. Does the person using a website belong to a valuable consumer population? What is access to someone from that population worth at this moment? How likely is it that this user will click on a certain ad? Has the user seen the ad before? Will this ad placement be visible for more than one second? What happened after the ad was served? Answers to these questions are just some of the predictions and claims that are the stock-in-trade of adtech. A tangle of standards and logistics helps produce them in the blink of an eye. Adtech is an assemblage of machines, markets, mechanisms, mathematics, and human workers, all of them generating, processing, and coordinating flows of information and commerce.

It is, admittedly, imprecise to speak of adtech as something singular and coherent, even though that is the fashion. A discernible cultural identity has been erected around a subset of companies and specialists calling themselves adtech, yet I suggest the term should include any application of information technologies to mediate advertising transactions, to manage the distribution of ads, and to execute and evaluate decisions about what ads to serve, to whom, and at what prices. I resort to the singular noun as a matter of convenience, but I intend for this discussion to embrace a wide variety of businesses and operations.

Arguably, too, a critical definition of adtech should focus not only on existing markets, technologies, and practices but also on ever-evolving promises of progress. Adtech stands for a flexible universe of perceived solutions to advertising problems and ambitions. The specifics of these solutions reflect the range of strategies across firms and market sectors, and a procession of technoscientific trends cuts the emperor's robes to whatever style is in season, from big data to blockchain to the next attractive buzzword.

Whatever the details, the fundamental and overriding purpose is this: to recognize and act on differences in value. Or, more precisely, adtech enacts claims about value, creating and exploiting new investment opportunities. Advertisers, intermediaries, and attention merchants all try to use advantages based on information, speed, or analytical muscle to buy or sell units of audience

attention and behavior more efficiently. Adtech comprises a loose cluster of difference engines trying to capture tiny but incremental fractions of surplus from billions upon billions of transactions. These engines are designed to identify or predict the profit potential of consumers and their moments of attention with competitive precision and thus increase the probability of achieving strategic objectives. In this sense, adtech is a kind of automated management science, an embodiment of decades-old dreams of remodeling advertising as a machinelike optimization system that orchestrates frictionless transactions and dynamic adaptations at the speed of light.

This chapter sketches the current adtech landscape. The sketch is necessarily incomplete. Adtech's economic and technical dimensions are vast, intricate, and constantly shifting. Even as I write, corporate and public policies related to data governance and privacy are forcing drastic changes in the business of collecting and monetizing personal information.³ After decades of advocacy by critical researchers and civil society groups, popular and legal resistance to surveillance advertising is not only altering adtech processes but also foreshadowing diminished growth for the industry, which may depress stock prices and affect companies' strategic plans as investors channel their speculative capital elsewhere. Rather than trying to anticipate the future, I focus on current conditions, and I keep the discussion at a general level because existing techniques may abruptly fall into obsolescence or non-compliance with the law. Still, I hope this chapter provides a descriptive map of adtech that helps readers situate themselves in the recent state of the art and navigate the historical contours traced throughout this book.⁴

BEHAVIORAL AND PROGRAMMATIC ADVERTISING

Adtech is typically associated with behavioral and programmatic advertising. Behavioral advertising means targeting media users with messages that have been selected or customized based on observations of or inferences about a person's actions and inclinations. It is sometimes called interest-based advertising, implying that ads reflect people's revealed or suspected interests, although it more accurately means that the targeted consumer is deemed interesting or relevant to the advertiser. The principal assumption here is that past behavior provides insight into an individual's habits and intentions and ultimately that person's value as a consumer, voter, or social actor. Data about lifestyle, location, buying patterns, psychographics,

and more are assembled into profiles that then factor into automated or computer-supported decisions that determine the distribution and pricing of ads. These profiles include personal details that users disclose voluntarily (such as through social networking accounts), as well as passively collected records of behavior (such as clickstreams). Most commonly and straightforwardly, advertisers target internet users based on their internet usage: the websites they visit, the ads they click on, the items they put in their shopping carts at e-commerce sites, and so on. Behavioral advertising decisions can also draw on data generated by offline activities, such as car or home ownership, purchases made with credit cards or recorded through a store's loyalty rewards program, and much more.

Marketers also use statistical techniques to try to identify attributes a person has not disclosed or exhibited, to detect patterns of behavior in large data sets, and to predict future considerations—such as a customer's profitability to a firm, a person's probable responses to certain stimuli, or an individual's passage into a life stage associated with big purchases and new brand loyalties, like getting married or divorced. Frequently, marketers act on behavioral data in time-sensitive situations, such as a momentary assessment of emotion or a consumer's presence near a retail location.⁵ The advertising industry has constructed elaborate technological scaffolds to help companies exploit micromomentary opportunities.

