#DIGITALDECEIT

The Technologies Behind Precision Propaganda on the Internet

DIPAYAN GHOSH & BEN SCOTT

JANUARY 2018

About the Authors

Dipayan Ghosh is Fellow at New America in Washington, D.C. and Joan Shorenstein Fellow at the Shorenstein Center on Media, Politics and Public Policy at the Harvard Kennedy School, where he works on digital privacy, artificial intelligence, and civil rights. A computer scientist by training, Ghosh until recently worked on global privacy and public policy issues at Facebook, where he led strategic efforts to address privacy and security. Prior, Ghosh was a technology and economic policy advisor in the Obama White House. He served across the Office of Science & Technology Policy and the National Economic Council, where he worked on issues concerning big data's impact on consumer privacy and the digital economy. Ghosh received a Ph.D. in electrical engineering & computer science at Cornell University and a bachelor's in the same field from the University of Connecticut.

Ben Scott is Senior Advisor at New America in Washington, D.C. and serves on the management board of the Stiftung Neue Verantwortung, a technology policy think tank in Berlin. Previously, he was Policy Adviser for Innovation at the US Department of State where he worked at the intersection of technology and foreign policy. Prior to joining the State Department, he led the Washington office for Free Press, a public interest organization focused on public education and policy advocacy in media and technology. Before joining Free Press, he worked as a legislative aide handling telecommunications policy for then-Rep. Bernie Sanders (I-Vt.) in the U.S. House of Representatives. He holds a Ph.D. in communications from the University of Illinois.

About New America

New America is committed to renewing American politics, prosperity, and purpose in the Digital Age. We generate big ideas, bridge the gap between technology and policy, and curate broad public conversation. We combine the best of a policy research institute, technology laboratory, public forum, media platform, and a venture capital fund for ideas. We are a distinctive community of thinkers, writers, researchers, technologists, and community activists who believe deeply in the possibility of American renewal.

Find out more at **newamerica.org/our-story**.

About the Public Interest Technology Program

New America's Public Interest Technology team connects technologists to public interest organizations. We aim to improve services to vulnerable communities and strengthen local organizations that serve them We are engineers, designers, product managers, and researchers. Our approach starts and ends with user needs. We believe in humanity, humans, and human-centered design. We design and deploy technology that serves people and solves problems, not technology for technology's sake.

Find out more at **newamerica.org/public-interest-technology**.

Acknowledgments

Contents

INTRODUCTION

Revelations last summer that a Russian-backed disinformation campaign used major U.S. internet platforms to interfere with the 2016 presidential election sent shockwaves across the nation. Evidence has since emerged from intelligence findings, investigative reports, and corporate forensics providing a more complete picture of how this happened. Facebook, Google, and Twitter have each now confirmed that Russian agents coordinated efforts to disrupt the American political process by spreading divisive messages over their platforms.[1] On Facebook alone, Russian influence reached over 125 million users.[2] But much of what happened still remains unexplored.

The intersection of political disinformation and internet platform technologies has spawned not only a public policy debate but also a broader national conversation about the integrity of the American democracy. The scale, opacity, and influence of political communications on the largest digital platforms have resulted in tremendous public scrutiny of the leading U.S. internet firms. Top lawyers from Facebook, Google, and Twitter testified before the U.S. Congress in a series of globally publicized hearings in late 2017.[3] During the course of those wide-ranging and often acrimonious sessions, congressional leaders made clear their view that these companies should have done more to detect and prevent Russian exploitation of their platforms.[4] Alongside their alarm was the recognition that what is known about Russian activities may only be the tip of the iceberg of political disinformation.

The companies have promised to take voluntary action. Among the most widely discussed plans are those to enforce the labeling of political advertisements shared on their platforms; to

inform users who have been exposed to Russian disinformation operations; and to create a searchable public database of digital advertisements including their content, sponsor, and targeting parameters.[5] More recently, Facebook announced plans to overhaul their flagship News Feed algorithm to prioritize content shared by friends and family and downgrade text and video posted by brands and news organizations. They also plan to survey users' trust in news sources, decreasing the visibility of untrusted outlets.

These are important, positive steps. But they will not be enough. This deep-rooted problem goes far beyond the scope of the current debate.

In this paper, we deliver a deeper examination of internet-based advertising and media platforms, and how they enable the dissemination of political disinformation. At present, the focus is largely on the role of Russian agents and the advertising technology stacks developed and operated specifically by Google, Facebook, and Twitter. However, the digital tools available to

disinformation campaigners are far from limited to the services offered by these three firms. These platform companies are at the center of a vast ecosystem of services that enable highly targeted political communications that reach millions of people with customized messages that are invisible to the broader public. In this analysis, we examine the entire toolbox of precision propaganda including:

- Behavioral data collection

- Digital advertising platforms

- Search engine optimization

- Social media management software

- Algorithmic advertising technology

This combination of interconnected tools is a brilliant technological machine that serves to align the economic interests of advertisers and the platform companies. The more successful the advertising campaign, the more money everyone makes. In this marketplace, all advertisers are essentially alike, whether they are pushing retail products, news stories, political candidates, or disinformation. When it comes to the application of these tools, all advertisers seek to emulate the most successful persuasive strategies. That means all the tools of behavioral data collection available for the purpose of targeting communications into highly responsive audiences—i.e. pre-filtered segments of demographically similar people that are easier to

> **Google, Facebook, and Twitter are at the center of a vast ecosystem of services that enable highly targeted political communications that reach millions of people with customized messages that are invisible to the broader public.**

engage and persuade—are applied to the task of political disinformation.

The problem is that when disinformation operators leverage this system for precision propaganda, the harm to the public interest, the political culture, and the integrity of democracy is substantial and distinct from any other type of advertiser. Our thesis is that we must study the entire marketplace of digital advertising and disentangle the economic alignment of interests in order to find the best ways to constrain bad actors and minimize harm to the public.

Doing so presents a more complex, and perhaps more disturbing, picture of the problem. Absent this wider perspective, we cannot prepare policies to effectively deter disinformation operations.

The incident of Russian disinformation operations that sparked this controversy illustrates these assertions. The technologies the Russians employed are industry-standard digital advertising and marketing tools that can be readily adapted to any disinformation campaign operated by any actor. It is tempting to view Russian digital propaganda exclusively through a national security lens, rather than a political economic one. But political disinformation constitutes a public harm regardless of the propagator, and most disinformation in America is distributed by domestic actors.[6] This more general problem is far more difficult to solve, not only because what constitutes "disinformation" is often contested, but also because many disinformation activities are likely entirely legal— and even protected under the First Amendment— in the United States. This creates a tremendous challenge with which internet platforms must actively grapple.

Based on the analysis presented here, we believe effective solutions must address the political economy of digital information markets. The financial interests that drive the core technologies of the leading internet platforms and the objectives of disinformation campaigners are often aligned. This is clearly not welcomed by the platforms,

but the shared interests are nonetheless present because the form of the advertising technology market perfectly suits the function of disinformation operations. These campaigns often deploy sensational themes and polarizing politics. This content draws and holds consumer attention, which in turn generates revenue for internet-based content.[7] A successful disinformation campaign delivers a highly responsive audience that drives forward engagement on the platform and ultimately delivers more revenue for all parties.

This problem cannot be solved by simply blocking content attributed to Russian agents, or otherwise playing a game of censorship whack-a-mole with blunt-force attempts to delete particular content or bar particular speech. Such approaches will fail because these systems are highly adaptable and ephemeral, and even a well-resourced multinational will not be able to whack the moles fast enough. But more importantly, it will fail because much of the disinformation is legally protected political expression. Transparency about the sponsorship, reach, and targeting parameters of advertising will help. But transparency is most effective when it exposes a small number of bad actors whose behavior stands out in contrast to everyone else. If transparency reveals that the problem is systemic,

it could have the effect of normalizing all but the most egregious behavior. Will it create disciplinary pressure on disinformation operators if a database of online advertisements that is infrequently viewed by users demonstrates that propaganda is epidemic on the platforms?

The path to finding better remedies necessitates a deeper understanding of the problem at hand. This paper seeks to broaden the focus from an analysis of digital ad buying to an analysis of the entire market for products and services related to driving sentiment over the internet. Our purpose here is to document the argument that political disinformation succeeds because it follows the structural logic, benefits from the products, and perfects the strategies of the broader digital advertising market. If that is correct, then we need to understand exactly how this ecosystem works. We need to consider how to differentiate political from non-political campaigns on internet platforms in order to focus on the kinds of practices that have an impact on democracy. We need to find creative ways to empower citizens to discern, expose, and discredit media manipulation. And we need to think critically about how to solve not just today's problems of disinformation, but those of tomorrow.

> **When disinformation operators leverage this system for precision propaganda, the harm to the public interest, the political culture, and the integrity of democracy is substantial and distinct from any other type of advertiser.**

BEHAVIORAL DATA COLLECTION

Data is the lifeblood of online commerce. Every post, click, search, and share is logged to a user profile, grouped into a segmented audience, and fed into machine learning algorithms. This data allows advertisers to infer an individual's preferences, behaviors, and beliefs—all of which inform highly targeted digital advertising campaigns. The power of accumulated data is also the driver of disinformation campaigns that leverage the social graph of individuals not to drive purchasing decisions but to influence sentiment, political views, and voting behavior through precision propaganda.

Advertising is at the center of the internet economy, and its value is increasingly derived from the collection of behavioral data used to segment and target audiences. The business has evolved rapidly from a digital version of conventional ad placement involving agencies and publishers, to what is now a data-driven market focused on audience segmentation and targeted messaging.[8] Algorithmic technologies determine the content, timing, and consumer audience for the delivery and display of online advertising.[9] By vacuuming up huge quantities of personal information, firms across the digital advertising ecosystem—be they social media platforms, advertising exchanges, programmatic advertising agencies, or brand companies—can begin to infer an individual's preferences and predict responsiveness to different kinds of ads.[10]

It cannot be understated how important personal data is to the long-term sustainability and success of the digital advertising ecosystem. Data drives commerce on the internet; every consumer-facing internet company that has a major presence in online advertising collects and shares information about individuals to help their advertising clients

succeed.[11] It is similarly true for political campaigns that use voter files and demographic data to construct complex profiles of constituencies. And, it is also true for political disinformation campaigns.

While there is nothing technologically novel about tracking and profiling internet users for ad targeting, these practices nevertheless take on a different valence when the intelligence gleaned from behavior monitoring is used not to sell products but to manipulate and deceive voters. But how do companies in the digital advertising ecosystem collect this data? In the rest of this section, we will detail three categories of technologies in widespread use today that enable identification and profiling.

Web Tracking

When an internet user navigates to an article at a news website—www.bbcnews.com, for instance—hundreds of digital transactions take place behind the scenes. When a web browser retrieves and

loads a page, it delivers much more than just the substantive content. Most obviously, it loads third-party advertisements that have paid to appear on that page. But more importantly for our purposes, and invisible to the user's eye, the page loads any inbuilt web tracking technologies. The most basic tracking technology is the standard web cookie.

Web cookies can either be "first-party" or "third-party." Typically, first-party cookies are developed and placed by the owners of the website itself, while third-party cookies are developed and placed by other entities in partnership with the website owner.[12] Often, websites deploy cookies to help the user seamlessly log in to the page and load an account. For instance, Barnes & Noble or Amazon might utilize a first-party cookie to help the online shopper log in to the site and maintain a real-time list of the items in a user's digital shopping cart.

Just as often, however, websites deploy cookies, whether first- or third-party, to track a user's activity on the site and potentially infer personal routines and behaviors.[13] This kind of behavioral tracking is typically done so that the website or its affiliates and associates can, over time, infer such things as the preferences, interests, and beliefs of the user. That information can then be used, among other things, to target users for specific advertisements—about, say, the newest line of Hondas or Chanel perfumes, depending on personality and behavior as exhibited through past and real-time internet use.[14]

> **The more behavioral data they are able to collect on users, the better they are able to serve them targeted ads that cater to their unique interests.**

