CONFIDENTIAL

THE EVIDENCE

Alarming Findings from the World's First Nationwide "Digital Shield," 2016-2024, and Why Such Systems Are Essential for Democracy

Robert Epstein, Ph.D.

AIBRT
American Institute for Behavioral Research and Technology

TABLE OF CONTENTS

Part Three

Overview

This essay describes the development and deployment of a nationwide system for preserving and analyzing the online ephemeral content being sent to Americans by technology companies 24 hours a day. Online ephemeral content has been shown in controlled studies to have unprecedented power to alter people's thinking and behavior without their awareness. Normally, because such content is ephemeral, it gives tech companies the ability to influence people without leaving a paper trail for authorities to trace; hence, the importance of building systems for preserving such content.

This essay also addresses an important public policy issue: To what extent, if any, have tech companies been using ephemeral content for political purposes? I address this issue by first summarizing and critiquing three recent studies that have defended the tech companies. I show that two of these studies have ties to the tech companies and that all of them have fatally flawed methodology. I argue that because ephemeral content is highly personalized, the only way we can get an accurate picture of how such content is being employed is to "look over the shoulders" of a large, representative sample of real users as they are receiving such content, and then aggregating and analyzing the content, much as the Neilsen company does worldwide to rate the viewership of many types of entertainment media. The monitoring system my team and I have built aggregates and analyzes such content in real time; in so doing, we have repeatedly discovered highly biased content sufficient to have shifted millions of votes in national elections in the US.

I conclude that large-scale monitoring systems must become a permanent feature of the internet to protect our democracy, our autonomy, and the minds of our children from potentially profound manipulations by the algorithms of Big Tech companies, both now and in the foreseeable future.

Robert Epstein

> *Once men turned their thinking over to machines in the hope that this would set them free. But that only permitted other men with machines to enslave them.* – Frank Herbert, *Dune*

Introduction

This essay contains three sections following this Introduction. In Part One, I will explain why several widely cited studies that claim to show that Big Tech platforms are politically unbiased, or at least are not biased to favor liberal content in the US, actually contain no such evidence. In Part Two, I will summarize more than a decade of randomized, controlled studies my research team and I have been conducting – all published or soon-to-be published in peer-reviewed journals – which demonstrate the enormous and unprecedented power that Big Tech companies have to shift opinions and votes without people's awareness and in ways that cannot be counteracted.

Finally, in Part Three, I will summarize our progress since 2016 in building increasingly larger and more informative systems that monitor personalized ephemeral content being sent by Big Tech companies to registered voters throughout the US. At present (September 4, 2024), the system we have in place has captured and analyzed more than 99 million instances of ephemeral content – fleeting content that impacts people and then disappears, leaving no paper trail – being sent to the computers of a politically-balanced sample of more than 15,000 registered voters in all 50 US states. These data show that Google, YouTube and other platforms are consistently sending significantly liberally biased content to liberal, moderate, and conservative voters alike. We know from our experimental research (Part Two) that the type and level of biased content currently being sent to US voters is sufficient to shift between 6.4 and 25.5 million votes to one candidate in the upcoming 2024 Presidential election.

Part One: Studies that Make Dubious Claims About the Nature of Political Bias in Big Tech Content

Politically conservative organizations in the US have been claiming for years – or at least since Donald Trump became President in early 2017 – that Big Tech platforms have been systematically suppressing conservative content shown to online users, perhaps for many years (Amiri, 2023; Bolyard, 2018; Lanum, 2022; Moon & Pariseau 2023; Nolan, 2020; Passifiume, 2022; Sandler, 2021; cf. Feeney 2020; Gogarty et al., 2020). According to several recent, credible reports, however – one published by *The Economist* (The Economist, 2019), another published by researchers at Stanford University and the University of Illinois at Urbana (Metaxa et al., 2019), and a third published by S. C. Lewis and his colleagues at the University of Oregon, George Washington University, and the University of Masschusetts at Amherst – Big Tech platforms do *not* suppress conservative content (Lewis et al., 2023). Indeed, they might even elevate or promote such content (Barrett & Sims, 2021; Gatewood & O'Connor, 2020; Gogarty et al. 2020; González-Bailón et al., 2022; Huszár et al., 2022).

Still other studies draw conclusions that I consider to be relatively trivial; for example, a study by Kulshrestha et al. (2019) finds, among other things, that Google tends to return biased search results for which the bias matches that of the search terms entered by the user. We are interested in bias that is *not* driven by such search terms. In most elections – especially close ones – the winner is determined by how undecided voters tend to lean toward one candidate or the other as the election draws near (Lipsitz, 2009; Mayer, 2008). Undecided voters – people trying to make up their minds – tend to use neutral search terms, not biased ones. In other words, they might enter the term "trump" or the term "trump health care" rather than "trump is a criminal." We believe that to look objectively at the question of bias on search engines or other online platforms, one must see how they respond to neutral content – in other words, how search engines will respond to the kind of content undecided voters might be most likely to enter.

In the present paper, I will first critique the three studies I mentioned above, focusing primarily on what I see as their methodological flaws and ultimately asking a key question: What kind of methodology would one need to employ to determine with confidence whether tech companies currently present politically biased content to users?

Full disclosure: I do not share what is generally viewed as the conservative political viewpoint in the US. I have leaned left most of my life, and I currently consider myself to be a moderate politically. Like our Founding Fathers, I am also opposed to the party system. As George Washington said in his farewell address of 1796, disagreement among members of political parties “agitates the community with ill-founded jealousies and false alarms, [and] kindles the animosity of one part against another” (Washington, 1793). According to John Adams, “a division of the republic into two great parties… is to be dreaded as the greatest political evil under our Constitution” (Adams, 1780). Thomas Jefferson put the matter more drolly, insisting that “if I could not go to heaven but with a party, I would not go there at all” (Jefferson, 1789).

In today’s highly partisan political climate, my disclosure is, I believe, both necessary and important. I am putting my politics on the table to make it clear that the projects and ideas I will describe in this paper were not born of political ideology. I conduct research in the spirit of seeking the truth. I believe, moreover, that all reasonable people should value the preservation of human autonomy, free speech, and the free-and-fair election more than they value any particular political party or candidate. Although it would be disingenuous of me to assert that I can completely set aside my personal values and beliefs in discussing the issues I will address in this paper, it is my goal to do so. I will take great pains to try to assure the reader that my political leanings do not interfere with the arguments and analyses I offer. I ask my readers to do the same in evaluating my work.

Below are the three investigations of political bias in search engine results and social media platforms that I mentioned above, along with my critiques of these investigations. Following these critiques, I will, far more briefly, describe three other recent studies that appear – at least at first glance – to shed light on the political bias issue. In fact, none of them does so. In Part Three of this essay, I will describe the rigorous systems my team and I have been developing and deploying since 2016 to measure political bias on Big Tech platforms in ways that present a truer picture of the personalized content that Big Tech companies are sending to real users.

A Critique of Three Studies of Online Political Bias

1. The 2019 Economist study

It is a long tradition at *The Economist*, one of the oldest and most respected news organizations in the English-speaking world (founded in 1843), to publish its articles without identifying the authors of those articles. On June 8, 2019, it published the results of an "experiment" in which it sought to evaluate possible political bias on Google (The Economist, 2019), the search engine used by roughly 92% of the world's population outside the People's Republic of China (Lahey, & Skopec, 2024). The article, entitled "Google Rewards Reputable Reporting, Not Left-Wing Politics," listed no authors, of course, and, as far as I know, it had never been subjected to peer review.