Most digital advertising depends on those scaffolds, using computers, algorithms, and auction mechanisms to automate and optimize decisions and transactions. This bundle of automation, optimization, and auctions is called programmatic advertising. "Programmatic advertising technology promises to make the ad buying system more efficient, and therefore cheaper," an editor at *Digiday* explains, "by removing humans from the process wherever possible. Humans get sick, need to sleep and come to work hungover. Machines do not."⁶ Besides replacing "unreliable" human workers with stone-cold rational calculators, programmatic advertising also promises to unlock new efficiencies by dynamically assessing and extracting the precise value of each advertising opportunity. Instead of placing ads next to content that attracts certain types of people, advertisers buy individual *impressions*—opportunities to serve ads to users who are defined by some set of measured or estimated properties—wherever they come up for sale, irrespective of the surrounding content. As one executive put it, "Ads are being delivered individually, one at a time, to consumers based on data that identifies them as members of some

ideal target audience an advertiser is seeking.⁷ The capacity to single out consumer targets and to control the distribution of ads to exclude individuals or classes of people who do not fit a desired profile is called *addressability*.

Typically, what digital advertisers buy with their money is either an amount of exposure among a target audience or some measurable user action. In the first case, advertisers pay when an ad has been served and/or viewed by a user, and the price is set in terms of cost per mille (CPM), or the cost per one thousand impressions. Costs vary across advertising formats (e.g., display and video), but in general, programmatic advertising is quite cheap. With typical CPMs of just a few dollars, the winner of a programmatic display auction often pays only a few cents (or less) per impression. In contrast to these exposure-based transactions, some ads are sold on a cost-per-click (CPC) or cost-per-action (CPA) basis, where advertisers pay when a user clicks on an ad, downloads an application (app), signs up for an account, buys a product, visits a store, donates to a political campaign, and so on. In other words, advertisers are paying for documented events.

Different pricing models have implications for the distribution of ads. Google search ads, for example, are sold on a CPC basis. When an advertiser bids in a keyword auction for the opportunity to have its ad appear next to search results, Google's auction mechanism factors into that bid an estimate of the probability that a user will click on the ad. Because Google and other search engines (e.g., Microsoft's Bing) get paid when users click, these companies are keen to learn how to predict click-through rates. They employ and consult small armies of researchers who work at the intersection of computer, data, and management sciences to study and fine-tune the design of automated auctions. This is also true of Facebook. Advertisers pay the platform to orchestrate a wide variety of actions, from clicks to likes to customer acquisition. Facebook's primary service is using algorithmic decision systems to distribute ads into situations with an apparently high probability of achieving specified outcomes. Inflected by the causal language of attribution, these outcomes are often called conversions, and they can be selected as the objectives Facebook will try to optimize when making ad delivery choices for its customers. Facebook even allows advertisers to target users that the platform's machine learning technology predicts will generate the highest return on advertising dollars spent.⁸

Programmatic advertising can refer to a number of different configurations, depending on how deals are brokered, who has access to a supply

of inventory, and which professionals and budgets are activated in the buying process. Auctions and algorithms are used to place addressable ads within social media feeds, on mobile apps, between songs on Spotify, and surrounding videos on YouTube, Hulu, and other streaming services. Even inventory on digital billboards and other “out-of-home” advertising screens is sold using similar methods.⁹

Perhaps the best-known type of programmatic advertising is the real-time bidding that occurs when internet users navigate to websites and thereby create opportunities to serve ads. In the few hundred milliseconds between the time a user clicks on a link or enters a URL into a browser and the site loads on the user’s device, an auction takes place to determine which ad to insert. This process is facilitated by intermediary technology platforms that circulate information about the impression for sale, execute programmed or learned decision rules, and keep a portion of the money advertisers spend. Publishers—meaning any website, big or small, with advertising inventory to sell—have relationships with intermediaries that help them auction their impressions. For instance, most publishers employ an ad server, almost always run by Google, to manage the routing, insertion, and logging of advertisements. When a user visits a website, the user’s browser contacts the site’s server to request the content; it also contacts the site’s ad server, telling it to populate any available ad slots. Publishers and ad servers often plug into supply-side platforms (SSPs) to access more sources of demand for their advertising inventory than would be feasible with just an internal sales staff. Initially, SSPs were decision support systems that predicted which advertising sales network was likely to return the best price for the publisher’s inventory. These ad networks, operated by companies like Google and Facebook, function similarly to the way sales representatives in other media industries have functioned for generations—simplifying the logistics of making ad inventory from a large number of sellers accessible to many potential buyers.