Behavioral data is stored by companies in large databases that record user activity over time. This data can be linked to an IP address, the unique identifier assigned to a particular computer or mobile device. However, many people often
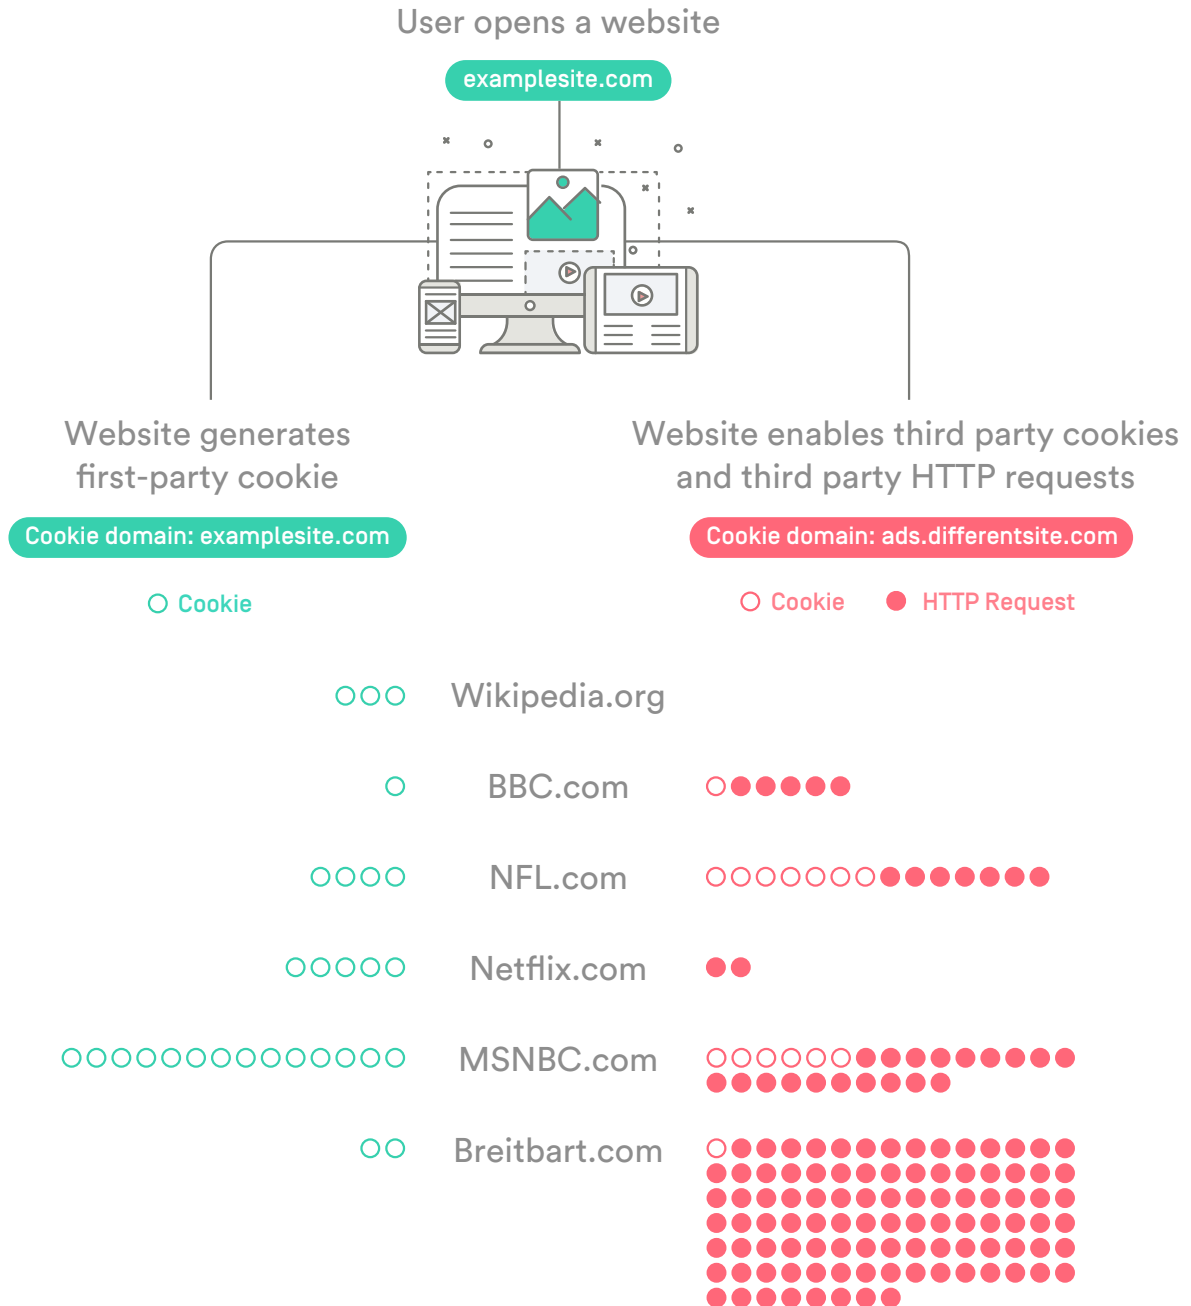
volunteer more personal information by making a credit card purchase, providing an email address, or connecting with that site using their personal online accounts, thus tying their behavioral data to their individual profile more closely.[15] Some companies use this data exclusively for marketing their own products and services. Others sell it to data brokers that combine it with other data and sell it along to other marketers. Over time, behavioral data profiles can become large and detailed.

For internet companies that operate at a global scale, including search engines and social media platforms, there is tremendous potential value in maintaining behavioral data on users. There is a virtuous cycle in the collection of behavioral data for two primary reasons. First, the more behavioral data they are able to collect on users, the better they are able to serve them targeted ads that cater to their unique interests. Second, the better they are able show the most relevant content to the user, the longer they can keep that user on the platform, thereby maximizing potential advertising space for the user. This vertical integration of the practice of behavior tracking and the business of online advertising is central to the market power of global internet platforms.

In the case of platforms like Google, Facebook, and Twitter, user engagement with digital content—including ads, page likes, clicks on individual search results, or interactions with news feeds—can be recorded and compiled into behavioral data profiles, which can further empower these companies to target individual users with the content and ads most relevant to them. Typically, this type of behavioral data is privy only to the platform on which those interactions occurred. For reasons primarily of user privacy, competitiveness, and protection of intellectual property, it is in the platform's interest to share the data with no one—including the original user whose actions generated the data. Among others, Google and Facebook have achieved tremendous commercial success through this vertical integration of behavior tracking and ad targeting.

Comparing Cookie Use on Major Websites

User opens a website

examplesite.com

Website generates
first-party cookie

Cookie domain: examplesite.com

○ Cookie

Website enables third party cookies
and third party HTTP requests

Cookie domain: ads.differentsite.com

○ Cookie　● HTTP Request

○○○	Wikipedia.org	
○	BBC.com	○●●●●●
○○○○	NFL.com	○○○○○○●●●●●●●
○○○○○	Netflix.com	●●
○○○○○○○○○○○○	MSNBC.com	○○○○○○●●●●●●●●● ●●●●●●●●●
○○	Breitbart.com	○●●●●●●●●●●●●●● ●●●●●●●●●●●●●●● ●●●●●●●●●●●●●●● ●●●●●●●●●●●●●●● ●●●●●●●●●●●●●●● ●●●●●●●●

Data retrieved on January 22, 2018 with cookie-checker.com.

Because their purpose is tracking online activity and behavior, cookies are regarded by many as severely damaging to the typical internet user's privacy. In Europe, for example, regulatory authorities have required that websites flash a browser warning to visitors if their site uses cookies.[16] In light of this, some internet users choose to delete them. But websites and advertisers have over the years caught on to this; many websites deploy technologies called local shared objects (LSOs), also known as Flash cookies, to reinstall cookies that have been deleted by the user.[17]

Another common practice is embedding cookies in emails. Email cookies enable organizations to determine when and where an email was opened. Organizations deploying email tracking in this way can additionally determine the type of device on which the email was opened. These services are often implemented in much the same way that standard web-based cookies are deployed through insertion of a 1x1 pixel. The tracking service provider can see if this pixel has been downloaded on the email recipient's device, or if it has been forwarded and downloaded onto other devices. These services are offered through ready-made software packages developed by a wide variety of firms.[18] A recent report from one of these companies estimated that over 40 percent of all 269 billion emails sent every day is tracked.[19]

Apparel outfits like Nike or retailers like Macy's may have an interest in email tracking to understand the level of customer engagement that different types of marketing emails receive. But perhaps more interestingly, web-based services like Amazon and Facebook have widely deployed email tracking, including through the emails they send their users to increase user engagement with their services. This form of email tracking can contribute to the vast stores of information these companies have about their users. Perhaps most critically, email tracking can help refine the service providers' understanding of the user's location habits.

Another internet tracking tool is called the web beacon. Also simply called beacons, these tools are programmed objects that are inserted into a web page, typically by a third-party partner that has paid for the privilege. When a user loads the page, the beacon loads too. The third party operator of the beacon receives the signal that a user has visited the website, also known as a server call. Upon receiving the call, the third party sends the web beacon to the site owner so that it can be rendered on the user's browser.[20]

But the key to web beacons is that they are typically invisible or otherwise hard-to-detect for the user, because their purpose is not to serve content, but rather to know that the user has visited the page. Depending on what that page contains, the third party can begin to make inferences about the user's preferences using web beacons, especially as they are often used to track real-time user actions on the page—such as a mouse clicks, hand movements, or hovering cursors. An advertising exchange that has installed web beacons across a network of websites can wield extraordinary power as it amasses a wealth of information after tracking the online movements of any visitors to those pages. All of this data can be associated with an individual's IP address, personally identifying information, or persistent identifier. As advertisers gain access to that data to inform and drive their ad campaigns, they can execute highly targeted outreach to individuals based in part on data gathered from web beacons.

Location Tracking

Modern smartphones are tremendous technological marvels. They are also extraordinarily effective tracking devices. Each handset contains a satellite-powered technology developed by the U.S. government for military applications—the Global Positioning System. The GPS technology involves a network of over 30 satellites in orbit around the Earth, each of which beams its real-time position through a simple communications protocol to Earth's surface. So long as a GPS-integrated device can receive positioning signals from at least four of

GPS Location Tracking on Modern Smartphones

In 2011, two independent UK-based researchers discovered a file on the iOS called consolidated.db

consolidated.db
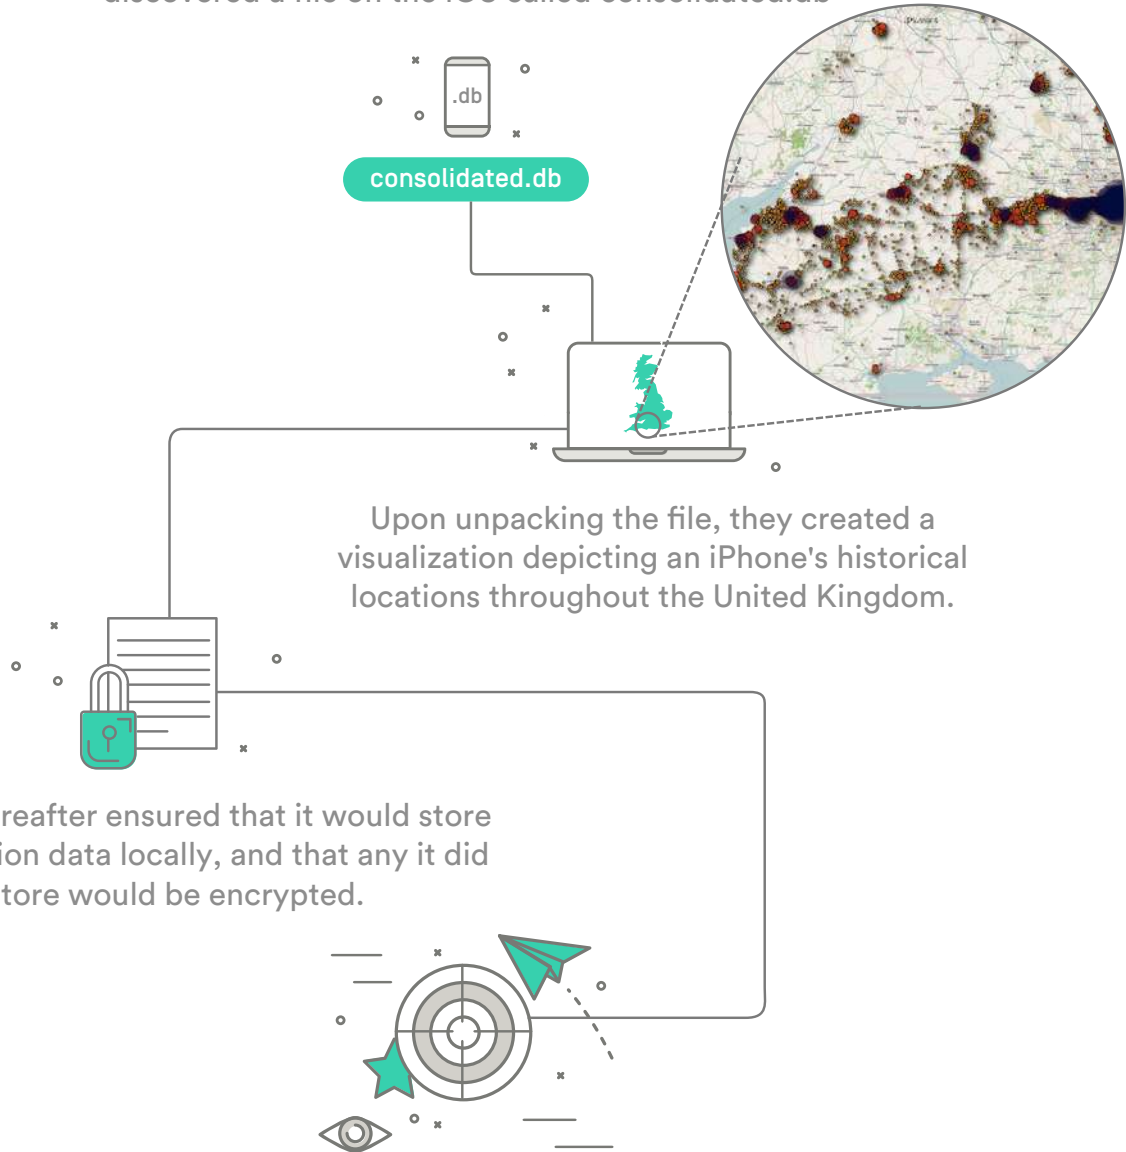
Upon unpacking the file, they created a visualization depicting an iPhone's historical locations throughout the United Kingdom.

Apple thereafter ensured that it would store less location data locally, and that any it did store would be encrypted.

The finding reminded the public just how sensitive GPS location data could be, particularly as a company that possesses it could share it onward with data brokers, the advertising industry, and government agencies.

the GPS satellites, the device's unique position can be triangulated with great accuracy and precision.[21] This is the system that powers the smartphone mapping applications that are widely used and pre-loaded on nearly every device.