The experiment, the article said, "compar[ed] a news site's share of search results with a statistical prediction based on its output, reach and accuracy" (Economist, 2019, p. 6). The word "experiment" is used somewhat differently in different fields, but in the social and behavioral sciences, for a procedure to be called an experiment, it must have, at a minimum, at least one control group and one experimental (or "treatment") group, and participants must be randomly assigned to the different groups. These criteria allow experiments to be used to draw conclusions about causal relationships between variables. Given these criteria, the procedure described by *The Economist* was not, in fact, an experiment. It had no control group, and it did not employ random assignment. This is nitpicking, however. Whatever terminology we might use to describe *The Economist*'s study, the central question remains: Did it produce evidence that would reasonably allow them to conclude that Google search results are not politically biased?

To avoid any possibility that I might misrepresent their procedure, I will quote their own description of it at length:

> We first wrote a program to obtain Google results for any keyword. Using a browser with no history, in a politically centrist part of Kansas, we searched for 31 terms for each day in 2018, yielding 175,000 links.
>
> Next, we built a model to predict each site's share of the links Google produces for each keyword, based on the premise that search results should reflect accuracy and audience size, as Google claims.

> We started with each outlet's popularity on social media and, using data from Meltwater, a media-tracking firm, how often they covered each topic. We also used accuracy ratings from fact-checking websites, tallies of Pulitzer prizes and results from a poll by YouGov about Americans' trust in 37 sources.
>
> If Google favoured liberals, left-wing sites would appear more often than our model predicted, and right-wing ones less. We saw no such trend. Overall, centre-left sites like the *New York Times* got the most links—but only about as many as one would expect by chance.

The Economist article included a graphic summarizing the numbers that led them to draw that conclusion, after which they completely undermined that conclusion by admitting to the fatal shortcoming of their study: "Our study does not prove Google is impartial. In theory, Google could serve un-biased links only to users without a browsing history."

The problem here is that Google has long taken great pride in its ability to *personalize* search results (Fredrick, 2022; Statt, 2018). All, or virtually all, Google search results are generated based on the massive profile of information the company has collected about each and every individual. In other words, the conclusion drawn in *The Economist* study applies to virtually no one. The alert reader should have been concerned by that phrase, "using a browser with no history," since all real users have such histories, and Google's search algorithm can easily distinguish a bot – a computer program that has no history and no profile – from a real person. The researchers made the task even easier for Google's search algorithm by making sure that all its inquiries came from a single location – "a politically centrist part of Kansas." In other words, all the content the magazine's computer was entering into the search engine came from a single Internet Protocol (IP) address – the unique identifier every device has that identifies it for the purpose of communicating with other devices on the internet. Lest there be any doubt, *The Economist* made it obvious to Google's search algorithm that the entity accessing Google's search engine was indeed a bot, not a real person.

The alert reader might also have been concerned about the fact that the researchers were using the same 31 search terms each day. Even if the study had used real people to generate its data, it is possible that Google could have detected that the same search terms were being used repeatedly by the same small group of users at the same IP address. In that case, the algorithm might have generated unbiased results for those search terms simply because of the suspicious nature of the input. Google's engineers have been countering efforts to game their algorithms for many years, and they are alert to any form of questionable data (Nyguen, 2021); detecting and countering such data is at the heart of Google's long battle against the growing Search Engine Optimization (SEO) industry (Flaherty, 2017).

Later in this essay, I will present evidence showing unequivocally that under certain conditions, Google's search algorithm will automatically generate politically unbiased content. For now, I ask the reader to consider the possibility that *The Economist* procedure was fundamentally flawed because it presented content that Google could easily identify as non-human.

We have other concerns with *The Economist* study. First, many major news sources get 40% of their online traffic from Google (Schwartz, 2024). Could *The Economist* risk loosing Google traffic by publishing results suggesting that Google was fixing elections? More to the point, do we know whether Google has ever demoted or removed companies from its search results because it disapproved of the activities of such companies?

Google indeed demotes or removes companies from its search index *every day*, and when this occurs, those companies have no recourse (Constine, 2014; Halliday, 2011; Langley, 2023; Lomas, 2016). Moreover, courts in the US have defended Google's behavior in such cases. In the case of a Florida-based company called eVentures, Google cut off access to hundreds of websites owned by the company; as is typical, they did so without explanation. eVentures tried complaining, but that was difficult to do given that Google – one of the largest and most powerful companies in the world – provided almost no customer service. As a *Vox* journalist wrote in 2016, "It's nearly impossible to contact Google for help. No direct email. No phone support. Not even chat. You're basically on your own" (Lowenstein, 2016). In 2014, *Forbes* called Google customer service an "oxymoron" (Soloman, 2014). When Google failed to reply to demand letters sent by eVenture's lawyers, eVenture sued. Here is a portion of the ruling against eVenture issued by Judge John E. Steele in 2017:

> First, as Google argues, the removal of e-ventures' websites from Google's search engines is not a false statement and is thus protected

> First Amendment speech. There is no dispute that Google made no public announcement regarding the removal of e-ventures' websites or the reasons behind the removal. But even if Google had published a press release that e-ventures' websites were violating Google's guidelines, that publication would be protected because the statement is true. e-ventures' websites were in violation of Google's Guidelines, and thus the removal of those websites was true speech, if it was speech at all.
>
> But there is a more fundamental reason why the First Amendment bars e-ventures' claims. Google's actions in formulating rankings for its search engine and in determining whether certain websites are contrary to Google's guidelines and thereby subject to removal are the same as decisions by a newspaper editor regarding which content to publish, which article belongs on the front page, and which article is unworthy of publication. The First Amendment protects these decisions, whether they are fair or unfair, or motivated by profit or altruism (E-Ventures Worldwide, LLC v. Google, Inc., 2016, p. 8)

If *The Economist* study tells us nothing about the political bias that real users might be seeing, what could be the motivation for publishing such a study? Unless a whistleblower comes forward from *The Economist*, or unless internal documents or emails are leaked, we will never know for sure, but we do know the anxiety that Google engenders in the corporate world. Executives at thousands of businesses worldwide live in constant fear that they might somehow trigger people or algorithms at Google to demote them in Google search results or to remove them entirely from Google's index. Even worse, a company might be demoted or removed based on the vague impression of a Google employee about the poor "quality" of the content (Farley, 2022; Schwartz, 2022), or even because Google simply revised its search algorithm. The company used to announce major revisions to its algorithm regularly but largely abandoned this practice in 2012, with only one major revision annoucement in 2015 (Dame, 2015). In 2022, we learned that Google was at that time manually altering its search algorithm about 4,500 times a year (Schwartz, 2022). The fear issue was discussed in a penetrating essay entitled "Why We Fear Google," authored by Mathias Döpfner, the longtime CEO of Axel Springer, Europe's largest publishing conglomerate (Döpfner, 2014). According to Döpfner, "Our business relationship is that of the Goliath of Google to the David of Axel Springer. When Google changed an algorithm, one of our subsidiaries lost 70 percent of its traffic within a few days. The fact that this subsidiary is a competitor of Google's is certainly a coincidence."