SSPs now make publishers’ inventory available to buyers by feeding into the next link in the chain—the advertising exchange. Ad exchanges conduct auctions by interconnecting with demand-side platforms (DSPs) that carry out bidding and campaign management functions for advertisers and ad agencies. Ad exchanges contact DSPs with bid requests that convey data about impressions and the users they represent (e.g., location, device type and operating system, and, potentially, a consumer profile attached to a

unique user or device identifier). The DSP tries to recognize and buy the most valuable or efficient advertising opportunities, given its programmed objectives and parameters—such as whom to target, how much to spend, and what outcomes to optimize. As one industry expert put it, DSPs “receive on the order of 1 million opportunities to bid every second. And for each one of those opportunities they make a judgment: how much money is this opportunity worth to [the client’s] brand?”¹⁰ Multiple DSPs bid in programmatic auctions, and in many cases, each of them first conducts an internal auction to determine which client’s bid to submit to the exchange. Ad exchanges quickly collect and adjudicate bids, select winners, and report the outcomes back to the ad servers, which then distribute the appropriate ads into the slots on the websites and trigger billing and record-keeping actions. Small advertisers also participate in programmatic auctions using self-service buying tools that function like DSPs.

Inputs for these valuations and decisions come from a variety of tracking and data integration tools. The cookies and pixels that websites and marketers use to monitor traffic, profile visitors, and measure advertising outcomes are the workhorses of surveillance advertising. A variety of additional tools operate at other layers of the infrastructure, such as application programming interfaces (APIs) and software development kits (SDKs), which enable the circulation of information about what people do on mobile apps and social media platforms. These tools facilitate integrations across data and analytics services, helping to stitch the seams of an ecosystem that aspires to envelop a multitude of sites and contexts.¹¹ Programmatic transactions thus depend on data brokers, data management platforms (DMPs), and identity resolution vendors to connect bits of information held by different parties. For example, many advertisers have first-party data stored in customer relationship management systems; DMPs match those data sets with the information publishers have about their audiences. Essentially, data and identity vendors help marketers determine and recognize the value of profiled consumers and thus calculate probabilities and profit potentials in particular auctions. The same is true on the other side of the auction, as suppliers of advertising inventory use DMPs to discern the likely market value of what they sell. By doing so, they may be able to set an auction reserve price—the “floor” below which they will not sell an impression—to maximize revenue.

The trading process is largely automated, but it still relies on human workers for direction and extensive help. For example, programmatic buying

agencies employ people to specify and adjust bidding parameters, which then guide DSPs' algorithmic decisions. This work involves selecting options from drop-down menus or populating fields in a software interface. Despite the blanketing pretense of rationality, programmatic buyers sometimes exercise informal priorities, such as deviating from prescribed bidding guidelines in order to liquate any unspent portion of a client's budget before the scheduled close of a campaign. Staff at buying agencies also prepare, curate, or review documents that report (with varying degrees of precision, accuracy, and honesty) what clients got for their money. Elsewhere in the pipeline, people work at "cleaning" data sets, maintaining software products and system integrations, and managing business-to-business relationships. Dreams of payroll savings via automation have often been frustrated by the labor required to animate programmatic advertising.¹²

Traditional advertising agencies occupy an important but conflicted place in these markets. Agencies, and particularly media buying organizations, act like investment managers, deciding how to spend clients' money. Many agencies built or bought their own trading-desk systems for buying programmatic impressions. Over the last decade, companies providing advertising and media buying services, including the largest conglomerates, have branded themselves around commitments to data science and artificial intelligence. As the *Wall Street Journal* reported in 2018, "Ad giants such as WPP, Omnicom and Publicis have gone on acquisition sprees, bringing legions of information-technology experts into their ranks."¹³ Caught up in the big-data frenzy of the 2010s, advertising companies (and media companies, too) poached data scientists from other industries and padded their rosters with PhDs in physics, mathematics, and neuroscience.

But agencies are also playing defense. With the profusion of automation and analytics services, other intermediaries have absorbed traditional agency roles, and some advertisers are managing more of their media planning and buying in-house.¹⁴ A few large retailers, such as Walmart and Walgreens, even launched their own DSPs and programmatic marketplaces to take advantage of all the shopping data they collect, including the personal profiles amassed through data-extractive customer loyalty programs. Meanwhile, management and accounting consultants such as Accenture and Deloitte leverage their competencies in data analytics to chase a bigger share of the market for advertising services.¹⁵

The largest adtech companies play multiple roles in this commodity chain. In addition to owning highly valuable advertising inventory on its search engine and on YouTube, Google controls an integrated stack of adtech utilities, serving both the demand side (advertisers) and the supply side (publishers) of the market and serving the two sides simultaneously as an auctioneer. This multilevel control allows Google to exert enormous pressure on advertisers and publishers. Auction design, access to information, bidding speed, selective interoperability, and preferential partnerships all help manifest power asymmetries.¹⁶ The company uses that structural power to drive and lock-in adoption of its tools. For example, advertisers that do not use Google's DSP are disadvantaged by design in certain circumstances. Because its buying and exchange platforms can share exclusive data that is unavailable to rivals, Dina Srinivasan explains, "Google-owned buying tools . . . more frequently make informed decisions about the value of inventory for sale" on Google's ad exchange.¹⁷ Google has also leveraged its near-monopoly in the ad server market to channel publishers' inventory through its own exchange, taking a cut of the resulting trades.¹⁸ There are even allegations that the company intervenes in its auctions to influence outcomes and prices.¹⁹ Essentially, Google provides market-making infrastructures as for-profit services, which ad buyers and sellers depend upon to access routes of trade and to utilize the calculative devices that seem necessary for rational optimization. Other "walled garden" platforms such as Instagram, Amazon, and TikTok also contain dedicated adtech stacks.