GPS location is not the only type of location data that phone manufacturers and system operators collect. Particularly in urban areas where GPS signals may be blocked by buildings, this data is often combined with cellular network triangulation, Wi-Fi SSIDs and Bluetooth connectivity to pinpoint the physical location of a user. Fine-grained, highly resolute location data is a valuable component of the advertising technology ecosystem because it can reveal a great deal of useful context for a person's interests, preferences, beliefs, and behaviors. Consider the types of information that can be gleaned from address-specific location data. A smartphone manufacturer or developer of a mobile operating system can predict with great confidence where a person lives, where he works, how he gets to work, who he spends time with, where and how he spends time with those friends, which brick-and-mortar business establishments he frequents, and what he does in his free time whether he plays baseball, watches movies, visits the local bar, or canvasses residences on weekends.[22] This information about routine activity provides important signals about how to deliver targeted digital ads optimized to his personal interests.[23] Notably, the U.S. Supreme Court is presently considering a case that will determine whether law enforcement agencies must treat a person's location data as so sensitive that it requires a warrant to access.[24]

Location data is captured by numerous commercial parties, which may include a consumer's network operator and device manufacturer[25], as well as the operators of a individual's mobile operating system and applications. Granting access to location data has become a standard part of enabling many popular mobile services, not limited to obvious ones like mapping services. Many modern apps—navigation, shared ride, gaming, and social media applications among them—do not include full

functionality unless access to location data is granted. Users are often pressured to share their location data with a range of organizations, which can further opt to share user data or inferences gleaned from it with a variety of other third parties.[26] Users typically have no knowledge of this sharing of data.

Location service, including GPS functionality, is typically an opt-out service on smartphones. Users can turn it off through their mobile operating system settings. But this is uncommon and often reversed because many popular applications like Yelp or Uber persistently request access to location data if the user disables location service or suspends the sharing of location data with those particular applications.[27]

Once a user has activated location services and shared location data with applications, a series of further privacy issues can arise. Though applications ought to collect location data solely to enable the technical functionality of their service, there are no clear rules or regulations over the collection of location data in the United States. Instead, companies usually need only do what they commit to their users in the fine print of a privacy policy. This leads to situations where some mobile apps ask for or even require location data even though it is unclear how that data would benefit the functionality of the app.[28] This is particularly alarming given the thriving data broker ecosystem and huge demand for historical location data. And even if users do disable location services, mobile phone operating systems like the iOS and Android may still track location, including through triangulation of precise location using cell tower data.[29]

Newer methods of tracking an individual's location will likely continue to emerge in the coming years, both because of advances in technologies like virtual reality and artificial intelligence, but also because of run-of-the-mill oversights in security and data access. In late 2017, one researcher illustrated how any app that gains access to an iPhone user's camera roll can also see where each photo in the
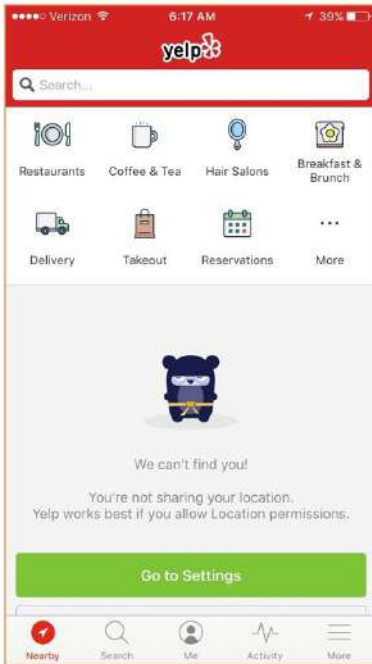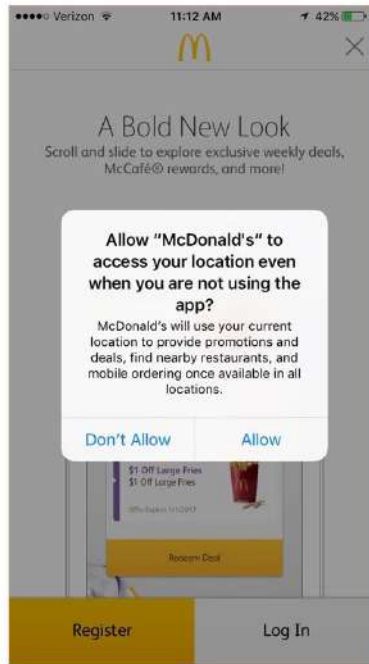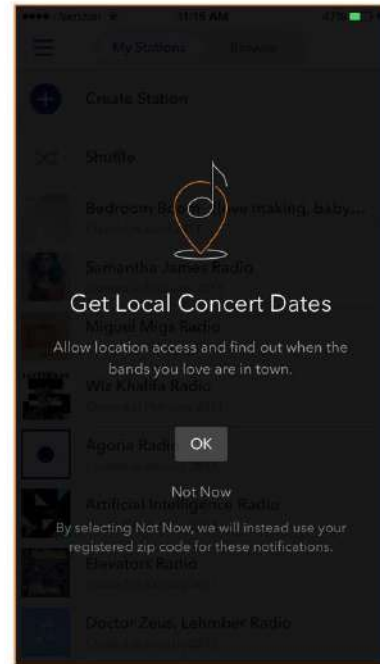
Yelp	McDonald's	Pandora

Location requests from Yelp, McDonald's, and Pandora mobile applications.

camera roll was taken if the user has enabled geotagging of photos.[30]

Cross-device Tracking

Consumers in the United States typically access the internet from not just one but several devices around the home and at work. Last year, the average North American household had more than seven internet-connected devices.[31] For the digital advertising industry, this situation can be problematic. Advertisers often do not want to send duplicate advertisements to the same person over multiple devices.[32] At a minimum, advertisers wish to track the number of times a user has seen a given advertisement, whether on web or mobile. While seeing the same advertisement can have a positive engagement effect on the potential customer, the advertiser will want to control when and where these ad deliveries occur.

To this end, the digital advertising industry has developed a suite of robust technologies that enable cross-device tracking. Some of these signals include information gleaned from IP address tracking, location data, and web tracking.[33] While regulators and advocates have raised various concerns with this practice, particularly for individual privacy, consumers currently have few options but to accept whatever voluntary standards internet companies develop.[34]

Once these cross-device inferences are made with high enough confidence, many companies—Apple, Facebook, and Google among them—may choose to associate what is known as a "persistent" or "unique" identifier with the user.[35] This identifier then becomes the central anchor of user data collected across multiple applications, platforms, and devices. In the end, this can mean that those who try to maintain a certain level of privacy by using different devices cannot actually do so if they

log in to certain web services or otherwise transmit signals that service providers can decode.

A close relative to the practice of cross-device tracking is that of browser fingerprinting. A user's browser fingerprint is typically a set of data about the set-up of their browser or operating system. It may include information about the user's browser provider and version, operating system and version, language of preference, list of browser plugins, Do Not Track settings, ad blocker usage, time zone, and browser fonts among other information.[36] Because there are so many possible configurations of browser settings, and users often tweak them to their liking, a browser fingerprint may be used to help identify an individual internet user. For instance, a widely-cited project developed by the Electronic Frontier Foundation in 2015 found that 84 percent of internet users who participated in the experiment produced unique browser fingerprints.[37] Critically, one group of researchers recently showed that websites could track an internet user's browser fingerprint even if the user were to use different browsers.[38] Users may be able to limit their exposure either through the use of privacy-enhancing software like that offered by Ghostery and other firms, or by minimizing their use of such features as Flash and Javascript, among others.[39] But all told, the advent of browser fingerprinting has in effect vastly increased the risk to user privacy.

Implications for Disinformation Campaigns

Disinformation propagators, like any other type of advertiser, benefit greatly from behavioral data collection. Their goal is to collect as much data about potential audiences as possible in order to tailor organic content—text, audio, or video that, unlike digital advertising, is posted on social media but which no entity necessarily pays to place or promote online—across media channels and target advertisements on internet platforms.

This can be done in a variety of ways. If the operator has one or more web properties, data can be collected directly through first-party cookies. Purveyors of disinformation often run fake news sites or blogs that carry content similar to that which they push over social media platforms. Behavioral data reflecting interaction with this content is particularly valuable for subsequent message development and ad targeting to specific demographic groups. Operators can also drop third-party cookies on other websites that carry content relevant to the campaign in order to expand the reach of collection. And they have a wealth of options for buying access to behavior and location data from commercial vendors like data brokers. In addition, the operator could link behavioral data from first-party cookies to personally identifiable information by collecting email addresses for subscription services or mobile phone numbers for text alerts.

Facebook also offers a popular service permitting it to track individual activity on third-party web sites that include Facebook engagement objects on their pages (e.g. "Like" buttons). Facebook is able to capture behavioral data—including data associated with both Facebook users as well as people who do not have a Facebook account—and link it back to its Audience Network service that allows clients to target advertisements outside of the Facebook application and on other mobile applications.[40] This service does not deliver personal data to the advertiser, but it allows them to integrate it into Facebook's advertising services. This is an effective way to segment various demographic groups and focus on individuals highly responsive to particular messages even if they do not have a Facebook account. Some regulators have in the past sought greater clarity on these practices and made certain requirements of the firm.[41]

From all of this data comes a nuanced picture of current and potential audiences. The more disinformation operators know about their target audiences, the easier it is to find, manipulate, and deceive them.

ONLINE AD BUYS

The vast collection of behavioral data on the internet has one central purpose: targeting ads more effectively to drive exposure to content, sales, sentiment, and persuasion. Due in part to advances in computing and storage capacity over the last few years, the sophistication and precision of ad targeting has increased dramatically, utilizing automated experimentation on the effectiveness of thousands of message variations in conjunction with user profiling. The analytics of audience segmentations offer a cost-effective way for disinformation campaigns to achieve two key goals: 1] to reach and cultivate responsive audiences beyond their core constituencies; and 2] to drive popular messages into viral phenomenon by flooding channels with promoted content.

The trove of user data acquired through digital tracking tools is applied to the cause of audience segmentation and targeted advertising, practices that over the past ten years have revolutionized the advertising industry. Advertisers have always attempted to create messages that would resonate with particular demographics: income, education level, gender, age, location, sexual orientation, language, and ethnicity. In the pre-internet era, this was an imprecise science. No longer. Major advertisers have shifted away from print and television. Last year, digital ad revenue surpassed television for the first time.[42] The most desirable "publishing" outlets to maximize engagement are now internet platforms like Facebook, Google, and Twitter.[43]

Over the past several years, internet companies have built increasingly effective tools for intelligently targeting digital ads toward individuals depending on their personal preferences. The business case is simple—the more relevant the ad, the more the user will engage with the advertisements they see and the longer they will stay on the platform.[44] Real-time ad dispatch powered by programmatic ad buying carries so much value for advertisers that it has become the industry standard.[45]

The targeting process is driven by data that comes from user-generated content or by directly tracking user behavior—either on the properties of the advertising network itself or from other websites. These datasets become large and detailed over time, permitting internet companies to compile specific profiles of various types of users of their services whose attention can be sold to advertisers. In addition to detailed demographic profiles, advertisers can over time infer the interests, preferences, beliefs, and behaviors of individuals as they interact with online content to create refined user profiles.[46] They typically retain large databases of these profiles, which sometimes might even include people who are not active users of their services. One can imagine this database as

AI Applications in Social & Digital Advertising

1 **Ad mediation** allows publishers to maximize ad revenues by searching different ad networks for the most relevant and best paying ads. Ad mediation software takes into account factors like the end user's region and device, and AI is increasingly being used to optimize ads being served.
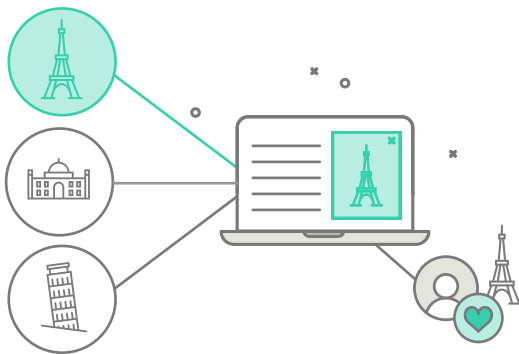
2 **Dynamic Creative Optimization** allows advertisers to optimize ad content for an end user. For instance, an airline could test how social media users react to factors like an advertised price or destination, or the timing of an ad, and then tailor ads to different groups of social media users.

3 Advertisers' primary goal is to engage people to look at their brands, products, and campaigns. Machine learning and AI are deployed across the digital ecosystem to help advertisers with **audience development**: uncovering the customer groups that are likeliest to react to advertisements.

4 **Social Media Management** firms are integrating AI to offer advertisers an ever increasing set of tools to start, develop, and target social media campaigns.

a spreadsheet, one user per row, with columns sequentially indicating every user's demographics, preferences, beliefs, and so on. We would note that while some of this data may be tied to personally identifiable information (PII), much of it is tied to individual user accounts maintained by internet companies and is generally unavailable in de-anonymized form to outside parties.

This database is then linked to what is known as an advertising technology platform. Through the platform, advertisers can select the communities they most want to reach through an advertising campaign. For instance, a small business that specializes in selling men's sneakers in New York City might wish to target males aged 18-35 who are interested in basketball, have visited a major shoe manufacturer's website in the past seven days, and live in that metro area. Modern advertising technologies allow them to do exactly this with tremendous accuracy and at low cost.[47] Sophisticated advertisers will test variations of the same message across a wide variety of user profiles to optimize for the content and target profile that delivers the highest response. This kind of conditional testing can sometimes represent thousands of ads placed by the advertiser each day.