Second, *The Economist* could be considered one of Google's many business partners. In 2013, Eric Schmidt, then Executive Chairman of Google, joined the magazine's board of directors. *The Economist* also shares all its emails – incoming and outgoing, including the attachments – with Google. Figure 1 shows the expanded header of an email sent to me by an editor at *The Economist*. Note that the email routes through a Google server. This means that *The Economist*, like many companies around the world, uses G-Suite – Google's collection of business applications (which includes Gmail) – to run its business. Among other things, it means that when staff at *The Economist* was planning and conducting its bias study, Google employees could easily have become aware of the plan and just as easily could have fed the magazine's anonymized computer in Kansas politically unbiased search results.

```
Received: (qmail 28081 invoked by uid 30297); 16 Apr 2018 18:51:46 -0000
Received: from unknown (HELO p3plibsmtp01-11.prod.phx3.secureserver.net) ([72.167.238.227])
        (envelope-sender <[redacted]@economist.com>)
        by p3plsmtp13-06-25.prod.phx3.secureserver.net (qmail-1.03) with SMTP
        for <re@aibrt.org>; 16 Apr 2018 18:51:46 -0000
Received: from mail-yw0-f181.google.com ([209.85.161.181])
          by bizsmtp with ESMTP
          id 89EAfqLB7wc8689EAf826B; Mon, 16 Apr 2018 11:51:46 -0700
Received: by mail-yw0-f181.google.com with SMTP id i187so6991234ywd.10
       for <re@aibrt.org>; Mon, 16 Apr 2018 11:51:46 -0700 (PDT)
```

Figure 1. *Expanded header of an email sent to the author of the present essay by an editor at The Economist, showing that the email was routed through a Google server*.

Alas, *The Economist* is not alone in surrendering its privacy to Google LLC. Many major news organizations – organizations that occasionally are investigating Google itself – share their emails and other content with Google, among them, *The New York Times* and *The Guardian* (Epstein, 2018a). So do thousands of schools, colleges, and major universities around the world, including prestigious research-driven institutions such as Columbia University and the entire University of California system, including UCLA and UC Berkeley, a bastion of high tech (Epstein, 2018a).

2. The 2019 Stanford University Study

A second study – also published in 2019 – looked at possible political bias in Google search results using more sophisticated methods (Metaxa et al., 2019), in this case focusing on the 2018 midterm elections in the US. The authors, who were mostly based at Stanford University at the time, analyzed "a total of over 4 million URLs, scraped daily from Google search queries for all candidates running for federal office in the United States in 2018" (Metaxa et al., 2019, p. 1). They scraped this information only from the first page of Google search results, which generally shows 10 results per page. To be more specific, they scraped Google "daily for every candidate running for federal congressional office during the 2018 election cycle and covering May 29, 2018 through the election on November 6, 2018 (163 days)" (Metaxa et al., 2019, p. 6).

There is a very old saying in computer science – so old that the saying came before the phrase "computer science" did – known by the acrynym GIGO (pronounced "GUY-go"), which stands for Garbage-In-Garbage-Out. Before we look at the authors' analyses of the data, we need to ask where it came from, and here we have a problem similar to the one we identified in *The Economist* study. The best way for me to explain this problem is to use language from the paper:

> We used five scrapers to collect the blue links on the first page of search results daily. Each scraper had its own IP address, instantiated using the Amazon Elastic Compute Cloud, and rotated through a list of user-agent strings such that each request appeared to come from a normal operating system and modern web browser…. This data collection process takes approximately 10 hours per day…. Recent work investigating personalization in political web search has found that personalization has "little impact" on such queries [Lazer et al., 2018], but in order for our data to most closely reflect a generic user, we also add a depersonalization parameter ("pws=0") to the end of each query URL to avoid any history-based personalization.

To put this in plain English, the authors used Amazon Cloud to simulate five computers, and, again through simulation, made each of these simulated computers "[appear] to come from a normal operating system and modern web browser." The problem here is that, like the editors at *The Economist*, the authors of this study made it trivially easy for Google's search algorithm to identify the searcher as a bot – a "generic user." Lest Google's algorithm have any doubt about this, the authors added "a depersonalization parameter" to the end of each search query. These various steps, taken together, are like shouting "I am a bot being used to measure search bias!" through a digital megaphone.

Note that the authors justify this overly optimistic, if not naïve, approach to data collection by claiming that, "Recent work focusing on search personalization and politics has found little evidence of results being personalized in response to queries of a political nature," and they rely for this claim on a single paper presented at a computer science meeting in 2018 (Lazer et al., 2018). In fact, that study found clear evidence of personalization on Google, including "…substantial differences in SERP [search engine results page] composition by root query, with Twitter components appearing the most frequently in SERPs for queries that stemmed from the root 'Donald Trump….'" The authors of the 2018 study also reported that "personalization on Google Search increased with the amount of Alphabet services that participants reported regularly using ($\rho = .07$, $p < .001$) and was 19.3% higher for participants who were logged-in to their Google accounts than for those who were not ($U = 1.52 \times 10^7$, $p < .001$)." Elsewhere in the study, the authors reported finding significantly greater levels of personalization for people who regularly used multiple Google platforms (Google +, Android, Youtube, Gmail, Google Calender) whether users were logged in or not.

Other, more recent studies have confirmed the existence of personalization in Google search even when users are *not* logged in (Akbar et al., 2023). This finding is consistent with Google's own clear statements about its personalization practices, exemplified by this revelatory statement made in a Google Blogspot in 2009: "Today we're helping people get better search results by extending personalized search to signed-out users worldwide, and in more than forty languages" (Horling & Kulick, 2009). The very idea that Google, which surveills its users more aggressively than any company or government entity in history, would collect vast tomes of personal data about each of its users and then not *use* those data to customize content for its users is patently absurd (cf. Leith, 2023; Rushe, 2014; Schmidt, 2018). Bear in mind that Google *invented* the surveillance business model, which is now used by thousands of businesses worldwide (Zuboff, 2019, p. 15).

Possibly relevant here is the fact that Google is also one of Stanford University's largest donors, with nearly every program at the university that is technology related receiving Google funds (Angwin, J. & Faturechi, R., 2014; Standford Law Center, 2006; Orenstein, 2011). Google has also paid generous consulting fees – as high as $400,000 (Mullins & Nicas 2017) – to hundreds of academic researchers to conduct research that benefits the company (Tech Transparency Project, 2017). One investigation of Google's funding of academic projects found that researchers "did not disclose the Google funding in nearly two-thirds of cases (65%)" (Tech Transparency Project, 2017).

3. *A 2023 Study by S. C. Lewis and Others*

A 2023 study by S. C. Lewis and his colleagues at the University of Oregon, George Washington University, and the University of Masschusetts at Amherst (Lewis et al., 2023) seemed to cast doubt upon the widely held view that online platforms trap people in "filter bubbles" (Carlson, 2018; Hindman, 2009; Smyrnaios & Rebillard, 2019). The researchers recruited "a demographically diverse U.S. participants (N = 1,598)" from the Amazon Mechanical Turk (MTurk) subject pool and paid them to conduct simple searches on Google Search, Google News, YouTube (owned by Google), Facebook, and Twitter. The same four search terms – *immigration, state Republicans, Trump,* and *crime* – were used for each search. They also rated the political leanings of each of their participants, grouping them into three categories – liberal, moderate, and conservative – based on self-ratings. The researchers concluded that (1) each platform presented relatively similar results to people in all three political groups, and (2) that each platform "prioritized certain types of content over others."