The proliferation of adtech intermediaries has divided advertising transactions into sequences of discrete events, each of which is an opportunity for intermediaries to collect fees or gather data that can be commodified or used to render future services. Critics have noted that with so many brokers holding out their greedy hands, a diminishing portion of ad spending reaches "quality publishers."²⁰ Sometimes the programmatic business even aspires to internal perpetual motion. You have likely encountered advertisements that look like clickbait headlines in the "chumbox" placed by companies such as Taboola and Outbrain at the bottom of news pages. These ads, disguised as news recommendations, often redirect the user to "made for advertising" websites that exist solely to monetize this paid traffic by surrounding anemic, outrageous, or plagiarized content with heaps of ads sold by other programmatic vendors.²¹ In many ways, this represents the

culminating logic of a business organized to produce evidence of attention and behavior. Caitlin Petre has aptly likened digital publishers to “traffic factories.”²²

Programmatic advertising is an elaborate and opportunistic solution to the informational and administrative burdens involved in selling the American people. For all its apparent novelty, though, it revolves largely around the same interconnected operations that Oscar Gandy recognized in the discriminatory system he called the “panoptic sort”—identification, classification, and assessment.²³

ADTECH’S PANOPTIC SORT: IDENTIFICATION, CLASSIFICATION, AND ASSESSMENT

There are many methods for identifying people on the internet. The best known is the cookie, a small text file that companies use to track web users and build behavioral profiles. Third-party cookies, left by someone other than the website itself, have been widely used in programmatic advertising to recognize individual consumers and their perceived value whenever and wherever advertising opportunities arise. But third-party cookies are being phased out, leaving many adtech companies scrambling to find new ways to persistently and uniquely identify consumers. Pixels are another tool for tracking and identification. These tiny bits of code collect detailed records of what people do on websites—where they click, what they add to shopping carts, how long they spend on a page, and so on. Facebook has placed conversion pixels on millions of websites to connect evidence of advertising exposure with evidence that the user bought an advertised product. Fingerprinting is a third technique that creates a unique identifier for a device and ostensibly the humans using that device. It is basically a distinguishing inventory of features such as time zone and language settings, hardware and software specifications, how fonts are rendered, and more.²⁴

All these techniques are used to recognize people (or, more precisely, devices and browsers) across situations. But it is no small challenge to assemble the data an individual generates across devices, accounts, and online and offline shopping; associate the data correctly with an identifiable persona; and then recognize that persona when configuring or documenting marketing events. A variety of specialist companies and software systems have emerged to deal with issues of “identity resolution.” Amazon Web Services,

for example, markets an “identity graph” that promises a “360° view of customers.” The graph, which lets marketers “establish persistent identifiers” for individuals, “is used for real-time personalization and advertising targeting for millions of users.”²⁵ This is typical jargon for descriptions of identity solutions. The goal is to make better-than-random guesses about the value and probable behavior of users whose impressions are for sale—although, in many cases, assigning identity to web users is itself a guess generated by a probabilistic model. Identification represents a particular advantage for walled garden platforms where users log into an account.

Identification does not mean that advertisers can recognize a person by name or other human-readable markers. Adtech targets individual users not because of an interest in their personhood but based on some recognition of how a user, or even just a measurable feature the user exhibits, can be classified. A major function of adtech is sorting people into dynamically adaptive categories and populations. Adtech systems classify users, moment to moment, based on patterns or propensities that seem predictive of someone’s responses to certain stimuli and their value for specific marketers. Adtech produces calculated or algorithmic publics, which may exist for only an instant as marketing opportunities appear and disappear or as a system ingests new information.²⁶ The inscrutability of classification decisions—the fact that people have very little insight into why, how, and with whom they are being grouped and targeted—suggests that these publics might as well be called “privates.” This also gestures to the feeling that adtech classifications are dubious and malformed, hidden in proprietary databases where the sun doesn’t shine.