Because these services are delivered over advertising platforms that are typically powered by algorithmic sensoring technologies more so than human review, targeted advertisements can be dispatched automatically across thousands of websites and social media feeds simultaneously. The engagement statistics are logged immediately and serve to train both the advertiser and the advertising platform algorithm on the most successful targeting strategies. The more money and time an advertiser spends on testing market segmentation and message responsiveness, the more effective they become at both refining and growing their lists of high response constituencies.[48]

Advertisers today have a plethora of options in designing a digital ad. One key dichotomy, though, is that of content that appears in panel displays versus that which appears as promoted content in,

for example, the user's social media news feed. The former is the descendent of the traditional display ad, and might appear in a side panel next to search results or the news feed. The latter includes content that the advertiser might post to their social media account, and pay to promote so that it appears with more frequency in users' social feeds.

The leading internet platforms have taken targeted advertising to another level by offering advertisers the ability to target ads at customer audiences.[49] These audiences—or specific groups of people using a given internet service—typically comprise people who are known customers of the advertising brand. Customer audiences can be composed by advertising brands in a number of ways, but perhaps the easiest way is to collect their personal information—such as name, email address, or phone number—when they visit the brand's website or brick-and-mortar store. For instance, Breitbart, the news and commentary platform, asks for visitors' PII repeatedly throughout its website. Capturing personal contact information is standard practice for many online services. The purpose of this is twofold: first, to directly send them messages and offers over email, but perhaps more importantly, to know who their audience is so that Breitbart can target them on advertising platforms like Facebook that enable segmentation of audiences that match PII previously gathered by the advertiser.

Advertisers utilize still cleverer mechanisms to compose customer audiences. One of the best examples of a tool that enables this is the Facebook pixel, a 1x1 image that some web developers choose to insert on their pages to enable targeted advertising via Facebook. When an individual visits a website that has installed the pixel, Facebook can automatically receive data about that person's visit to the page. Though transfer of PII to the website operator may not be a part of these informational transactions, the website owner may thereafter be able to target Facebook users who have visited their website and loaded the Facebook pixel within the past several months.[50]

Scroll Popup	**Shadow Popup**	**Side Panel**

Persistent email address list sign-up requests on breitbart.com

Customer audiences are often only the gateway to more expansive, automated outreach to a wider group of people. The leading internet platforms now enable advertising clients to upload customer audience lists and "remarket" their advertising campaigns to a wider group of like-minded people using those platforms.[51] In the case of Facebook, this advertising feature is known as reaching a "lookalike" audience.[52] In concept, an advertising brand could have a list of its 5,000 best representative customers, and wish to find more people like them. The brand could upload that list of email addresses or phone numbers of those 5,000 people to Facebook's advertising platform. These email addresses are then matched to Facebook accounts to form a custom list. Through the 'lookalike' function, Facebook might then identify 50,000 more people who are similar to the client's original list of 5,000 customers—with the total number of identified lookalikes depending in part on the size of the ad spend. These similarities could be based on any of the information the advertising platform possesses about them, including gender, ethnic affinity, or, for example, a special interest in a particular sport. This can help the advertising client reach a much larger audience of people, many of whom may have never interacted with the client's brand before. If and when one of those new potential customers visits the client's website, the client could collect their personal information and continue building its base customer audience.

Over time, the advertiser and its associates— sometimes including advertising agencies or

advertising platforms—can continue to refine the advertiser's outreach list and use algorithms to better target potential customers and identify lookalike audiences. Further conversion tracking technologies—web tools that help advertisers measure the success of their digital advertising, which might include web beacons, mobile app monitoring scripts, or actual sales information— can also be applied to better target future ads and engage customers as effectively as possible.

Implications for Disinformation Campaigns

The tools of audience segmentation and targeted advertising are the most important features of digital advertising for disinformation campaigns. All of the behavior tracking tools, collection of personal information, and audience profiling that disinformation operators might utilize are largely aimed at increasing the effectiveness of targeted advertising. Organic content posted on social media can only stretch the reach of disinformation to those users who follow the feeds of the disinformation provider or see the posts of friends who recirculate that type of content. To reach new audiences with different demographics and preferences or to blanket a social media platform at scale with particular messages requires display advertising or paid promotion of content that is otherwise posted organically. These are increasingly common tactics among disinformation campaigns—profiled for example in a recent report on campaign tactics used

ahead of national elections in Kenya.[53]

These ad buys deepen engagement with known audience segments and broaden engagement to new ones. They also boost content that is clearly popular over a threshold where it might go viral and enjoy wide distribution to tens of millions of users via the network effects of large internet platforms. On top of these benefits, the disinformation campaign gets the support of the ad platforms' own built-in intelligence about targeting (e.g. taking custom lists and building lookalikes) and in turn contributes to making that algorithm smarter and smarter with each new message and audience segment that is tested. In other words, the more successful the ad buy (including disinformation), the more effective the successive ad buys will be because the ad platform has learned more about the best targets. This is the virtuous cycle of advertiser and ad platform; both benefit from the success of the other.

SEARCH ENGINE OPTIMIZATION

The search algorithm sits at the center of the internet economy, steering billions of internet users each day to sites that appear at the top of the results page. A multi-billion dollar industry has developed that is dedicated to optimizing search engine results for all kinds of commercial websites through careful observation and reverse engineering of the Google search page rank algorithm. Most of the tools and tactics in this industry are above-board and dedicated to lifting company websites to the top of the rankings for particular search terms. However, other techniques—known collectively as "black hat SEO"—are designed to trick the algorithm and dominate search results for a few hours of the news cycle before Google corrects the distortion. This is a critical weapon in the arsenal of the precision propagandist.

An important tactic of disinformation campaigns is the strategic manipulation of search engine results pages (SERP). Numerous reports published since the 2016 election have highlighted very unusual search results on sensitive topics. For example, a week after the election, the top news listing in a Google search for "final election results" was a link to an obscure blog called "70 News" that claimed Donald Trump won the popular vote.[54] In January 2017, search terms related to U.S. intelligence reports on Russian interference in the American election returned a majority of top links pointing to Russia Today content denying the allegations.[55] In early October, after the horrific terrorist shootings in Las Vegas, top-ranking search results the next morning on Google were conspiracy blogs that claimed the shooter was an anti-Trump liberal with links to ISIS.[56]

How is this happening? Google says that these are glitches in their algorithms that will be swiftly remedied. That may sometimes be true. But these incidents are potentially the result of intentional efforts to manipulate commercial search algorithms to steer users towards particular content. Google has been dealing with these problems for years. Most incidents are competitive attempts by marketers to steer potential customers to one product rather than another. The efforts of disinformation operations to manipulate search is a newer phenomenon. It has caused repeated embarrassments for Google.

Allowing that Google is working hard to combat this kind of distortion in search results, this cat-and-mouse game nevertheless continues to produce disturbing incidents that privilege conspiracy and falsehood over news and facts in search results. That matters because without question, search results on news topics play a role in shaping public opinion. And because Google owns over 75 percent of the search market in the United States, the integrity of its search algorithm is directly connected with the integrity of public debate. The fact that search manipulation succeeds despite Google's wealth and technical capacity is a testament to how much effort is going into gaming the system.

> **Companies are spending almost three times as much to optimize organic search results as they are to buy search ads.**

It is not clear who is doing this, but it is clear what they want. The primary objective for those trying to shape opinion with search engines is to drive their content to the top of the search results page. Studies show that about a third of internet surfers click on the first non-advertising link in the search results. The top five search results get around 75 percent of the traffic.[57] The first page of links (top 10) gets 95 percent.[58] Hence, there is huge incentive to be at the top, whether you are selling tennis shoes or publishing political news. This has spawned

an industry known as search engine optimization (SEO) that is big enough to go toe-to-toe with Google and win with surprising frequency.[59] The entire internet content industry is in many ways organized in response to Google's algorithm.[60]

SEO is big business. To get a sense of it, consider that Google's total revenue from all search advertising in 2016 was $24.6 billion.[61] That is what advertisers paid to serve ads on search results pages. The size of the industry—SEO—that seeks to occupy the top unpaid spots on the search results page is estimated to be $65-70 billion in annual revenue.[62] If this estimate is correct, companies are spending almost three times as much to optimize organic search results as they are to buy search ads. That is extraordinary given that Google's search ad business is an historic money-maker.[63]

The costs involved with SEO are a function of the complexity of the problem that optimizers are trying to solve. Google's search algorithms—known collectively as Hummingbird—utilize a legendary set of 200 different (secret) ranking signals to determine search rank. Included in this system is machine learning software known as RankBrain. It is constantly updated. By one estimate, Google changes its search algorithm five to six hundred times each year. These updates impact search results and therefore revenues, and consequently they are tracked carefully by the SEO industry.[64] In addition, Google publishes a lengthy set of guidelines for its search quality evaluators. These are thousands of people around the globe that Google contracts to evaluate search results based on numerous factors. The results submitted by the evaluators are integrated into the search algorithm, but they are not a direct proxy for it.[65]

The entire mission of SEO is essentially to reverse-engineer the moving target of Google's search algorithm in order to modify websites to achieve a higher search rank. But of course, they don't have clear guidance from Google about how it works. What they do have are the results of innumerable trial-and-error searches that can be painstakingly examined in A/B tests to identify which features

of a website are favored by Google's search rank. It is like measuring shadows on the wall to draw the objects that made them.

Over time, the industry has settled on standard practices that form the basis for any conventional SEO strategy. These techniques include standard methods of website architecture, content formatting, and link building (i.e. getting as many other sites as possible to link to your site). They are very common across the commercial web. This "white hat" SEO work adjusts websites to fit how Google crawls, indexes, searches and ranks. These tactics are consistently applied and designed to put a company website in the top ranks and keep it there.

But there are other tactics that are episodic and ephemeral, sometimes known as "black hat" SEO. They are designed to put a particular webpage or set of pages in the top ranks for a short period of time (e.g., a news cycle). These tactics are especially important to disinformation campaigners. Some of these tools are "black hat" SEO operations that attempt to trick the Google search algorithm into assigning a search rank that does not correspond to quality content, source reputation, or even topic relevant responsiveness to the query.

Boiling it all down, most of the SEO business (and in turn the search business) is about three things—content, links, and reputation. These are the common denominators of the numerous SEO guides and handbooks available to companies.[66] And they are the basic components of both white hat and black hat SEO techniques.

- **Content:** Google's algorithm crawls the semantic structure of entire pages to determine how well search terms are matched. The richer and more specific the content, the better the rank. Keywords matter to some degree, especially in the anchor text of a URL, page titles, and article headers. A logical link structure between sites within a domain, mobile friendly display, and a speedy load time are also now standard SEO recommendations.

- **Links:** The centerpiece of SEO is link building—getting other websites to link to yours (backlinks). Backlinks from highly credible sites increase the score; backlinks from spammy sites damage it. Social media links (through Facebook and Twitter) appear to also count, but not as much as organic backlinks.

- **Popularity:** SEO success breeds more success because click-through-rate (CTR) is also a factor in search rank. Every time a user clicks through a search result, the page and domain reputation improves. Fresh articles on similar topics that also draw attention increase the rank further. Ad buys on the search page do not contribute to CTR reputation. However, they do help zero-in on the keywords that deliver the best conversion rates and contribute to optimized content.

None of these elements of SEO is inherently illegitimate, nor do they indicate any inbuilt bias in the search algorithm. However, they can all be gamed to some degree—particularly if the objective is to dominate search results pages for a short period of time.

With so much money and technique pouring into SEO, it is fair to conclude that anomalous and distorted search results are not always attributable to glitches. They are there because someone spent money and effort to put them there—through both legitimate and illegitimate means. And judging by the frequency of reports about distorted search results, it appears to be working.[67] It is worth repeating that intentional distortion of search results is not a new phenomenon. But its relevance to the emerging practice of comprehensive disinformation campaigns does appear to be novel.

And Google's usual reaction—to change the algorithm to frustrate the cheaters—is trickier to implement. When SEO results in a search page include disinformation that appears to lean in a partisan or ideological direction, it puts the search engine operator in a much more difficult position. If they "fix" it, they will be accused of bias by one

side. If they don't, they will be accused of bias by the other. In the meantime, they endure the discomfort of periodic explanation about why search results appear to support a disinformation campaign. Whether or not these incidents are the result of underhanded SEO is almost never transparent to the user.

We can see evidence of these "black hat" SEO tactics at work when search results on a particular news topic are dominated by blogs with similar (usually extreme) viewpoints that push more credible news sources out of the first page. The content on each page is typically rich—diving deep into a specific topic. It is then published and republished in slightly varied forms with high frequency to keep the freshness factor high. Then, a coordinated set of domains publishing similar stories link to one another as often as possible, generating a robust backlink economy that is hard to match for other stories in a short, noisy news cycle. These links are then aggressively promoted through social media and advertising spends, pushing up the number of social media links as well as the CTR. One example of how this is done—reputed to be a popular current tactic—is coordinated posting of a particular URL (or URLs) on Reddit. Hundreds or thousands of posts across relevant Reddit sub-threads are crawled and indexed by Google's search algorithm and may play a role in driving up search rank before Reddit moderators intervene or Google spots an anomaly.[68]

From a technical perspective, this flurry of interlinked activity on a breaking news topic will be hard to distinguish from a genuine media event, until users start complaining that the search results are dominated by Russian propaganda, conspiracy blogs, and other forms of disinformation.