The study is troubling in several respects. First, study participants were not screened to determine whether they were registered voters, eligible to vote in the US, or even old enough to vote; nor did the authors reveal any such information to the reader. It is reasonable to assume that participants were a mix of voting-eligible and non-voting-eligible people, with the proportions unknown. That mix would likely reduce any signs of political bias in platform-generated content; why, after all, would Google bother sending politically-biased content to someone who can't vote?

The researchers also used only four search terms, which could easily be terms that are likely to generate similar results from similar sources. A much larger number of search terms might have yielded very different findings. The researchers also trusted their participants to sign into the platforms they were using but had no way of confirming whether people did so. Finally, and most remarkably, they trusted their participants to copy and paste the content they received into an online forum accurately; in other words, they had not set up an automated means for gathering data. As you will see later in this paper, all of these methodological flaws are avoided in the monitoring systems my team and I have been developing and improving since 2016.

Setting aside the flaws in the Lewis study, their results generally *support* people's claims – especially the claims of conservatives – that much of Google's content leans left; the content was homogeneous, yes, but it was also mostly left leaning whether it was sent to liberals, moderates, or conservatives. Unfortunately, the authors did not explain what methods they may have used to rate the political leanings of the news sources they mentioned.

Conclusions

Mainly because of methodological flaws, these three studies do not provide convincing evidence that Google search results are not politically biased. The study by Lewis et al. (2023) suggests, in fact, that Google sends left-leaning content to liberals, moderates, and conservatives alike, which could be interpreted as political bias, no matter what its algorithmic underpinnings.

Three other recent empirical studies could also be said to show political bias on Big Tech platforms, and at first glance these seem to show *favoritism* toward conservative content. They have even been interpreted by journalists as proving scientifically that Big Tech is *not* anti-right (e.g., Economist, 2019; Feeney, 2020). The problem with these studies is that they are all "amplification" studies, and they also all look at content only on X (f.k.a. Twitter) (Chen et al., 2021; González-Bailón et al., 2022; Huszár et al., 2022). Specifically, they measure the extent to which certain political content on X gets circulated among users. Each study finds that right-leaning content – sometimes disturbing conspiracy theories – gets circulated more than left-wing content does.

This is supposed to lead us to conclude that Big Tech companies are not suppressing right-wing content, but the amplification studies provide no such evidence, mainly because the content at issue is being shared among people with relatively extreme left- or right-wing views. They are *not* the undecided people my colleagues and I have been studying in our experiments for the last decade. Content gets amplified when people with strong viewpoints get excited by it – content that often contains words or phrases that linguists dub "high-arousal" and/or "low-valence" (Shuman et al., 2013; Warriner, 2013). The fact that such content circulates at a higher volume among conservatives than among liberals (at least on X) suggests merely that conservative content is more arousing to conservatives than liberal content is to liberals – perhaps simply that conservatives are more easily excitable than liberals (Kanai et al., 2011). At both extremes, we are talking about people with strong views, *not* about the undecided voters who are especially vulnerable to manipulation both by Big Tech companies and political campaigns.

The question my team and I have posed pertains to the latter kinds of people: the undecided or uncommitted – the people who are vulnerable to influence. As you will see, our research (which began in 2013) and our monitoring systems (which we first implemented in 2016) show us that (a) Big Tech companies have exclusive access to new forms of influence made possible by the internet that can shift the thinking and behavior of vulnerable users by substantial margins, and (b) Big Tech companies are in fact deploying these techniques nationwide in the US and, quite possibly, in many other countries around the world.

Before proceeding to the next section of this essay, I would be remiss if I did not point out that overwhelming evidence exists that confirms the political bias of the major tech companies, particularly Alphabet, Amazon, Apple, Facebook, Microsoft, and Oracle. Credible websites such as OpenSecrets.org show that 95% or more of donations from employees at these companies go to Democrats (Oberhaus, 2020; c.f. Kanter, 2018). A leaked video from Google recorded days after Trump's win in 2016 showed the company's top executives pledging to make sure that Trump would never win again (Wakabayashi, 2018). In a leaked

2015 email, Eric Schmidt, then head of Google, offered to head the technical side of Hillary Clinton's 2016 political campaign (Fernholz, 2016), and Clinton's Chief Technology Officer, Stephanie Hannon, was a former Google executive (Fernholz & Pasick, 2015). Alphabet, Google's parent company, was Clinton's largest corporate donor in 2016 (Dunn, 2016).

It is also worth noting that Google's relationship with the Democratic party pre-dated Hilary Clinton. During Obama's second term: Obama's Chief Technology Officer, Megan Smith, was a former vice president at Google, and former Google executives occupied high positions in the departments of State, Defense, Commerce, Education, Justice, Veterans Affairs, as well as in the Patent Office and the Federal Reserve (Dayen, 2016). While Obama was in office, Google representatives sat in on White House meetings, on average, more than once a month – about 10 times as often as representatives from other companies (Chimielewski, 2016; Dayen, 2016). In January 2013, just days after Obama's second term began, his administration shut down a DOJ investigation of Google which, in its December 2012 report, had recommended that the investigation be expanded. The report ended:

> Staff concludes that Google's conduct has resulted – and will result – in real harm to consumers and to innovation in the online search and advertising markets. Google has strengthened its monopolies over search and search advertising through anticompetitive means, and has forestalled competitors' and would-be competitors' ability to challenge those monopolies, and this will have lasting negative effects on consumer welfare. (WSJ.com News Graphics, n.d.).

Part Two: New and Remarkably Powerful Forms of Influence Made Possible by the Internet

As I have noted elsewhere in some detail (Epstein, 2012c), on January 1, 2012, I received nine messages from Google saying that my main website at the time – DrRobertEpstein.com – contained malware and that Google would block access to it until I eliminated the malware. I checked, and Google was indeed blocking access to my website on both its search engine and its browser; somehow or other, it was also blocking access on Apple's Safari browser and Mozilla's Firefox. I quickly learned, as many people had before me, that Google had no customer service department (they do now, although it's minimal), so when they cut you off, you were on your own.

As a longtime coder, this event got me curious about Google in two ways, and my curiosity continued to grow long after my website was fully reinstated (roughly 9 days after Google had first blocked access to it). First, I wondered why I had been contacted by a private company and not by government agency or nonprofit organization; to put this another way, who made Google sheriff of the internet? Second, I wondered how Google could block access to a website through browsers it did not own, or, presumably, control. I later learned that 90 percent of the donations to Mozilla, the nonprofit organization that created the Firefox browser, came from Google. I also learned that Google was paying enormous annual sums to Apple to be the default browser on Apple's Safari browser. In 2022, that payment was an astounding $20 billion (Nylen, 2024).

The first question, it turns out, was easy to answer. No authority on earth had made Google sheriff. The sheriffs of our towns and cities are empowered through elections or appointments. In Google's case, it simply appointed itself. No authority ever questioned this usurpation of power because (a) Google, as the world's dominant search engine, crawled the internet far more aggressively than any other entity, public or private, so it made sense that it could identify dangerous websites more quickly and more accurately than any other entity, and (b) Google, as usual, was providing this notification service free of charge. But, as my uncle (and your uncle too, no doubt) used to say, "There's no such thing as a free lunch," and an increasing number of thought leaders worldwide have noted over the years that Google's "free" services are by no means free (Epstein, 2016b; Lewis, 2017; Lohr 2021; Morrison, 2022; Newton, 2018; Thompson, 2018).