Finally, assessment is the evaluation of individuals and populations to determine the opportunities available to them. Often it means valuation, or assigning prices or estimates of economic worth. Programmatic bidding is an interesting site of assessment. DSPs use algorithms that can be trained to discern the value of users and impressions over time. As an industry consultant explains, a DSP “starts to identify what are the characteristics of impressions that make revenue, and tries to optimize to create the most favorable ratio possible of revenue to cost. That’s what DSPs do. . . . It’s really sophisticated technology that optimizes on the fly to achieve your business objective.” DSPs do this by first bidding randomly “in order to acquire data about which types of impressions lead to good outcomes and which ones don’t,” and then using the “training dataset” compiled through this random bidding to “build its predictive model,” which renders assessments about whether an

impression is likely to yield the desired outcome.²⁷ Moreover, assessments of people and advertising opportunities are adapted to different situations. For instance, someone who is classified as wealthy might fetch rich prices from advertisers selling mutual funds, whereas someone judged to be in a desperate financial situation might be highly valued by predatory companies advertising high-interest loans. Assessments about the worth of these two consumer types, and the likelihood of them clicking on these respective ads, influence what offers they see and perhaps their own sense of self.

This is the bedrock of adtech—assembling machine-readable profit projections to discriminate between good and bad bets.

PROBLEM GAMBLING

Programmatic advertising has been enormously lucrative for some companies. But it exhibits some well-known pathologies. It is notoriously incompatible with any serious definition of privacy.²⁸ Requests for bids on impressions can communicate personal information about a user to hundreds of companies.²⁹ These “bidstream” and other adtech data are sometimes recycled in situations where disclosures of sensitive details can have life-changing or life-threatening implications, such as immigration enforcement, military targeting, and persecution of minority or criminalized persons.³⁰ Right-wing assaults on reproductive and civil rights, for example, are intensifying the danger that the location tracking used for mobile advertising could facilitate state and extrajudicial violence.³¹ The imperative to recognize the value of consumers and advertising opportunities powers a massive business in surveillance, profiling, and data brokering.

Paradoxically, by selling impressions in isolation from the surrounding media content, adtech leads to perverse decisions about advertising distribution. Advertisers’ willingness to delegate bidding and ad serving to black-boxed systems has helped monetize hateful speech and disinformation,³² and it has caused companies to advertise next to content detrimental to their brands.³³ No less perverse, when advertisers *do* dictate the types of content they want to avoid in the interest of “brand safety,” they often withdraw support from news about public health, racial justice, and climate change, as well as from content creators who are deemed “risky” because they deviate from straight, white, middle-class norms.³⁴

These pathologies are endemic to adtech. Programmatic advertising is emphatically not the answer to questions about how to fund democratic communication systems, how to support diverse stories and intrepid journalism, or how to improve user experiences. Programmatic advertising is a solution to logistical and epistemological problems of identification, classification, and assessment. The solution, essentially, is to make advertising markets into machinelike configurations awash in computer-processed claims and predictions, all motivated by the desire to discriminate profit potential more finely, more rapidly, and more massively—to optimize private values.

Understanding that adtech is about recognizing value helps us see that the advantages of scale and the industrial pressures toward consolidation and monopoly are systemic features. Despite the baffling number of companies under adtech's umbrella, an oligopoly rules most of this business. Google, Meta, and Amazon are head and shoulders above the rest of the field, leveraging dominant positions in search and web browsing, social media, and e-commerce. They enjoy expansive views of user behavior, easy and legally permissible access to tokens of user identity, and widespread interconnections for facilitating information flows. Platforms like Facebook record virtually everything known users do within their walls, and they exploit their integration with other websites, apps, and data brokers to connect their own observations with records of what the same users do elsewhere. Furthermore, as intermediaries that process advertising transactions, these platforms have unmatched access to marketplace information, such as bidding patterns and clearing prices.

Whatever the quality and usefulness of big platforms' big data, these companies convincingly present themselves as the actors best positioned to make credible claims about the value of each consumer and each advertising opportunity. This is why "privacy" maneuvers by Google and Apple have partly been power plays to sabotage competitors.³⁵ By preserving their own abilities to collect and cross-reference user data and denying competitors those privileges, these companies gain advantages in identifying, classifying, and evaluating consumers and in attributing marketing outcomes to the events they stage and sell. Marketers gravitate toward these platform companies because they promise, quite persuasively, to measure and optimize return on investment. Bigness in adtech affords truth-making and infrastructural power.

This brings us back to the chapter's central proposition: adtech is about epistemology and logistics. The term *epistemology* seems inapt for systems that may be held together by duct tape or are, in some cases, full-blown scams. But the point is that the business produces evidence, claims, and predictions that organizations and professionals act on. Adtech develops management science machines for enacting and exploiting differences in value and probability. Whether or not they work in ways that are legible to humans and outsiders matters less than the apparent fact that companies trust them to work or feel compelled to imitate competitors for fear of suffering some disadvantage. The knowledge infrastructures that support particular definitions of audience attention and behavior produce what Jérôme Bourdon and Cécile Méadel call a form of "procedural truth."³⁶ They organize the realities admitted into the frame of rational calculation. We often refer to advertising markets as markets for attention, but we should also recognize them as markets for particular types of evidence.