> **The evidence of reported anomalies in search results suggests that SEO has been used as a disinformation tactic in the past**

In response to periodic reports of disinformation in search results, Google has pledged increased vigilance in updating its algorithm to outfox SEO mercenaries. They have been highly successful in the past in identifying and shutting down common "black hat" tactics, and have even created search-rank penalties for those that use them.[69] It is unclear whether these are a substantial deterrent. Meanwhile, in a new initiative called Project Owl, Google has pledged to offer users the ability to flag when autocomplete (when search fills in a suggested search term based on the first words typed) or snippets (content featured as highly relevant on the search results page) produce a problematic result.[70] There is also the Trust Project, an initiative spearheaded by media researchers and news organizations dedicated to measuring and describing trust indicators for news content and working with Facebook, Twitter, and Google to make those descriptions available to users.[71] Each of these initiatives intends to address the prevalence of disinformation by steering search results to more credible/authoritative content based on different forms of user input. But in a sign of just how challenging it is to parse good content from bad, Google has been criticized for marginalizing certain views and perspectives.[72]

Implications for Disinformation Campaigns

The evidence of reported anomalies in search results suggests that SEO has been used as a disinformation tactic in the past. And it is reasonable to assume this will continue as long as it is effective. More inquiry is needed to determine how widespread this phenomenon might be and whether it is differentiated according to market, language, or topic. Further, SEO may play an outsized role less in the promotion of clearly false content but rather in the gray zone between malicious disinformation campaigns and the digital media channels that profit (either commercially or politically) from promoting these stories and themes. These sites may have built legitimate credibility and popularity (with click-through-rates and a backlink economy

to validate that status) but leverage SEO on occasion for the cynical purpose of amplifying misinformation.

We should also watch for evidence of SEO tactics being used in ways that are analogous to how zero-day attacks (cybersecurity breaches using previously unknown flaws in popular software) are used by hackers. Disinformation campaigners may identify ways to inappropriately amplify search rank and then deploy that tool in a one-off effort to promote a story during a particular news cycle, betting that they will achieve viral distribution before Google can respond and shut it down. Notably, there is virtually no alignment of interests between black-hat SEO and Google (or Bing and Yahoo for that matter). Distorted search results damage the search engine brand and send users looking for alternatives. This should be an area where the companies are inclined towards collaboration with other stakeholders to stamp out practices that cause public harm. Combatting search manipulation may be an early opportunity to apply a common strategy across corporate policy, media consumption research, and user engagement.

SOCIAL MEDIA MANAGEMENT SOFTWARE

A new breed of digital marketing company—the social media management service (SMMS)—may represent the most promising intersection of machine learning algorithms and advertising technology. The SMMS offers advertisers a fully-integrated solution that pre-configures campaigns with multiple messages for different targeted audiences across both standard social media posts and paid content. The software draws on complex behavioral data analytics, engages in real-time social media listening to place the right message at the right time, and coordinates across multiple channels simultaneously and automatically. It tests and learns to optimize persuasive power for every dollar spent. It is a brilliant innovation for advertisers. It is a finely tuned disinformation machine for the precision propagandist.

Over the past ten years, a wide variety of social networks have emerged as platforms for people around the world to connect with others. While the largest and most visible social network is clearly Facebook, social networks come in many forms for many different groups of people. YouTube allows anyone to openly share and view videos online; Twitter is the most noted micro-blogging

site on the web; and Snapchat grew exponentially in popularity among adolescents who wanted to connect with each other over their own private ephemeral messaging service. The list of leading social media websites, however, extends far beyond those few leading American brands. There are dozens of online social networks and virtual communities that can claim at least 1 million users, and each of them offers a platform over which ideas and content can be shared at a quick pace.

Given the importance and reach of social media, advertisers have quickly recognized the tremendous value that can be gained out of executing well-timed and well-targeted advertising campaigns over social networks. But traditionally, they have had one major problem: How can they simultaneously manage effective campaigns across several different social networks in real time in a way that optimizes for impact at low cost?

Over just the past few years, a thriving new industry of "social media management platforms" has emerged to help brands address the headache of managing ad campaigns and content sharing over multiple social media channels.[73] Led by brands such as Hootsuite, Sprinklr, Hubspot, and Sprout Social, this sector includes firms that help customer brands efficiently manage their social media accounts and coordinate campaigns.

Imagine a global soft drink brand like Pepsi that wishes to develop a marketing campaign for a new line of seltzers. Instead of having its marketing team manually manage the brand's Facebook, Twitter, Instagram, and other social media accounts individually on each of those social media websites, Pepsi would most likely wish to empower the marketing team to manage all of those accounts through one global interface.

That is precisely what the social media management platforms offer. By enabling clients to create content and share it in just a few clicks through Hootsuite's commercial web application, brands are able to better control their marketing messages, including through the content they develop, the platforms

where they choose to share it, the timing of when they wish to push it, and the audiences they wish to target.

But the management of a client's social media accounts is not limited just to creating efficiencies in sharing or promoting content to targeted audiences. As the social media management industry becomes more competitive and newer players emerge, platform providers are increasingly trying to feature more automation in their services to empower their clients to drive media campaigns with even less human involvement.[74] Critically, social media management platforms are now integrating machine learning algorithms into the client workflow to power recommendations about ad audience, content, timing, and other factors.[75]
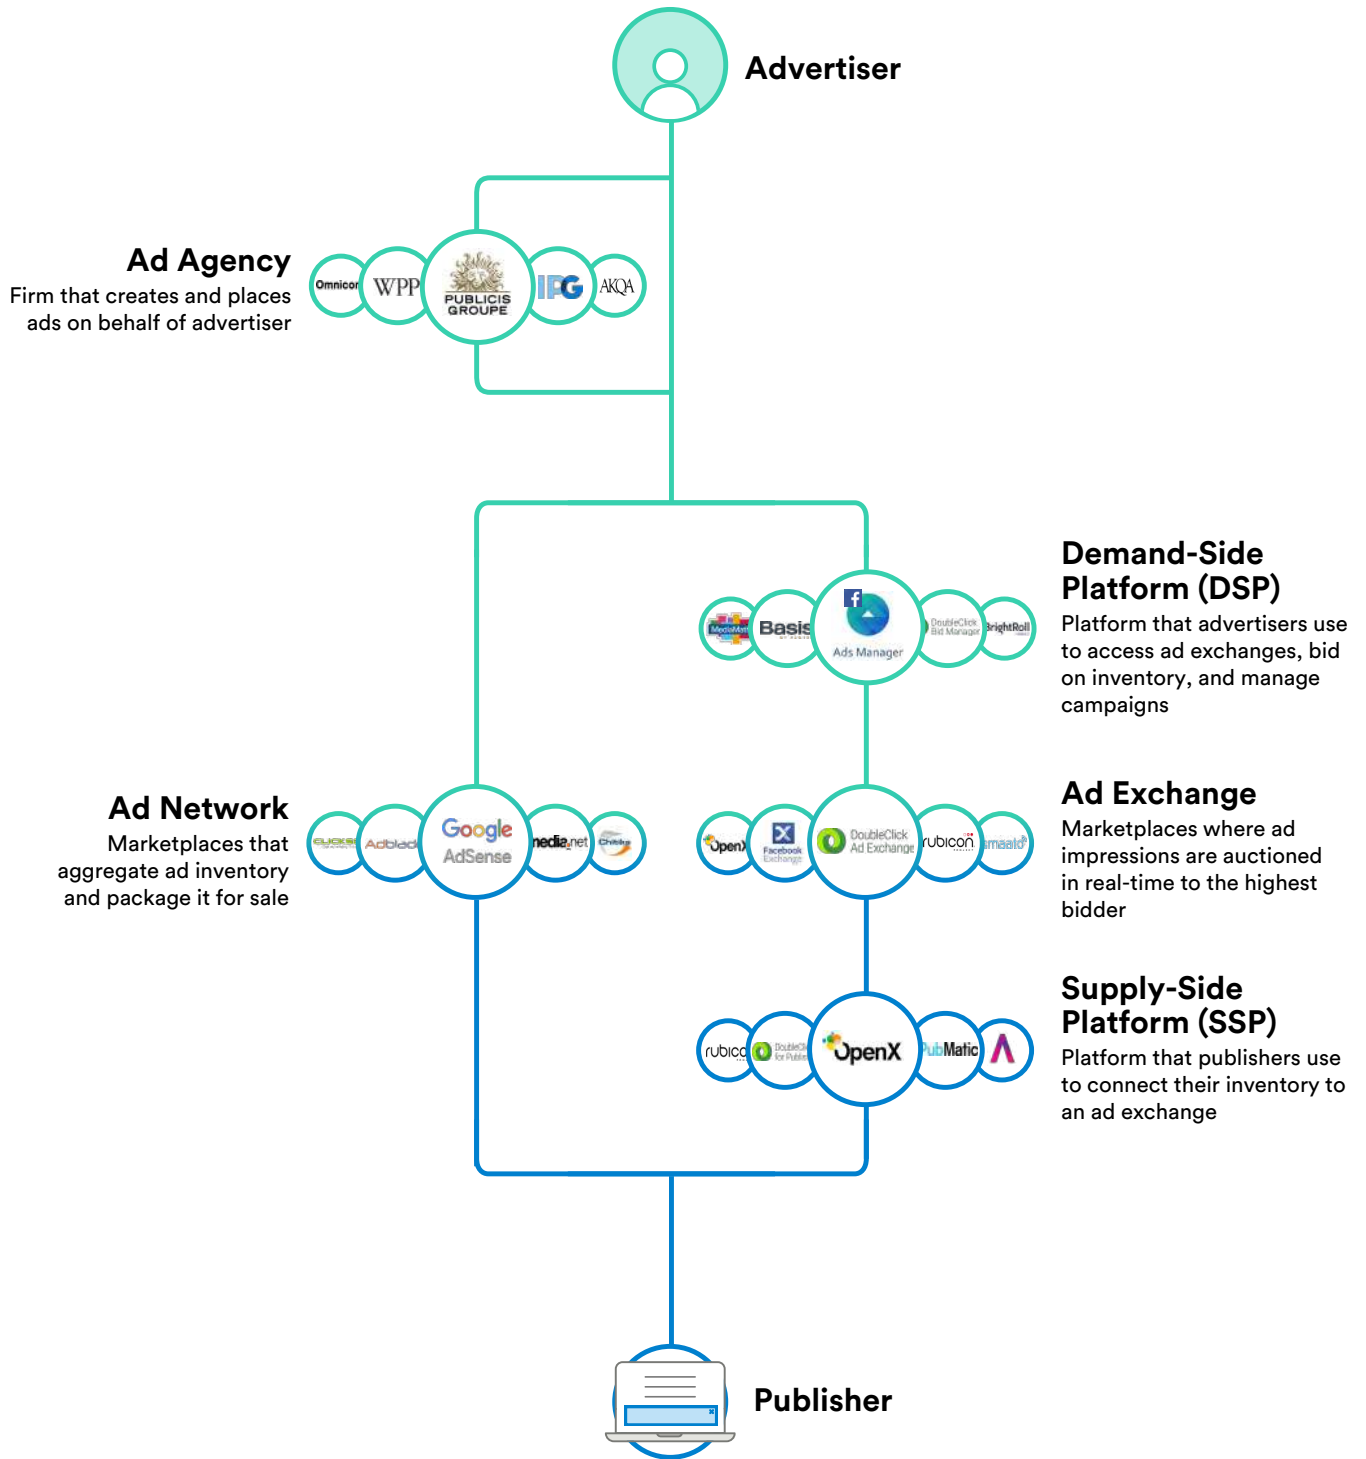
> **Platform providers are increasingly trying to feature more automation in their services to empower their clients to drive media campaigns with even less human involvement.**

This kind of technology can have tremendous commercial impact for brands that wish to develop a robust marketing campaign. Notably, social media management platforms now enable clients to create contingency-based outreach strategies. In such scenarios, clients can set their accounts to initiate a particular advertising approach or campaign if and when certain events happen.[76] For instance, a customer-facing brand like Sprite could create a contingency to trigger promoted advertisements about their drink anytime LeBron James appears in a Sprite-related hashtag or whenever the Cleveland Cavaliers are playing on television.

This full suite of advertising technologies can come together to prodigious effect in the context of political communications. Political advertisers might start by deciding to attempt to reach certain

The Digital Advertising Ecosystem

Advertiser

Ad Agency
Firm that creates and places ads on behalf of advertiser

Demand-Side Platform (DSP)
Platform that advertisers use to access ad exchanges, bid on inventory, and manage campaigns

Ad Network
Marketplaces that aggregate ad inventory and package it for sale

Ad Exchange
Marketplaces where ad impressions are auctioned in real-time to the highest bidder

Supply-Side Platform (SSP)
Platform that publishers use to connect their inventory to an ad exchange

Publisher

online audiences with a persuasive message about why a political candidate is or isn't optimal for individual voters. To determine which groups of people are best to target online, they could begin to collect personal and behavioral data from them—which could be purchased from political entities and data brokers, or harvested online through alternative means. Through the use of SMMS platforms and other services, political communicators can then begin to segment these various audiences by demographic background, geolocation, and other signals. SMMS platforms can then help political advertisers blend the tools we have heretofore described to help the communicator create a persuasive campaign that engages voters at once with promoted and organic content on social media platforms like Facebook and Twitter, in addition to display ads on YouTube and Google Search. Certain features of SMMS can also allow advertisers to manage contingency-based, automated content placements, such that the client is enabled to automatically begin an online ad and content campaign according to event contingencies, be they a statement by the U.S. president, an announcement by a foreign actor, an electoral candidate's victory, or a sudden public outcry on a particular issue like an incident of local hate crime. And over time, they can continue to refine their knowledge of the individual profiles of each person they are tracking through a combination of real-time listening and other online services.