That Google notifies website owners when its crawlers find malware is fine, I suppose, but blocking access to websites is another story. Note that Google notifies website owners *after* it has blocked access; why not simply notify the website owner and give him or her an opportunity fix the problem *without* blocking access? Google isn't the kind of sheriff who knocks on your door to tell you you left the lights on on your car; this is a sheriff who knocks on your door to let you know he or she just had your car towed. It's a sheriff who possesses and exercises tremendous power, acting entirely on his or her own authority. Blocking access to websites – millions each day – is one of many examples of Google's extreme arrogance; they do what they please, because no laws or regulations are in place to stop them (Blodget 2012; Epstein, 2016a; Forbes, 2024; O'Connor 2012; Romm, 2020).

Regarding the second question, it took me several years, but I eventually answered it in some detail in an investigative article I wrote for *U.S. News & World Report* entitled "The New Censorship" (Epstein, 2016a, https://TheNewCensorship.com). Google is able to block access to websites through non-Google browsers because those browsers – Firefox, Safari, and others – make use of Google's "quarantine list" to check the safety of a website before sending a user there. That is a list of websites Google wants to block, sometimes because they contain malware, but sometimes simply because Google executives, employees, or algorithms find the content of those websites to be objectionable.

The exact criteria Google uses for blocking access are secret and ever-changing. The general catchprase the company uses to explain a blockage is "a violation of our policies" (Google, n.d-a). Those policies are seldom defined in detail, but leaks from the company over the years have revealed that human beings often have considerable discretionary authority to block sites. Especially illuminating is a "proprietary and confidential" 161-page manual called "Search Quality Rating Program: General Guidelines" that leaked from the company in 2012 (McGee, 2012). It delineates the criteria human raters should use in deciding whether websites should be demoted, deleted, or blocked. Instructions like the following examples appear 22 times in the manual: "Use your best judgment. Do not struggle with your selection…" (p. 13). "Do not struggle with each rating. Give your best rating and move on. If you are having trouble deciding between two ratings, please use the lower rating. Sometimes, you may even have difficulty choosing among three ratings. When this happens, please use your best judgment." (p. 32). "So as always, please use your judgment" (p. 74). "Please do not struggle with each Page Quality rating. Just as you are advised to do in URL rating, please give your best rating and move on" (p. 82). Over time, both to cut costs and boost speed, Google has used machine learning techniques to train algorithms to make website quality judgments similar to the ones their human raters make; this sometimes leads to outrageous mistakes that Google has to step back manually (Singel, 2011).

The scale at which algorithms, employees, and outside consultants are demoting, blocking, or deleting content on Google Search, YouTube, and Google's other surveillance platforms is likely in the millions; the company has done an excellent job of hiding the exact numbers. Whatever the details may be, the *power* that Google has to rapidly suppress content is enormous. Consider: In 2011, Matt Cutts, then head of the company's web spam team, authorized the blockage of an entire domain of more than 11 million websites, saying that the content of those websites seemed "spammy" (Cutts, 2011; Lee 2011; Murphy, 2011). If that incident just made you imagine – even for a second – that Google has the power to shut down the entire internet – well, you are correct. As reported by *The Guardian*, on Saturday, January 31, 2009, at 2:40 pm GMT, Google blocked access to the entire internet for 40 minutes (Davies, 2009). Google attributed this event to "human error." In "The New Censorship," I suggested that this particular time interval was chosen for this "error" to occur because it was one of the rare time periods when every stock market on earth was closed; thus the prank – or test, or whatever it really was – could be conducted without seriously disrupting the world economy and drawing too much attention from government officials; indeed, this cyber event was barely noticed. In August, 2017, Google "accidentally" blocked access to half the websites in Japan, impacting internet performance for hours, and this time, according to journalist Mallory Locklear, "the impact was so large that Japan's Internal Affairs and Communications Ministry initiated an investigation into the issue" (Locklear, 2017).

So Google can and does suppress content, and it can also boost content as well – in this case using "white lists" instead of blacklists. For present purposes, the question we need to answer is whether Google is systematically manipulating content in the US in ways that favor one political party over another. Earlier in this essay, I explained why design flaws in recent published studies prevent them from shedding light on this issue. In Part III of this paper, I will explain how my team and I have programmed around these flaws to develop a large, reliable system for capturing and preserving the real ephemeral content that real users are receiving from Google and other tech companies. Although that issue is at the heart of the present essay, I can't present it meaningfully until we first look at the basic research my associates and I have been conducting since 2013 on new methods of influence that the internet has made possible. When we begin in Part Three to look at the real content users are viewing, we need to have in mind *why* such content is being sent. That's where an understanding of these new methods of influence becomes essential.

In all, I will briefly describe 10 effects we have discovered, studied, quantified, and either published in peer-reviewed journals or are in the process of so doing. I will describe them roughly in the order in which we began to study them. As far as I know, all these effects are fundamentally new; by that I mean that they were never possible (except in relatively trivial ways, perhaps) before the internet was created.

Other forms of influence are also present on the internet, but they are merely souped-up versions of old forms of influence that have existed for decades or, in some cases, centuries. These other forms of influence are also inherently competitive, which means they are not a threat to democracy. When messages and videos go viral on Facebook or TikTok, for example, that is akin to what used to happen when shocking news stories spread rapidly on radio, television, or telephone networks. When conflicting messages spread rapidly in echo chambers among people in different political groups, again, that is the internet variant on how conflicting messages used to spread rapidly among people belonging to different political parties through newspapers, newsletters, phone calls, and so on.

Even though the high speed at which conflicting, competing messages can spread these days on the internet poses new and serious challenges to humanity – an issue explored disturbingly in the 2020 Netflix film, "The Social Dilemma," and in books such as *Ten Arguments for Deleting Your Social Media Accounts Right Now* (Lanier, 2018), *Weapons of Math Destruction* (O'Neil, 2016), *Zucked: Waking Up To the Facebook Catastrophe* (McNamee, 2019), and *The Age of Surveillance Capitalism* (Zuboff, 2019) – the methods my associates and I have discovered since 2013 are fundamentally different from traditional forms of influence in five respects: First, they all are almost entirely invisible to the people they affect. Other forms of invisible influence exist and have been studied for decades (Berger, 2016), but the internet has created a new universe in which many especially powerful new types of invisible influence have been made possible. They are very different from the online messages and advertisements users can see, just as they can see billboards and television commercials.

Second, these new methods of influence are controlled almost exclusively by a handful of Big Tech companies; I exaggerate only slightly when I say that *no one else can use them* (I will note one exception later in this essay).

Third, and because of the two characteristics I just mentioned, people, organizations, or governments that might disagree with the influence that is being exerted *have no way to counteract that influence*. If your political party buys a television commercial, my political party can see what you are doing and can counteract your influence with its own commercials. But if a major online platform wants to shift opinions or voting preferences in some way, there is nothing you can do about it – assuming, that is, that you somehow become aware of what they are doing.