EVIDENT ATTENTION AND BEHAVIOR

Scholars have long understood that commercial media commodify audience attention.³⁷ How that works, basically, is that advertisers pay companies that control the means of communication to execute distribution events, delivering commercial messages to a specified number and type of potential consumers. Advertisers hope audience members will interpret those messages and, better still, act in some desired way as a result. Dallas Smythe argues that, as people decode advertisements, they participate in their own socialization into lifestyles and purchasing habits that reproduce their labor power and capitalist social relations.³⁸ The details of this process are complex and contradictory, but Smythe's contention is useful for recognizing how commercial media are organized to produce consumers and consumption in action. This side of the "attention economy" is about how allocations of consciousness and cultural energy make and remake social worlds. It reminds us of the prompt Harold Innis left for media studies: "Why do we attend to the things to which we attend?"³⁹

At a more prosaic level, commercial media also produce consumers and consumption in information—as evidence, documents, statistics, and accounts. These operations are actually more material to the business of monetizing

attention (including via fraud), and they demand an understanding of attention that is sensitive to its conditions of commodification—that is, producing and packaging observable events that can credibly be claimed to represent human cognition or behavior. Advertising transactions revolve around particular ways of defining, operationalizing, and materializing media usage. Audiences and attention, in this context, are informational products, assembled through historically specific conventions, instrumentations, and property regimes. Manufacturing these products involves making certain user activities and attributes available for measurement and maintaining industrial relations that make the extracted data intelligible and legitimate to those who would trade or act on it, including machines. Salable attention exists via the platforms and devices that manifest it and the methods of processing it into a discrete, tangible thing. What's sold is *evident attention*.

Historically, evidence of attention has meant measures of media exposure manufactured by third-party research firms that act as “neutral” intermediaries between advertisers and attention merchants. In television, for example, Nielsen gathers behavioral records from a sample population and then uses those panel data to produce “ratings” representing larger viewing publics. Ratings have functioned as an authoritative “currency” for setting prices and converting scattered and endlessly diverse audience experiences into an actionable collection of standardized marketplace facts. Panel-based audience measurement has continued on the internet,⁴⁰ but digital media also allow the companies serving advertisements to log the details of which ads they deliver to whom, as well as to observe clicks, downloads, mouse movements, and other indices of relevant behavior.⁴¹ Advertising intermediaries have exploited the technical fact that the process of serving digital ads can automatically produce documentary records, including data about users.⁴² Platforms like Facebook single-handedly execute advertising events and generate the corroborating evidence of attention or behavior. Because mistrust can arise when these audience aggregators “grade their own homework,” many advertisers and publishers employ third-party verification companies, such as Integral Ad Science and DoubleVerify, to validate reported metrics and confirm that ads were loaded and viewed correctly. Despite the capacity for comprehensive record keeping, truth in digital advertising remains contestable, with companies making competing claims about ostensibly identical realities. And given that evidence of attention

can be easily fabricated (often falsely, accidentally, or by internet bots), companies have their hands full sifting a torrent of evident attention to distinguish the legitimate from the illegitimate.⁴³

Adtech is the name now given to “new media” infrastructures for materializing attention as a marketable commodity. “New” should be interpreted less as a matter of recency than as one of remediating relationships. Josh Lauer explains that a defining feature of “new media,” in whatever historical context, is that they “produce new types of evidence, each with its own material form and truth claims.”⁴⁴ Today’s adtech represents the latest configuration of audience manufacture as an evidential paradigm, negotiated around the affordances of evidence-producing media, the capabilities of information systems to store and analyze evidence, and the uneven power to authorize certain evidence and enforce certain claims. “The semiotic detritus of new media, past and present,” Lauer writes, “has contributed to an intensification of surveillance by introducing new forms of evidence . . . by which individuals might be identified, their motives and thoughts inferred, and future behavior predicted.”⁴⁵ Adtech exemplifies this accommodation of new media and their evidence into audience manufacture.

This discussion is not meant to deny the lived practices of people making meaning with media. Rather, it highlights some specific ways that audience attention and behavior exist as objects of knowledge and commerce. The commodity audience resembles, in some sense, what information scientists call a “document.” Librarian Suzanne Briet famously distinguished an antelope in the wild from one in a zoo. The caged antelope becomes a document; it can be used as a type of evidence, or index, testifying to the fact of antelope in the wild.⁴⁶ The enclosed commodity audience has a similar existence; it is documentary evidence, assembled within an organized discourse, attesting to the fact of attention or behavior. Further examination of people—where they go and what they buy—yields additional evidence about their worth and inclinations, which is drawn into the frame to value units of attention when they come up for sale. Advertising has always revolved around “centers of calculation,” where expert communities work to collect measures of attention, behavior, and value and translate them into institutionalized claims.⁴⁷ Adtech is a bustling center of calculation and commerce for today’s evident-attention economy.