> **The entire toolbox of advertising technologies can be packaged together into coordinated campaigns that utilize both human and machine intelligence to optimize marketing.**

The modern SMMS should be able to help the client develop and maintain an automated response for every single combination of components in this multi-dimensional grid of contingencies. More

critically, the SMMS service could begin automating actions on behalf of the political client based on different contingencies that happen in the political sphere in real time. For instance, as the SMMS continually manages the outlay of the content of the client's political campaign, the SMMS could automatically record past ad targeting parameters so that it can refine and adjust them over time for optimal impact. Additionally, the SMMS could empower the client to execute further ad buys that segment its customer audience in different ways—whether using aggregate geographic or other demographic data features. SMMS and other client services could even help the client monitor more subtle sentiments among the voting population using social media listening services. These services could furthermore trigger these internet-based campaign responses on an automated basis, so that if many people begin tweeting with negative sentiments about a statement made by, say, the Turkish president, the SMMS-determined campaign is triggered and targeted at those tweeters.

Implications for Disinformation Campaigns

The SMMS industry demonstrates a point we have been alluding to throughout this paper—that the entire toolbox of advertising technologies can be packaged together into coordinated campaigns that utilize both human and machine intelligence to optimize marketing. If disinformation campaigns are not yet making use of these services (or adapting them for in-house application), it is only a matter of time until they do. The increasingly automated, contingency-based nature of these management services make them an optimal fit for disinformation operations that must respond to current events with rapid efforts to steer users to a particular interpretation of the news story. The SMMS can digest large amounts of behavioral data tracked over time, produce demographic ad targeting parameters, and refine them as new visitors land on the site as a result of organic and promoted content. The service can interface with Facebook, Twitter, and YouTube simultaneously and leverage

both its own data set and the advertising technology algorithms offered by the platforms.

Listening to social media to map sentiment, pre-loading both organic and promoted content distribution, and coordinating across multiple platforms and sites in order to create a backlink economy that drives SEO—these are weapons of choice for disinformation. Once again, there is nothing inherently nefarious about these tools themselves. They are perfectly legal and for the most part even align with the economic interests of the platforms. All parties in this ecosystem benefit financially from successful advertising campaigns. They have developed brilliant tools to achieve more consistent persuasion. But they have also opened the door to abuses that harm the public by weakening the integrity of democracy.

Consider, for instance, a more sophisticated agent of the Russian government that wants to subvert a future U.S. electoral process. It could choose to set up a shell company that places ads into social networks. Through an SMMS platform, either developed in-house or acquired through a service provider, it could track social media sentiments across all major U.S. congressional districts in the weeks of early voting leading up to election day. Using the platform tools, it could infer how many people have gone to the polls in each district by election night. And finally, it could target a voter suppression content campaign at those districts in which not many voters have yet cast their ballots. While it is probable that propagators of disinformation did not go to these sophisticated lengths in 2016, they will very likely become well-equipped to do so not long in the future.[77]

All parties in this ecosystem benefit financially from successful advertising campaigns. They have developed brilliant tools to achieve more consistent persuasion. But they have also opened the door to abuses that harm the public by weakening the integrity of democracy.

ADVANCES DRIVEN BY ARTIFICIAL INTELLIGENCE

The digital advertising industry has undergone a technology-driven transformation over the past 25 years. As the industry moves forward over the next decade, we will likely continue to see further innovation that will drive greater profits for advertisers, agencies, and publishers alike—largely because of the tremendous impact of artificial intelligence (AI).[78] As we have discussed earlier, online advertising often involves highly complex, contingency-based decision-making processes to determine what type of digital content to send to which audience segments when various event contingencies occur. This complexity is heightened in the political context, as individual sentiments about political ideas or candidates are often particularly impressionable and therefore fickle. AI and its associated family of advanced algorithmic technologies are thus likely to play an enlarged role in the context of political disinformation propagated through online social platforms.
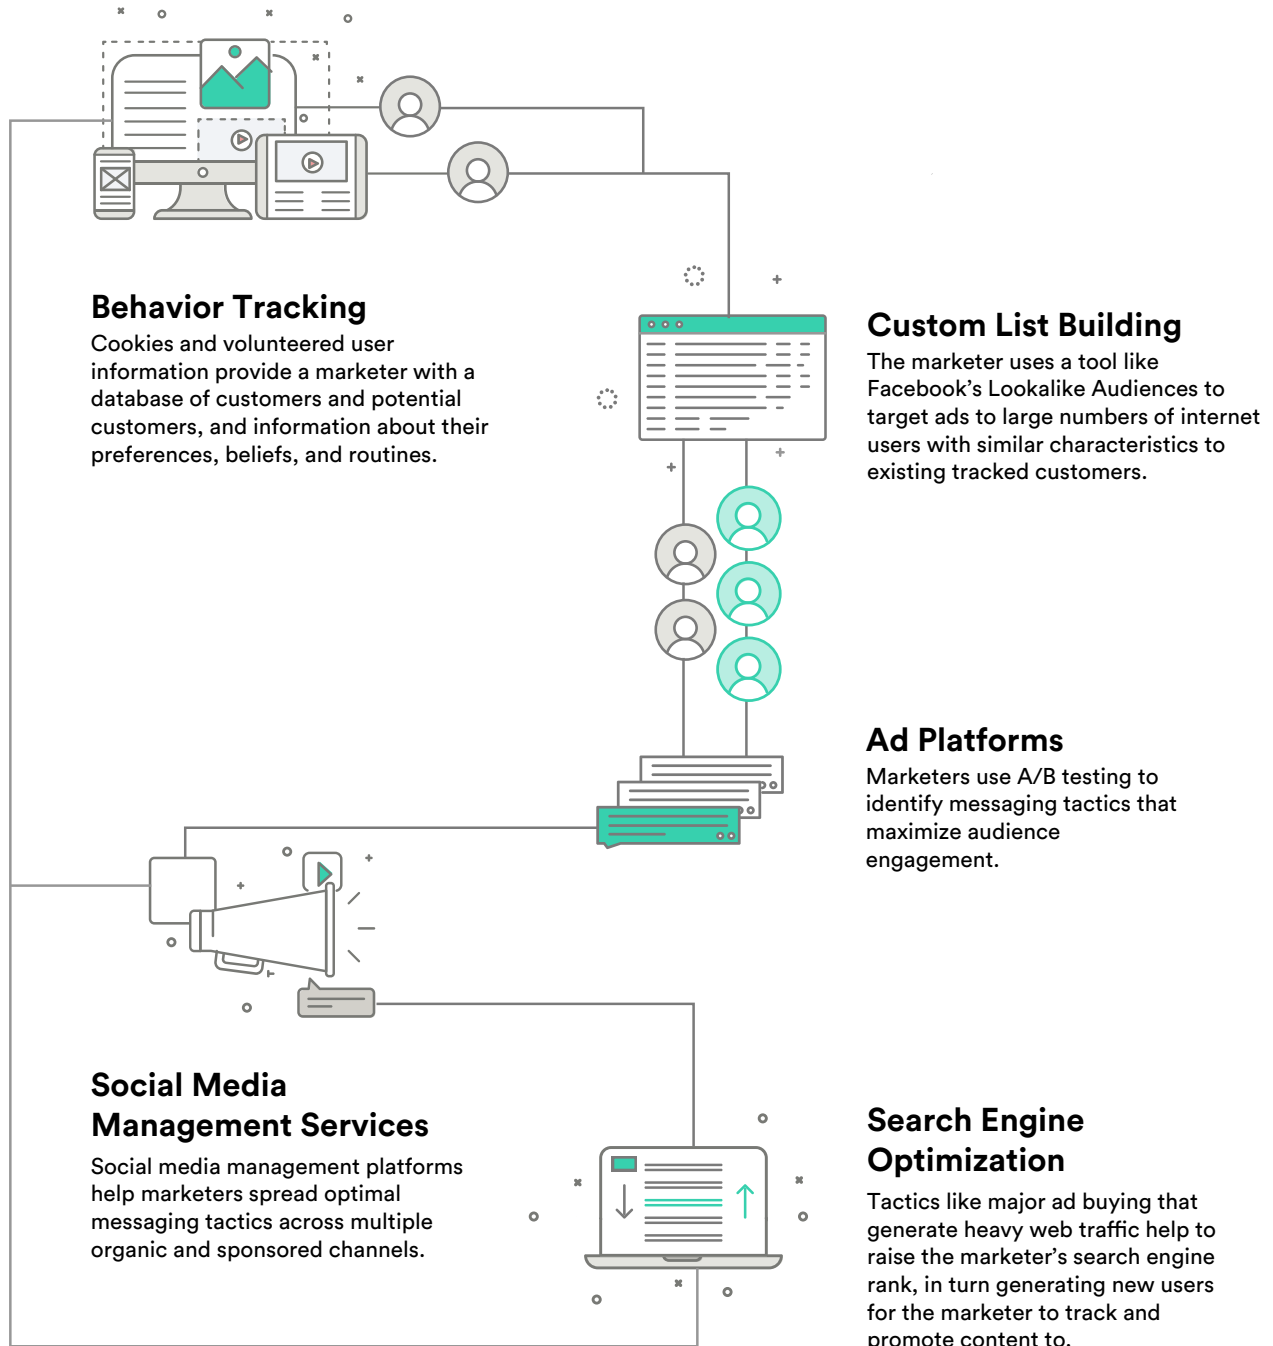
In addressing AI, experts often compare "strong" AI and "weak" AI.[79] Strong AI is so named because it has general applicability to many problems and situations, much the same as a human's intelligence. This conceptual form of AI is also called "artificial general intelligence," an example of which could be HAL 9000, the AI system that controlled the airship in 2001: A Space Odyssey. Such a sophisticated AI has not yet been developed; most scientists and engineers believe that we are still decades away from achieving strong AI.[80]

Weak AI, on the other hand, is all around us. Weak AI can take the form of advanced machine learning systems, which are deployed in such applications as self-driving cars, creditworthiness decision-making, page ranking in Google Search, and content ranking in social media news feeds. Another area where weak AI is taking hold is in advertising technology.[81]

Weak AI has the capacity to understand a narrow environment, typically with a degree of memory and computational power many orders of magnitude higher than average human intelligence. In an application like decision-making for financial credit, a weak AI that utilizes machine learning techniques can quickly study millions of past examples of creditworthiness decisions—also known as a 'training set"—and can learn about how to most effectively execute future creditworthiness decisions remarkably quickly. Such automation could potentially obviate the necessity for any sort of human intervention in the credit decision-making process.[82]

The value of weak AI, including machine learning, extends deep into the digital advertising space. Brands typically wish to programmatically reach the people who are most likely to react to their

A Full Suite of Advertising Technologies

Behavior Tracking

Cookies and volunteered user information provide a marketer with a database of customers and potential customers, and information about their preferences, beliefs, and routines.

Custom List Building

The marketer uses a tool like Facebook's Lookalike Audiences to target ads to large numbers of internet users with similar characteristics to existing tracked customers.

Ad Platforms

Marketers use A/B testing to identify messaging tactics that maximize audience engagement.

Social Media Management Services

Social media management platforms help marketers spread optimal messaging tactics across multiple organic and sponsored channels.

Search Engine Optimization

Tactics like major ad buying that generate heavy web traffic help to raise the marketer's search engine rank, in turn generating new users for the marketer to track and promote content to.

products or messages. Similarly, publishing and customer engagement platforms like social media services wish to help brands advertise as quickly and efficiently as possible. Given the wealth of personal data available to advertising networks, the wide variety of advertising clients and their needs, and the real-time nature of digital ad-serving, AI can connect brands to their optimal audience more effectively than any other existing technology. This new application of AI has already raised certain ethical concerns among experts and consumer advocates.[83]

One example of AI at work in this space is in digital ad mediation, the process of connecting a mobile publisher or consumer engagement platform (such as the *New York Times* or Twitter mobile applications) with a brand in real time to serve a display ad to the end user. Needless to say, given the volume of content and billions of users of mobile applications, mediation is most effectively accomplished through programmatic, or automated, ad serves. Mobile publishers typically work through any number of ad networks to place the ad that will provide highest revenue to the publisher. This market is set up so that the ad network knows as much as possible about the publisher's users and the brand's qualities, so as to maximize aggregate user engagement with the ads they see.[84] It is easy to see how a weak AI machine learning system can be trained over time to bring in optimal revenues for publishers while also serving the needs of advertisers.

Mediation can bring extraordinary efficiency to the advertising industry, but existing applications of AI do not end there. Firms are starting to explore combinatorial testing of digital ads, so that advertising agencies, publishers, and brands can understand how different personalities react to ads that might differ in content, timing of delivery, medium of delivery, or some other variable. There are strong similarities between this type of advertising industry AI and the Google-developed AI that recently defeated a Go master. The Go AI was trained to explore different combinatorial contingencies in the game by automatically

playing out many different scenarios and making a determination of what the next best move would be based upon its understanding of the contingency matrix.[85]

Another rapidly expanding use case of AI in advertising is with lookalike audiences. As discussed earlier, social media networks enable clients to upload customer lists to create an audience for an ad campaign—presumably consisting of people very likely to come back to the store for another future purchase. But companies like Facebook further enable the creation of lookalike audiences—comprising users the company believes share similarities to those in the client's original customer audience list. As social media companies search their user lists for similarities in personalities and preferences between different groups of people, it is natural that such functions will only be exponentially aided with the integration of AI-powered analysis and decision-making.

One of the most effective applications of artificial intelligence that is still evolving rapidly is in the space of social media management software. This is a sector in which leading service providers like Hootsuite, Sprinklr, and Lithium Technologies enable clients to understand public sentiments about their brand, often to a granular local level; queue content and advertising campaigns that can be pushed out at the click of a button; and match that sentiment analysis with prepared ad campaigns in real time, so that as soon as a certain incident or event occurs, the ad campaign can be automatically triggered for public release. As discussed previously, AI can aid the efficiency and impact of social media management to tremendous effect.