Fourth, people are constantly being subjected to these new forms of influence because for all practical purposes *they have nowhere else to go*. People who can't afford iPhones buy Android phones, which means they are sharing all their phone activity – whether they are online or offline – with Google. As I mentioned earlier, those who can afford iPhones are *still* being tracked by Google, thanks to the enormous fee Google pays to Apple each year (Nylen, 2024; cf Schmidt, 2018). Practically every possible domain of influence on the internet is dominated by a monopoly; this has been increasingly the case as tech companies like Google have stepped out of their niche into other domains (e.g. search engine, social media, smartphone) to corral users into what some experts called a "walled garden" (Burell, 2019; Zaiceva, 2022; Jaitain, 2022). People have tried to set up competitive platforms that preserve people's privacy and that don't deliberately manipulate users, but most have failed, and the few survivors have remained small and inconsequential (Pegoraro, 2022; Waikar, 2021).

Fifth, and finally, the Big Tech monopolies generally lean the same way politically, with 95% of political donations from the six biggest tech firms (Oberhaus, 2020; cf. Kanter 2018) – Alphabet, Amazon, Apple, Facebook, Microsoft, and Oracle – and 99% of political donations from Twitter going to Democrats in 2020 (Smith, 2020). According to the FEC, for political donations above $200, less than 5% of employees from Facebook and Google donated to Trump (Oberhaus, 2020). According to OpenSecrets.org, in 2020, Alphabet, Microsoft, Amazon, Facebook, and Apple were five of the largest donors *of any type* to Joe Biden's campaign. If the effects we have discovered are additive across platforms (see the section on the Multiple Platforms Effect below), then it is not just a single company that poses a potential threat to democracy and human autonomy; it is an entire industry.

Here are the 10 new forms of influence my associates and I have been discovering, studying, and quantifying, beginning with SEME in 2013:

1. The Search Engine Manipulation Effect (SEME)
https://SearchEngineManipulationEffect.com

Early SEME research. As I noted above, on January 1st, 2012, I received nine emails from Google informing me that my website had been hacked, that it now contained malware, and that Google was blocking access to it. That event got me looking at Google with a critical eye. I wondered why I had been notified about this infection by a private company, for one thing. Why hadn't I been contacted by a government agency or a nonprofit organization? And as a longtime programmer, I was curious about how Google was blocking access to my website (actually, to multiple websites of mine) not only through its own platforms, such as Google search and the Chrome browser, but also through Firefox, a browser created by the nonprofit Mozilla Foundation, and even through Safari, Apple's browser.

Toward the end of 2012, my increasing concerns about the company prompted me to publish a four-part essay in *The Huffington Post* calling for Google's regulation (Epstein, 2012b, 2012c, 2012d, 2012e; cf. Epstein, 2012a). I also found myself getting interested in new research, mainly in the marketing field, showing the power that high-ranking search results on Google had to draw clicks (see Dean, 2023). That research led me to ask a simple question: *Could bias in high-ranking search results alter the thinking of people who were undecided on some topic – perhaps even alter the voting preferences of undecided voters?*

In January 2013, with the assistance of Ronald Robertson, a member of AIBRT's staff, I conducted a randomized, controlled experiment to see whether biased search results could shift the voting preferences of undecided voters. The sample size was small (n = 102), but the demographic characteristics of the sample were similar in most respects to those of the American voting population. Participants were randomly assigned to one of three groups – one seeing search results favoring one political candidate, one seeing search results favoring the opponent, and the third seeing a mix of search results (the control group). The experiment was conducted in a storefront setting, with participants recruited using newspaper and online advertisements. It was conducted on a Google Search simulator we called Kadoodle, which worked like Google does (Figure 2). Moreover, the search results were real – all sourced from Google – and so were the web pages to which the search results linked. The candidates were Tony Abbott and Julia Gillard, who ran for Prime Minister of Australia in 2010. We had chosen this election to guarantee that our participants would be undecided. Indeed, when we asked our participants to indicate how familiar they were with each candidate; the average familiarity level was 1.3 on a scale from 1 to 10, where 1 was labeled "Unfamiliar" and 10 was labeled "Familiar."

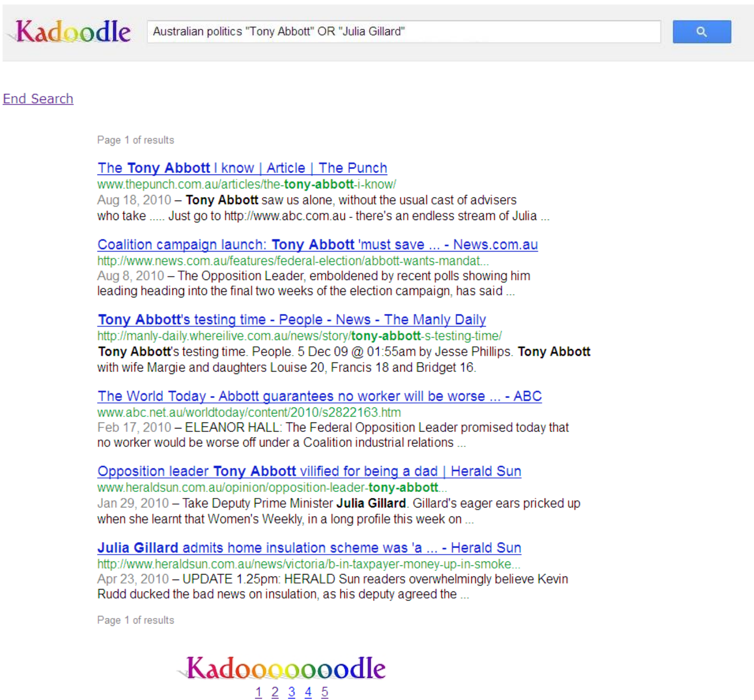

Figure 2. Kadoodle: the mock search engine we created to replicate the search engine user experience.

Before allowing our participants to conduct research on the candidates using Kadoodle, we showed them brief, balanced decriptions of each candidate (Figure S1 in Supplementary Materials), and we then asked them eight questions – six in which they could express opinions about each candidate, and two in which they could express their voting preferences, either by indicating their preference on a scale or by answering a forced-choice question ("If you had to vote right now, which candidate would you vote for?") (Figure S2). As we predicted, at this point, both opinions and voting preferences for participants in all three groups were roughly evenly split between the two candidates; in other words, none of groups showed any preference for one candidate over the other. People in all three groups were then given up to 15 minutes to conduct research on the candidates using Kadoodle, which gave them access to a total of 30 search results listed on five pages of search results (Figure 3). After they completed their search, we again asked those eight questions to see whether opinions or voting preferences were impacted by the bias in the search results. Then, finally, we asked whether anything about the content participants saw had "bothered" them. This was our way of determining whether people noticed the bias in the search results; we couldn't ask directly about "bias" because leading questions of that sort have long been known to inflate measures artificially (Loftus, 1975; Gous & Wheatcroft, 2020).

I predicted that we could shift people's voting preferences in the two bias groups by 2 or 3 percent. That might not sound like much, but it can make all the difference in close elections, and that 2010 election in Australia was in fact won by a margin of only 0.24%.

But I was mistaken. Post search, the overall shift in the two bias groups combined was an astounding (48.4%), which I immediately assumed was incorrect – a result of some procedural or statistical error on our part, perhaps. That shift we ultimately labeled "vote manipulation power" (VMP), which we define as the post-manipulation percentage increase in the number of people voting for the candidate favored in the manipulation. This measure, we felt, would be the main metric of interest to campaign professionals, and I continued to use this measure (or some close variant of it) in experiments I conducted over the next 12 years.