Here again, the singular noun overstates adtech’s uniformity. Different business models, such as online search and social media, configure evident

attention and behavior in their own ways, exploiting capacities and bottlenecks, with different implications for the organization of user experiences and the exercise of power.⁴⁸ The logic of selling the American people is implemented with variation, activating a range of adtech tools, techniques, and relations. Audience manufacture provides a fascinating case for looking at how documents and information technologies represent the market to buyers and sellers and thereby structure the basis of economic thought and action. What media researchers refer to as the “institutionally effective” audience—evidence of attention that becomes meaningfully incorporated into organizational routines—is constructed through a process of “framing.”⁴⁹ Framing sets boundaries around what objects and relationships are admitted into economic calculation and exchange, what is inside or outside a market or a transaction. Framing establishes the portion of reality taken into account and how it is “known” and represented. Because audiences and attention can be operationalized in myriad ways, this domain is open to constant crisis, opportunity, and negotiation, although in practice it tends to be fixed, at least temporarily, by institutional and infrastructural closures. Perceived openings—created, for example, by new ways of generating evidence or processing transactions—trigger contests among actors trying to cement a more profitable closure.

The historical construction of adtech illustrates these dynamics. To give perhaps the most significant example, digital advertising has taken shape around pressure to pay for evidence not of message exposure but of the marketing outcomes advertisers want to achieve. The long thrust of accountability has, in many ways, been about nudging the business to integrate evidence of *audience* behavior with evidence of *consumer* behavior (see chapter 8). Not surprisingly, then, adtech is part of a bridge between marketing and marketplaces. Digital advertisements almost always offer opportunities for e-commerce purchases. Leading advertising platforms such as Instagram, YouTube, TikTok, and, of course, Amazon are trying to facilitate more consumer purchasing within their own sites or apps. They are actively marketing themselves to advertisers, influencers, and other content creators as digital storefronts—self-contained shopping portals where users can buy whatever they see in ads and user-generated content. Through various partnerships, these platforms facilitate payment processing and order fulfillment in a seamless consumer experience.⁵⁰ The idea that digital media are shoppable media is taken for granted in advertising and cultural industries around

the globe.⁵¹ The production of consumers is treated, here, as a technical achievement of ubiquitous connectivity to online marketplaces. As Mark Andrejevic puts it, “the entire networked environment is a space of consumption (and labor) thanks to the affordances of the digital enclosure.”⁵²

Enthusiasm for shoppable media reflects the dream of advertising attribution. Attribution is essentially a claim about whether or to what extent certain advertising events influenced consumers’ behaviors or caused other relevant outcomes. These claims are often probabilistic, based on models that fabricate relationships between advertising events and marketing results through a mixture of observation, correlation, and estimation. And their limits get redefined by technical, legal, and cultural changes. The recent rush by Apple and Google to ensure privacy is impairing marketers’ ability to track the performance of their ads in generating sales. For example, Apple’s new restrictions on cross-app tracking “makes it harder to measure when an ad on one app leads to a sale on another.” This is pushing platforms to further develop their own e-commerce functionality. According to a writer in *Advertising Age*, “Without the ability to follow that digital trail [across multiple websites or apps], social media companies want consumers to go from discovering products to purchasing directly within their platforms. This would allow the platforms to prove their value, and retain customer data, too.”⁵³ Meanwhile, as media platforms are becoming shopping portals, retailers are becoming adtech companies. Limitations on third-party data sharing create opportunities for large merchants to exercise their massive customer data sets and their ability to document purchases by individual shoppers.⁵⁴

Whatever solutions marketers and marketing platforms develop to bolster their attribution claims and adapt to new privacy requirements, the underlying motivation will be to generate credible and institutionally effective evidence of audience attention and consumer behavior. That is the grist in adtech’s money mill. Importantly, though, that mill grinds a messy mash of prediction and intervention. It can be hard to tell whether apparently successful digital advertising caused sales or intercepted them. Adtech lets intermediaries stand beside the cash register and take credit for events that might have happened anyway. In other words, attribution may not measure the influence of an advertisement so much as an intermediary’s success at recognizing that someone was on their way to buying something. By betting on that propensity and documenting its consummation, adtech fabricates evidence of value, justifying service fees.

CONCLUSION

In one of the best-known histories of US advertising, Stephen Fox describes a recurrent cycling between two paradigms: one that creates cultural *imagery* about brands and their users, and one that issues *claims* about products, often appealing to science and rationality.⁵⁵ With adtech—and, as I argue, with broader processes of buying, selling, distributing, and evaluating advertising—it's claims all the way down. Or, rather, it is claims, conjectures, predictions, classifications, and investment decisions, all held to account by an avalanche of statistics and documents. A key task, then, is to uncover how certain types of claims and certain types of claim makers gain or lose legitimacy. Although the regimes of knowledge and power that operate in adtech today are novel in many ways, the fact that they were recognized as offering convincing solutions and desirable opportunities for advertising must be explained through longer histories of technology (including hype and failure), political economy, and business culture. Ironically, because it is so difficult to vet many of the claims at the heart of digital advertising, the backstop for judgment often rests on a company's ability to wield science and technology within its brand image.