Integration of AI will not end here. The industry is constantly innovating to bring AI into automated decision-making processes, particularly those that fuel the central profit-making functions of these companies—namely, advertising technology. Because of the challenges AI presents to consumer privacy and individual autonomy, many have questioned such integrations. The leading internet

companies have begun conversations with civil society about these issues,[86] but it is a massive public education challenge to alert communities about the inherent ethical challenges related to AI and engage them in discussion of solutions.

Implications for Disinformation Campaigns

The rapid advance of AI technology development is supercharging the digital advertising industry. Techniques that have been around for many years are now much more effective and scalable. AI will solidify the dominance of digital advertising over all other media channels. It will in turn greatly enhance political advertising online. What is possible from election cycle to election cycle in terms of precision voter targeting will leap forward. And with it, the troubling power of disinformation campaigns will grow as well. At its core, the marriage of advertising technology and political propaganda is nothing more than applying the tools of the industry—behavioral data analysis, audience segmentation, and tailored message targeting—to the task of exploiting prejudice. AI will create advance capability in each stage of the process. This prospect begs the question of what could or should be done to arrest this trend, or at least to manage its negative consequences for the public interest.

CONCLUSIONS

This analysis of digital advertising technology and its relevance to disinformation online is designed to broaden the focus in the current public debate beyond Russian operatives buying ads on social media. The problem is much bigger than that and the issues of concern are more diverse. Our analysis points to the core challenge of disentangling the alignment of interests between the commercial pursuits of digital platform companies and the success of disinformation-based political advertisers.

It is a mistake to fixate on Russia. Russia is one of many online disinformation operators targeting Americans. Future disinformation campaigns may just as likely be run by domestic operators as foreign ones. These operators will most likely leverage the most dominant U.S. internet platforms to reach tens upon hundreds of millions of Americans. The full range of these disinformation campaigns could produce a grave public harm. In particular, they can progressively weaken the integrity of our democracy by separating citizens from facts and polarizing our political culture.

The next generation of disinformation operators will use a robust toolset of digital advertising to pursue their goals. Understanding how this may work requires a thorough analysis of the dominant practices and technologies in the digital advertising industry. It is much more than buying ads on Facebook, YouTube, and Twitter. The true power of the advertising technology market lies in a combination of tools that prepare and enhance these ad buys. Vast efforts go into behavioral data collection to segment audiences with precision and target them with variations of different messages most likely to produce a response. The data analysis is the rocket fuel for targeted advertising on social media platforms and in the broader digital marketing industry. And increasingly, the data analysis and the targeting is handled by machine learning algorithms that grow ever more sophisticated.

Meanwhile, the campaigns themselves, while designed by the marketing operators, are implemented and optimized by AI. This technology is developing rapidly and is becoming increasingly effective at identifying, targeting, and persuading audiences. This AI-driven world will enable a single operator to deploy a disinformation campaign with real-time social media sentiment analysis, multi-channel automated content distribution through organic posts as well as ad buys, and contingency-based responses to current events. Add on top of this the potential of "black hat" SEO to punch a story or a topic through the noise floor of social media and make it go viral.

> **Their targeting efforts were reportedly unsophisticated; they simply took advantage of the basic tools of today's information markets that are designed to deliver targeted persuasive messages to tens of millions of people at low cost and with little transparency.**

Unfortunately, this confluence of AI-driven technology and advertising practice means that even poorly executed disinformation campaigns will achieve results because they will benefit from similar, better ones that have taught the algorithms. It is quite possible that the Russian disinformation campaign worked in spite of mediocre tradecraft. Their targeting efforts were reportedly unsophisticated; they simply took advantage of the basic tools of today's information markets that are designed to deliver targeted persuasive messages to tens of millions of people at low cost and with little transparency. Moreover, they benefited from the fact that there were many other domestic political actors doing similar things—running paid and unpaid content on social media to promote salacious, divisive, or emotionally manipulative political messages. Once AI-driven audience targeting has locked onto a successful combination of demographics, messages, and attention-spending user behavior, it will naturally steer all similar content into the same pathways. These platform economics are designed to help advertising succeed.

Disinformation campaigns are functionally little different from any other advertising campaign, and the leading internet platforms are equipped with world class technology to help advertisers reach and influence audiences. That is the business. As such, the economic incentives of the platforms and the political objectives of disinformation operators are aligned. We must grapple with this political economic linkage if we are to make progress.

We conclude by offering a set of high-level principles that we believe should guide the work ahead. We believe a combination of new corporate policy, public law, user expectations, and civic norms in an era of algorithmic disinformation can address this problem. Further, we suggest preliminary directions for regulatory inquiry for ad technology. These are exploratory conclusions in a field that requires substantial technical research and legal analysis. We intend to publish further analysis in the months ahead that builds on the problem statement presented here with more detailed proposals. For the purposes of this

conclusion, we reserve ourselves to offering a high-level structure for the next phase of work. Some of these ideas are reflected in current efforts and have been elevated by other analysts and observers.[87] But we believe the challenge is very large, and therefore our response must be commensurately ambitious from the halls of Capitol Hill to the C-Suites of Silicon Valley to the handsets of everyday internet users.

- **Transparency:** We need to conceive a thorough regime of transparency in political advertising that does not aim merely for parity with old technologies but rather to meet the needs of a digital public sphere. There are at least three angles on transparency to consider: how political advertisers label content for sponsorship disclosure; how users are notified about exposure to known disinformation after the fact; and how platforms make available information about political ads (e.g. sponsor, spend level and targeting parameters). Further inquiry might seek to refine current proposals about labelling political ads in digital media. But more importantly, we should examine how platform companies follow through on a pledge to create a searchable digital advertising database that includes information about sponsorship and targeting to determine its efficacy in curbing bad behavior.[88] Moving beyond voluntary transparency measures, however, any proposal to mandate transparency must be vetted against First Amendment and free expression concerns.

- **Security:** Although cybersecurity has not featured prominently in the disinformation debate, it should be understood as an essential part of the solution. Frequently, disinformation and cyberattacks are a paired phenomenon. Hackers breach email accounts or sensitive files and spill the contents into the media, often commingling truth and falsehood. If sensitive data is insecure and citizens are never sure whether disinformation campaigns are referencing real but stolen secrets or merely fabricated nonsense, we cannot easily navigate

out of this crisis. Moreover, we have no clear norms about how publishers, platforms and consumers will treat digital advertising that promotes this kind of information.

- **Public Education:** The dysfunction in our political media system is driven in no small part by the demand for salacious, divisive, and hyperbolic content. Entertainment will always trump civics lessons when it comes to consumer demand, but we can do more to prioritize the virtue of a well-informed polity in our education system, counteract the tactics of disinformation through digital literacy projects, and give platform users technical options to choose a healthier information diet. We observe numerous efforts along these lines.[89] The best results should be surfaced and scaled to reach as broad an audience as possible.

- **Public Service Journalism:** Our media system did not devolve overnight. It has been a decades-long process during which we have done little as the quantity of objective, public service journalism declined and the market replaced it with infotainment.[90] We must reverse this trend. Part of this project is to confront the role of advertising technology in shifting revenues out of traditional news businesses.[91] But this work also includes a concerted effort to build and sustain journalism—from the local to the global—that does not rely on advertising for its lifeblood. This starts by supporting numerous efforts to fill the void of community newsrooms, but also includes a special focus on investigative reporting.

- **Corporate Responsibility:** The leading internet platforms have responded to this crisis by pledging a number of voluntary moves to push back on disinformation. These include a transparent advertising database, increased human review of ad buys, and support for media literacy work. This should continue. Increased focus could be placed on making tools available to users that can help to verify the trustworthiness of content, to provide

multiple viewpoints on a topic, and to flag disinformation.

- **Consumer Empowerment:** The centrality of consumer data for disinformation campaigns begs the question of how we can empower consumers to gain more visibility and control over what data about them is collected, how it is used, and how they see content in social media based on this data. For example, consumers might be given the option to adjust the signals the algorithm prioritizes to customize a social media news feed. Or they might be provided an option to see the sequential flow of posts from social media contacts without the intervention of the filtering algorithm.

Preliminary Approaches to Ad Tech Regulation

This paper has identified a core unsolved public policy problem: How do we minimize the public harms that come from the exploitation of behavioral data for paid targeting of political content in the service of disinformation campaigns? The likely primary avenues for establishing legal restrictions to address this problem are election law, privacy regulations, and consumer protection law. How those areas of law can or should be applied will vary across national jurisdictions and legal systems, and will in part turn on their interaction with legal protections for free expression. We intend to explore these areas of potential regulation in the next phase of our work, but offer some preliminary thoughts here:

- **Political Campaigns and Elections:** The lowest hanging fruit in this agenda is to require more transparency for campaigns and other political actors that use internet platforms to advertise. These rules exist for campaigns (though they have not been effectively enforced), but the scope of inquiry here should consider the platforms' responsibilities as well. The FEC has recently taken a small step in this direction.[92]

- **Privacy:** The roots of disinformation campaigns draw on behavioral data collection to filter audiences into highly responsive segments that can be isolated and misled. The problem is not objectionable political speech, but rather the exploitation of social data to apply precision propaganda without the knowledge of the user. For this reason, we are focused on the question of whether and how to restrict data collection or ad targeting on political issues and elections-related topics. In addition, there may be useful reforms to the current practice of what constitutes informed consent for the collection and use of data. In the dawning age of AI and autonomous decision-making, it may be that the long absent political will to address the invasiveness of consumer data mining emerges not in response to the harms to personal privacy but to the damage inflicted by behavioral targeting on the body politic.

- **Consumer Protection and Competition Policy:** The sheer size of the user base for the largest internet platform companies—and their market dominance—has raised novel theories of how to analyze and shape their relationship to democracy. They are worthy of careful review. For example, the vertical integration of behavioral data collection and advertising networks in markets with little competition raises questions about how best to inform and protect consumers from harm. These questions, combining concerns about consumer privacy, consumer choice, and the absence of market competition, have recently been raised by European regulators.[93] It is a theme that has also been raised in recent commentary from prominent technology leaders.[94]

- **Freedom of Expression:** Despite our deep concerns about political disinformation, we must be mindful of the privileged role granted to political speech in American law—including anonymous and pseudonymous speech—by the First Amendment and the human right to free expression. There are clear civil liberties and human rights concerns with any regulatory

approach where the state attempts to require platforms to delete or block access to speech-- or to hold them liable for such speech--without due process of law, and in general we do not favor such censorship-based approaches.

Towards a New Political Economy for Digital Media

The simple fact that disinformation campaigns and legitimate advertising campaigns are effectively indistinguishable on leading internet platforms lies at the center of our challenge. They use the same technologies to influence people—reaching a share of the national market with targeted messages in ways that were inconceivable in any prior media form. But if the market continues to align the interests of the attention economy with the purposes of political disinformation, we will struggle to overcome it. The path forward is to explore effective ways to limit the exploitation of personal data—social profiles gleaned from online behavior—for the purposes of precision propaganda, isolating and manipulating audiences with commercialized political disinformation. This could be done through limits on data collection, rules about how it is applied, and measures to increase consumer transparency and control. Our task is to chart a course to a new social contract with technology. The technologies of precision propaganda do not distinguish between commerce and politics. But democracies do.

There are no easy answers, and this has not been done before. But the American political resilience has through the ages hinged on our implicit commitment that markets must take a backseat to democracy. A combination of new policies, corporate practices, technical product features, public education, data security, and citizen empowerment will all be needed to achieve this goal.

Notes

1 Craig Timberg and Elizabeth Dwoskin, Russian content on Facebook, "Google and Twitter reached far more users than companies first disclosed, congressional testimony says," *Washington Post*, October 30, 2017.

2 Testimony of Colin Stretch, General Counsel, Facebook, Hearing Before the United States Senate Committee on the Judiciary Subcommittee on Crime and Terrorism, October 31, 2017; Mike Isaac and Daisuke Wakabayashi, "Russian Influence Reached 126 Million Through Facebook Alone," *New York Times*, October 30, 2017.

3 "Seth Fiegerman and Dylan Byers, Facebook, Twitter, Google Testify before Congress," CNN, November 1, 2017.

4 Issie Lapowsky, "Congress asks tech to face hard truths about Russia meddling," *WIRED*, October 31, 2017.

5 Olivia Solon, "Twitter plans to make political ads more transparent amid Russia revelations," *The Guardian*, October 24, 2017; Issie Lapowsky, "Eight revealing moments from day two of the Russia hearings," *WIRED*, November 1, 2017.

6 See, for example, Yochai Benkler, Rob Faris, Hal Roberts and Ethan Zuckerman, "Study: Breitbart-led right-wing media ecosystem altered broader media agenda," *Columbia Journalism Review*, March 3, 2017.

7 For a broad treatment on the attention economy, see, Tim Wu, *Attention Merchants*, New York: Vintage, 2017.

8 Susan Young, "Getting the Message: How the Internet is Changing Advertising," *Working Knowledge*, May 16, 2000.

9 Keith Kirkpatrick, "Advertising via Algorithm," *Communications of the ACM*, February 18, 2016.

10 "Getting to know you," *The Economist*, September 11, 2014.

11 "The World's Most Valuable Resource is No Longer Oil, But Data," *The Economist*, May 6, 2017.

12 Tim Peterson, "A Google Cookie Replacement Could Upend Online Advertising," *AdAge*, September 19, 2013.

13 Annie Lowrey, "How Online Retailers Stay a Step Ahead of Comparison Shoppers," *Washington Post*, December 11, 2010.

14 Veronica Marotta, et. al., Who Benefits from Targeted

Advertising?, FTC Comment, October 8, 2015.