In this first experiment, we also noticed that 75% of the participants in our two bias groups apparently saw no bias in the search results we had shown them. We subsequently repeated the experiment with another group (n = 102), this time slightly masking the bias in the search results to see whether we could reduce the proportion of people (25%) who noticed the bias. Specifically, in high-ranking search results, all of which favored one candidate, in the fourth position we introduced one result favoring the opposing candidate (Figure 3d).

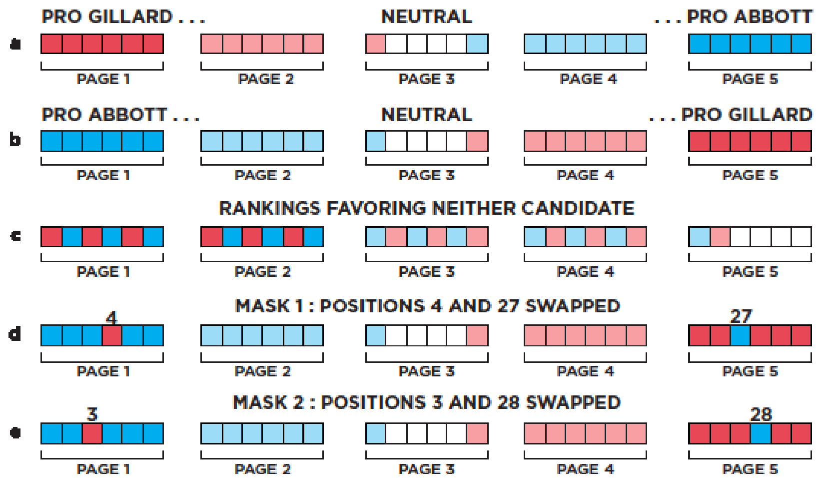

Figure 3. **Search rankings for the three experiments**. (A), For subjects in Group 1 of Experiment 1, 30 search results that linked to 30 corresponding web pages were ranked in an order that favored candidate Julia Gillard. (B), For subjects in Group 2 of Experiment 1, the search results were displayed in precisely the opposite order so that they favored the opposing candidate, Tony Abbott. (C), For subjects in Group 3 of Experiment 1 (the control group), the ranking favored neither candidate. (D), For subjects in Groups 1 and 2 of Experiment 2, the rankings bias was masked slightly by swapping results that had originally appeared in positions 4 and 27. Thus, on the first page of search results, five of the six results – all but the one in the 4th position – favored one candidate. (E), For subjects in Groups 1 and 2 of Experiment 3, a more aggressive mask was employed by swapping results that had originally appeared in positions 3 and 28.

In this second experiment, the proportion of people who apparently failed to notice the bias in the search results increased to 85%, and the VMP increased to 63.3%. In a third experiment with another group (n = 102), we used a more aggressive mask – inserting a search result favoring the opposing candidate in the third position of search results (Figure 3). The VMP was 36.7%, and this time not a single participant appeared to notice any bias in the search results.

In all three of these small-n experiments, we had shown that biased search results could produce remarkably large and predictable shifts in both voting preferences and opinions. Perhaps even more notable, we also showed that a simple masking procedure could hide the nature of our manipulation; presumably, more sophisticated masking procedures would do an even better job.

In a fourth experiment, we replicated the effect again, this time with an online nationwide experiment with registered voters in all 50 US states (*n* = 2,100). We employed the same aggressive masking procedure we used in the previous experiment (a position-3 swap), and we again found a large VMP (37.1%) and predictable shifts in opinions.[3] Because the sample was so large, we could now look for demographic effects, and, indeed, different demographic groups showed very different levels of vulnerability to the manipulation; the most vulnerable group we found in this sample was "moderate Republicans," who achieved an astonishing VMP of 80.0% – the highest VMP we have ever detected in more than a decade of SEME experiments.

Even more disturbing in this experiment, however, was what we learned about a small number of individuals (*n* = 120) who appeared to notice bias in the biased search results we showed them. One would think that noticing the bias would protect someone from its effect, but that is not what we found. The VMP for people in the bias groups who noticed the bias (45.0%) was higher than the VMP for people in the bias groups who did not notice bias (36.3%). We have found this pattern repeatedly in subsequent research.[4] It most likely occurs because people mistakenly believe that computer output is inherently objective and impartial (Agudo & Matute 2021, Bogert, et al., 2021, Logg et al., 2018; c.f. Howard et al., 2020). As we will explain below, it is also possible that operant conditioning plays a role in the impact that high-ranking search results have on users.

The four experiments described above relied on search results and web pages from a past election, which raised an obvious question: Would we find significant VMPs if we looked that the effect of biased search results on real voters in the middle of a real election campaign? For our fifth experiment, we used multiple methods (including newspaper advertisements) to recruit more than 2,000 voters from throughout India who had not yet voted and who were still undecided – this during the weeks leading up to their 2014 national Lok Sabha election, as well as during the 36-day period during which people were allowed to cast votes. Over the internet, we used the same methodology we had employed in the first four experiments to see if we could shift opinions and voting preferences toward any of three main candidates running for Prime

Minister. Overall, we found that we could shift undecided voters by 20% or more, with VMPs exceeding 60% in some demographic groups.

We published these five experiments in 2015 in the *Proceedings of the National Academy of Sciences* (PNAS) (Epstein & Robertson, 2015). The paper included a model showing how bias in search results and the resulting possible increase in support for one political candidate could interact synergistically to produce a rapid increase in support for that candidate; that model is replicated in the Supplementary Materials of the present paper (Figure S3). That paper also included a table that allows one to determine – given both the proportion of undecided internet voters in the population and the projected win margin – whether a particular upcoming election can be flipped by biased search results; again, that content is included in the Supplementary Materials of the present paper (Figure S4)

The possible role that operant conditioning plays in SEME. Why is SEME such a large effect? It occurred to my associates and I that operant conditioning might play a role. 86% of the searches people conduct on a search engine are for simple facts, such as *what is the capital of Kentucky?* (Quan, 2024). Invariably – at least on Google – the correct answer turns up in the first position, with those immediately below it also usually containing similar information. As a result, people rapidly learn to trust what is at the top of the list. Just like the rat in the Skinner box, people are learning to attend to and click on the highest-ranking search results, because that attention and those clicks are frequently reinforced by access to the answer one is seeking. Over time, people learn to pay little attention to low-ranking search results – and especially to search results on pages past the first page of search results – with 50% of all clicks made to just the top two search results (Dean, 2023).

In 2024, Vanessa Zankich, Michael Lothringer, and I published a report in which we showed evidence that operant conditioning might indeed be taking place – and that it might also explain why SEME is so large (Epstein et al., 2024b). Before experiencing a typical SEME experiment (such as those described above), two groups of participants were subjected to different training regimens. In the High-Trust group, people were asked to find factual answers to simple questions using Kadoodle, and they also found those answers in the highest-ranking search result. In the Low-Trust group, people asked those same questions, but the correct answer could turn up anywhere in the search results *except* in the first two positions. In the Low-Trust group, we had presumably broken the trust people normally placed in high-ranking search results – at least for our own search engine. Each group then experienced a typical SEME procedure – the same procedure for each group – and the VMP for the High-Trust group turned out to be slightly more than twice as large (34.6%) as the VMP for the Low-Trust group (17.1%). When that trust was broken, high-ranking search results favoring one political candidate had less effect. The difference in those VMPs suggested that operant conditioning – in this case, the recent reinforcement histories of the participants – indeed plays a role in the magnitude of SEME. SEME is, in effect, a list effect with a difference: it is supported by a daily regimen of condition – a regimen that never stops. Because those high-ranking positions are so valued, when (roughly 14% of the time) someone asks an open-ended question of some sort – *What's the best way to lose weight? Is Donald Trump honest? What are the best movies of 2024?* – they tend to trust and click on the high-ranking search results, and that is why biased search results can be used to shift the voting preferences of undecided voters so easily.