A focus on claims helps us position adtech within the development of informational, surveillant, and technoscientific elements in capitalist marketing. The rest of this book can be considered a contribution to what Adam Arvidsson calls a "prehistory of the panoptic sort." Arvidsson and other historians point to a shift in marketing around the 1950s, as computerization, segmentation, manufacturing capabilities, and innovations in statistical methods all facilitated a focus on specialized populations and a commitment to tracking and adapting to more flexible patterns of consumption.⁵⁶ I suggest, in particular, that this paradigm was inflected by management science, which is, in many ways, a set of techniques for enacting differences of value in the service of optimization. Adtech promises automated management science, combining logistical infrastructures with quantitative formalisms for decision making and evaluation. When we talk about programmatic advertising or the business mechanics of an advertising platform like Facebook, that is what we are talking about.

This introduction to adtech has focused on logistics and epistemology. Adtech applies techniques of identification, classification, and assessment to recognize costs and profit opportunities, to predict behavior, and to

act on those predictions. My critique, therefore, starts by defining adtech as ostensible solutions to problems of distribution and management. This is different from attacks on adtech that gather around opposite poles: at one end, adtech is an irresistible manipulative force; at the other end, it is pure marketing fiction. Neither of these positions captures what's really happening: rival computer systems dueling over pennies and probabilities. The issue is not about mind control. It is about the accumulation of power by companies that administer bottlenecks and spin all sorts of evidence into claims about value and the future. Certainly, many advertisers and adtech companies dream of controlling consumer behavior. The expansion of surveillance and the proliferation of dynamic choice architectures used for "nudging" are both part of marketers' efforts to manage consumption. But, for the most part, adtech materializes elaborate designs to engineer, orchestrate, and capture value from new investment opportunities.

Management science's major contribution to adtech is formalizing the decisions involved in distributing ads and allocating investments, thus translating ad buying, selling, and serving into operations that can be acted on by computers and algorithms. A critical insight from this history is that, in chasing dreams of optimization, advertisers and their agents have been required to state explicitly what they care about and how they value different types of consumers and different types of advertising opportunities. For all its complexity and opacity, adtech exposes a simple truth: advertisers are trying to maximize private economic value, not democracy, justice, or joy in public life. Equitable politics and lively arts and culture are not the objective functions that advertising models are designed to optimize. A close look at how advertising came to emulate management science shows that counting on marketers to fund truly social media is like counting on a wish and a prayer.

This points to a key conclusion. While adtech today is undeniably grotesque, it is an embodiment of deeper problems. It is a set of strategic responses to the prospects and challenges arising from within the business of selling to the American people. Over the second half of the twentieth century, a variety of actors tried using technoscience to optimize the twin goals of producing consumption and commodifying evident attention. Optimization has been such a productive ideology for advertising because it implies a definitive achievement of the best possible outcome, yet it is endlessly deferrable, open to renegotiation whenever conditions seem to change. Even as privacy rules restrict the field of play in adtech, for example, companies

will promise to deliver whatever new optima are possible under the circumstances, likely leaning to an even greater extent on claims of technoscientific and logistical power.

Adtech represents a particular expression of what I call advertising's calculative evolution. Despite the breathless talk of historical rupture surrounding digital media, the discourse related to optimization in advertising has displayed remarkable consistency over time. The rest of this book reveals how the world performed by contemporary adtech resembles the futures imagined by certain advertising professionals at moments when information technologies appeared to open paths for progress, permitting better optimization—new claims of *more perfect*.

© 2023 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.
Subject to such license, all rights are reserved.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: McGuigan, Lee, 1986– author.

Title: Selling the American people : advertising, optimization, and the origins of adtech / Lee McGuigan.

Description: Cambridge, Massachusetts : The MIT Press, [2023] | Series: Distribution matters | Includes bibliographical references and index.

Identifiers: LCCN 2022037239 (print) | LCCN 2022037240 (ebook) | ISBN 9780262545440 (paperback) | ISBN 9780262374231 (epub) | ISBN 9780262374248 (pdf)

Subjects: LCSH: Internet marketing—United States. | Internet advertising—United States. | Advertising—United States. | Digital media—United States. | Social media—United States.

Classification: LCC HF5415.1265 .M367 2023 (print) | LCC HF5415.1265 (ebook) | DDC 658.8/72—dc23/eng/20230103

LC record available at <https://lcn.loc.gov/2022037239>

LC ebook record available at <https://lcn.loc.gov/2022037240>