15 Manoush Zomorodi, "Do You Know How Much Private Information You Give Away Every Day?" *Time*, March 29, 2017.

16 Olivia Solon, "A simple guide to cookies and how to comply with EU cookie law," *WIRED*, May 12, 2012.

17 Erik Larkin, "Are Flash Cookies Devouring Your Privacy?" *PC World*, Oct. 23, 2009.

18 Ana Gotter, "The 5 (+ 5) Best Email Tracking Services of 2017," AdEspresso by Hootsuite, May 2, 2017.

19 Brian Merchant, "How email open tracking quietly took over the web," *WIRED*, December 11, 2017.

20 John Kennedy, "Beacons: Better than display advertising?" *Marketing Tech News*, August 17, 2016.

21 Jeffrey Hightower and Gaetano Borriello, "Location Sensing Techniques," *IEEE Computer*, July 30, 2001.

22 Stephen Wicker, *Cellular Convergence and the Death of Privacy*, Oxford University Press, September 19, 2013.

23 Drew Fisher et. al., *Short paper: location privacy: user behavior in the field*, SPSM '12 Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices, October 19, 2012.

24 Alex Emmons, "Supreme Court Hears Arguments about Cellphone Location Tracking in Landmark Privacy Case," *The Intercept*, November 29, 2017.

25 Kelsey Finch, "Location Tracking: Now Coming to a Government, Employer and Retailer Near You," *The Privacy Advisor*, October 29, 2013.

26 Andrew J. Hawkins, "Uber Wants to Track Your Location Even When You're Not Using the App," *The Verge*, November 30, 2016.

27 David Kaplan, "Overwhelming Number Of Smartphone Users Keep Location Services Open," *GeoMarketing*, April 22, 2016.

28 Lisa Gutermuth, "How to Understand What Info Mobile Apps Are Collecting About You," *Slate*, February 24, 2017.

29 Keith Collins, "Google Collects Android Users'

Locations Even When Location Services are Disabled," *Quartz*, November 21, 2017.

30 Mix, "Googler proves any iPhone app with camera permission can secretly record you," *The Next Web*, October 26, 2017.

31 Laura Hamilton, "How Many Active Connected Devices Does a Home in North America Average?" *CED Magazine*, August 24, 2016.

32 Joshua Koran, "The truth about cross-device tracking," *AdAge*, August 1, 2013.

33 Justin Brookman et al., "Cross-Device Tracking: Measurement and Disclosures, Proceedings on Privacy Enhancing Technologies," *De Gruyter*, 2017.

34 The Federal Trade Commission, "FTC Releases New Report on Cross-Device Tracking," January 23, 2017.

35 Zach Rodgers, "With Atlas Relaunch, Facebook Advances New Cross-Device ID Based On Logged In Users," *AdExchanger*, September 28, 2014.

36 Lance Cottrell, "Browser fingerprints, and why they are so hard to erase," *Network World*, February 17, 2017.

37 Nick Nikiforakis and Gnes Acar, "Web advertisers are stealthily monitoring our browsing habits—even when we tell them not to," *IEEE Spectrum*, July 25, 2014.

38 Dan Goodin, "Now sites can fingerprint you online even when you use multiple browsers," *Ars Technica*, February 13, 2017.

39 Mark Stockley, "Browser fingerprints – the invisible cookies you can't delete," *Naked Security*, December 1, 2014.

40 Amar Toor, "Facebook begins tracking non-users around the internet," *The Verge*, May 27, 2016.

41 Lisa Vaas, "Belgium to Facebook: Stop tracking non-Facebook users or face $267K daily fines," *Naked Security*, November 11, 2015; Natasha Lomas, "Facebook Ordered To Stop Tracking Non-Users In France," *TechCrunch*, February 9, 2016.

42 George Slefo, "Desktop and Mobile Ad Revenue Surpasses TV for the First Time," *AdAge*, April 26, 2017.

43 Randolph E. Bucklin and Paul R. Hoban, "Marketing Models for Internet Advertising," *Handbook of Marketing Decision Models*, July 14, 2017.

44 Dhruv Grewal et. al., "Mobile Advertising: A Framework and Research Agenda," *Journal of Interactive Marketing*, May 2016.

45 Shuai Yuan et. al., "Real-time Bidding for Online Advertising: Measurement and Analysis," in *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, 2013.

46 Ewan Duncan, *Developing a fine-grained look at how digital consumers behave*, McKinsey & Co., Jul. 2013; Cade Metz, "How Facebook's Ad System Works," *New York Times*, October 12, 2017.

47 Aske Christiansen, "The Ultimate Guide to Facebook Ads Interest Targeting Research (Advanced Methods Exposed)," AdEspresso by Hootsuite, March 27, 2017.

48 Devin Guan, "Machine learning is helping martech lead the AI revolution," *AdAge*, June 19, 2017.

49 George Slefo, "LinkedIn Debuts New Targeting Feature for Marketers," *AdAge*, March 1, 2016; Ginny Marvin, "Google To Let Advertisers Upload And Target Email Lists In AdWords With Customer Match," *Marketing Land*, September 28, 2015; Josh Constine, "Facebook Lets Businesses Plug In CRM Email Addresses To Target Customers With Hyper-Relevant Ads," *TechCrunch*, September 20, 2012.

50 Greg Finn, "Facebook pixels get upgrade to track actions & page data," *Marketing Tech News*, April 27, 2017.

51 Ginny Marvin, "Google Rolls Out Similar Audiences for Search and Shopping," Search Engine Land, May 1, 2017; Facebook, Target Facebook Ads to People on Your Contact List.

52 Ingrid Lunden, "Facebook expands dynamic ad retargeting to Instagram and travel sector, ramps up 'lookalikes'," *TechCrunch*, May 10, 2016.

53 Reuters Staff, "Kenya president's election campaign used firm hired by Trump: privacy group," Reuters, December 14, 2017.

54 Philip Bump, "Google's Top News Link for 'Final Election Results' Goes to a Fake News Site with False Numbers," *Washington Post*, November 14, 2016.

55 Kaveh Waddell, "Kremlin-Sponsored News Does Really Well on Google," *The Atlantic*, January 25, 2017.

56 Kevin Roose, "After Las Vegas Shooting, Fake News Regains Its Megaphone," *New York Times*, October 2, 2017

57 See, for example: Jessica Lee, "No. 1 Position in Google Gets 33% of Search Traffic [Study]," Search Engine Watch, June 20, 2013; and Madeline Jacobson, "How Far Down the Search Engine Results Page Will Most People Go?," Leverage Marketing.

58 Lauren Kaye, "95 Percent of Web Traffic Goes to Sites on Page 1 of Google Serps (Study)," Brafton, 21 Jun. 2013.

59 We do not single out Google in this section for any other reason than their dominance in this sector. SEO trade journals rarely discuss any other search product.

60 Anil Dash, "Underscores, Optimization & Arms Races," Medium, November 29, 2017.

61 Tess Townsend, "Google's Share of the Search Ad Market is Expected to Grow," *Recode*, March 14, 2017.

62 Jayson DeMers, "The SEO Industry is Worth $65 Billion; Will It Ever Stop Growing?" *Search Engine Land*, May 9, 2016. This estimate may be too high. But even a much more conservative reading would place this industry on the same level as Google's ad revenue.

63 Erin Griffith, "Bad News for Google Parent Alphabet: The 'G' Will Still Foot the Bill," *Fortune*, August 10, 2015.

64 Moz, "Google Algorithm Change History."

65 See: Google Search Quality Evaluator Guidelines, July 27, 2017.

66 See, for example: Aleh Barysevich, "4 Most Important Ranking Factors According to SEO Industry Studies," *Search Engine Land*, February 3, 2017; Danielle Antosz, "Google Releases the Top 3 Ranking Factors," *Search Engine Journal*, March 25, 2016; SEO PowerSuite, "8 Major Google Ranking Signals of 2017," *Search Engine Land*, July 11, 2017.

67 See: Danny Sullivan, "A deep look at Google's biggest-ever search quality crisis," *Search Engine Land*, April 3, 2017; and Olivia Solon and Sam Levin, "How Google's search algorithm spreads false information with a rightwing bias," *The Guardian*, December 16, 2016; Roger Sollenberger, "How the Trump-Russia Data Machine Games Google to Fool Americans," *Paste Magazine*, June 1, 2017.

68 Kimberly Coleman, "This is how Redditors

Manipulated Google's Image Search Engine," Edgy Labs, December 9, 2016. We have also confirmed this technique with practitioners in the SEO industry.

69 Barry Schwartz, "Unconfirmed Google Algorithm Update May Be Better at Discounting Links and Spam," *Search Engine Land*, February 3, 2017.

70 Ben Gomes, "Our Latest Quality Improvements For Search," The Keyword, April 25, 2017; Danny Sullivan, "Google's 'Project Owl' — a Three-Pronged Attack on Fake News & Problematic Content," *Search Engine Land*, April 25, 2017.

71 Sarah Perez, "Facebook, Google and Others Join The Trust Project, An Effort to Increase Transparency Around Online News," *TechCrunch*, November 16, 2017.

72 Daisuke Wakabayashi, "As Google Fights Fake News, Voices on the Margins Raise Alarm," *New York Times*, September 26, 2017.

73 John Koetsier, "28 Social Media Management Tools, Rated, Scored and Reviewed," *VentureBeat*, April 21, 2015.

74 Rob Marvin and Alyson Behr, "The Best Social Media Management & Analytics Tools of 2017," *PC Magazine*, September 1, 2017.

75 See, for example, Ciler Ay Tek, "Why Machine Learning Is a Game-Changer for Social Media Managers," *AdWeek*, March 8, 2017; and Barry Levine, "This new AI-powered social marketing tool can predict engagement or write the post for you," MarTech Today, March 23, 2017.

76 Erna Alfred Liousas and Jessica Liu, "Social Media Management Solutions, Q2 2017," The Forrester Wave, June 12, 2017.

77 Nitasha Tiku, "Russia's Facebook Ads Will Remain Secret, For Now," *WIRED*, October 4, 2017; Hannah Kuchler, "Facebook says Moscow sought to sow doubt over Trump win," *Financial Times*, October 31, 2017.

78 Liz Morrell, "IBM Launches AI Online Advertising Offering With Watson Ads," *Marketing Tech News*, June 7, 2016; Richard Oldale, "How AI is Changing SEO, Marketing Tech News," June 30, 2017.

79 James D. McCaffrey, "Strong vs. Weak Artificial Intelligence," November 26, 2016.

80 Jack Copeland, "What is Artificial Intelligence?" AlanTuring.Net, May 2000.

81 Kris Hammond, "What is Artificial Intelligence?" *Computer World*, April 10, 2015.

82 Sean Illing, "Why Not All Forms of Artificial Intelligence Are Equally Scary," *Vox*, March 8, 2017.

83 Jason Jercinovic, "The Ethics of Using AI in Advertising," *AdAge*, June 26, 2017; AMA Triangle, Ethics in Advertising in the AI Age, American Marketing Association, October 25, 2017; Mark MacCarthy, "Ethical principles for algorithms," CIO, October 13, 2017; Jonathan Vanian, "How Powerful AI Technology Can Lead to Unforeseen Disasters," *Fortune*, February 6, 2017.

84 VB Staff, "Programmatic Ad Mediation: Stop Sending Your Ad Traffic Over the Waterfall," *VentureBeat*, December 21, 2015.

85 Cade Metz, "Google's AI Wins Fifth and Final Game Against Go Genius Lee Sedol," *WIRED*, March 16, 2016.

86 See, for example, the Partnership on AI and the AI NOW Institute.

87 Claire Wardle and Hossein Derakhshan, *Information Disorder: Toward an interdisciplinary framework for research and policy making*, Council of Europe, September 27, 2017; Alice Marwick and Rebecca Lewis, "Media Manipulation and Disinformation Online," *Data & Society*; Kelly Born and Nell Edgington, *Analysis of Philanthropic Opportunities to Mitigate the Disinformation/ Propaganda Problem*, William and Flora Hewlett Foundation, Fall 2017; and Anamitra Deb, Stacy Donohue, and Tom Glaisyer, *Is Social Media a Threat to Democracy?*, Omidyar Group, October 2017.

88 Facebook has proposed creating a searchable database of ads in advance of the 2018 U.S. elections, and they have made available a tool that permits users to check if they were exposed to known Russian disinformation campaigns in 2016.

89 See, for example: Jason Horowitz, "In Italian Schools, Reading, Writing and Recognizing Fake News," *New York Times*, October 18, 2017; Lindsay Stein, "The News Literacy Project, JWT Team Up to Combat Fake News," *AdAge*, April 10, 2017; and Sophia Boyd, "5 Ways Teachers Are Fighting Fake News," NPR, February 16, 2017.

90 Roy Greenslade, "Almost 60% of US newspaper jobs vanish in 26 years," *The Guardian*, June 6, 2016.

91 See Robert Kaiser, *The Bad News About the News*, Brookings Institution, 16 Oct 2014.

92 Fredreka Schouten, "Federal regulators approve narrow Facebook ad disclosure," *USA Today*, December 14, 2017.

93 See, for example, a recent German government inquiry: Bundeskartellamt, Preliminary assessment in Facebook proceeding: Facebook's collection and use of data from third-party sources is abusive, December 19, 2017.

94 See, for example, Tom Wheeler, *Taming monopolies in the digital age*, Biden Forum, December 20, 2017; and Pierre Omidyar, "6 ways social media has become a direct threat to democracy," *Washington Post*, October 9, 2017.