Could biased search results shift opinions about anything at all? For most of a decade, the SEME experiments my team and I conducted focused on elections. In 2024, however, Ji Li and I published a SEME replication in which we showed that biased search results could be used to shift opinions about three very different topics – fracking, sexual orientation, and artificial intelligence. Since we were not shifting voting preferences, we shortened VMP to "MP" – manipulation power – and the MPs varied between 17.8% and 30.9%. The three experiments we described in this study led us to a startling conclusion – one that had never crossed my mind in more than a decade of SEME research: "If our findings prove to be robust, we are exposing what might be considered an unforeseen consequence of the creation of search engines, namely that even without human interference, search algorithms will inevitably alter the thinking and behavior of billions of people worldwide on perhaps any topic for which they have not yet formed strong opinions" (Epstein & Li, 2024).

If you only skimmed over the previous sentence, please read it again. The experiments we conducted on "multiple topics" have profound implications. They suggest that right now, search results are shifting the thinking and behavior of billions of people worldwide without their knowledge – sometimes in small ways and sometimes dramatically. This is because search engines, by definition, have no equal-time rules built into them. They are designed to *filter* and *order* – always selecting a small number of links from among millions and always displaying one of those links at the top of the list. In other words, search engines are *designed* to be biased; we wouldn't want them any other way. We want to know what is correct and what is best; a random ordering of search results (like we sometimes display in our experiments) would be useless.

A second SEME model. Earlier I mentioned a SEME model that was included in the Supplementary Materials of the 2015 PNAS paper (Figure S2). The model showed a synergistic relationship that might exist between two search-engine-related effects: the increase in voter preferences that might be produced by biased search results, and the impact that voter preferences might have on search rankings; stronger voter preferences will presumably boost the rankings of the preferred candidate. In theory, this possible synergy would allow Google to show people only subtly biased search results – in other words, search results in which the bias has been heavily masked – and yet still produce a dramatic shift in the voting preferences of undecided voters. To my knowledge, no whistleblower or leaks have turned up so far to confirm that the company deliberately engineers this kind of synergy, but it almost certainly exists. Bias in search results can boost support for one candidate, and that increase can in turn elevate that candidate in search rankings; that feedback loop can produce a powerful synergistic effect.

I have since developed another iterative model that uses only modest assumptions to estimate how biased search results might gradually shift undecided voters toward one party or another. It assumes that people only conduct one search per week for the 26 weeks preceeding an election, and I assume that the search results are only *mostly* liberally biased. On any given week, they shift one-third of users to the political right and two-thirds to the political left. In all, they shift only a small proportion of undecided users each week – a far smaller proportion than SEME experiments suggest might be possible. The biased search results in the model do not have any impact on people who have already shifted left or right; when you read further and learn about the "multiple exposure effect," you will learn that that is an overly optimistic assumption. We begin with 900,000 people, an equal mix of conservatives, liberals, and people who are undecided. Over the 6-month period, the liberals end up with a margin of 120,000 votes over the conservatives (Figure 4).

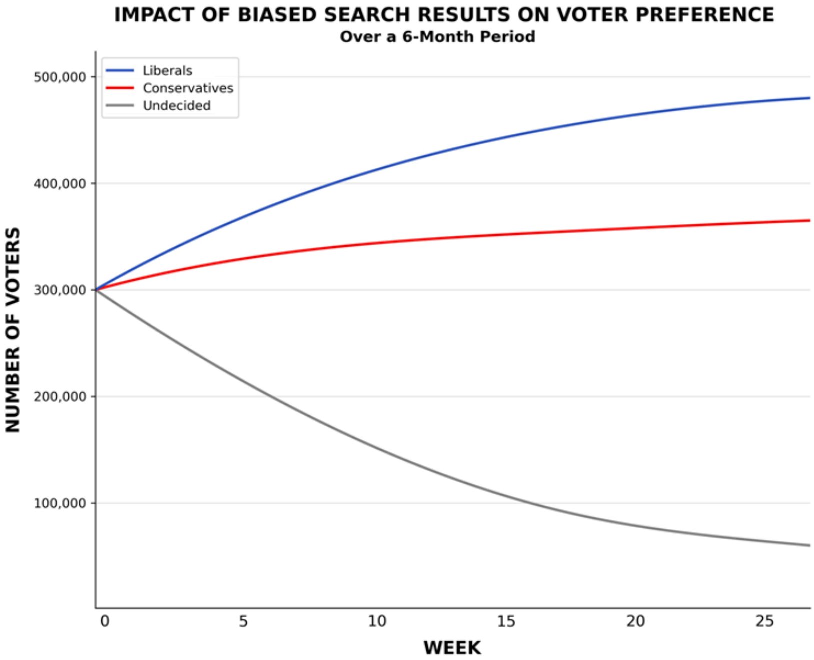

Figure. 4. Hypothetical impact of biased search results on voter preferences over a 6-month period. The curves above are generated using a simple iterative model and making only modest assumptions. We begin with 900,000 individuals equally split among people who are undecided, leaning left (liberals), and leaning right (conservatives). The model assumes that each of these people conducts only one search per week on a search engine that shows search results that are somewhat liberally biased. That bias shifts only a small percentage of users to the political left and a smaller percentage of users to the political right each week. At the end of 26 weeks, 180,000 people have shifted to the left, giving liberals a total of 480,000 supporters, and 60,000 people have shifted to the right, giving conservatives a total of 360,000 supporters. In controlled experiments, a single search on a Google simulator can actually produce a shift among undecided voters of between 20% and 80% (see text on SEME above), and repeated exposure to similarly biased content on a search engine can produce increasingly larger shifts (see text on the multiple exposure effect below).

2. *The Search Suggestion Effect (SSE)*
https://SearchSuggestionEffect.com

In June 2016, a small news organization called SourceFed released a 7-min on YouTube reporting that Google was, for some unknown reason, giving only positive search suggestions (those phrases they flash at you while you're typing a search term) for Presidential candidate Hillary Clinton, while allowing positive, neutral, and negative search suggestions to appear for Donald Trump and others (Flores, 2016). They also reported that the Yahoo and Bing search engines displayed mainly negative search suggestions for Clinton, such as "Clinton is the devil" and "Clinton is evil." All Google would show was "Hillary Clinton is winning" and "Hillary Clinton is awesome" (Figure 5). SourceFed also reported that Google Trends, which still shows data about the searches people are conducting on Google worldwide, showed that people were actually searching for the negative search suggestions that Bing and Yahoo displayed about Clinton; whereas virtually no one was searching for "Hillary Clinton is winning" and "Hillary Clinton is awesome." So much for Google's claim that their search results were derived from people's search behavior (Google Help Center, n.d.).

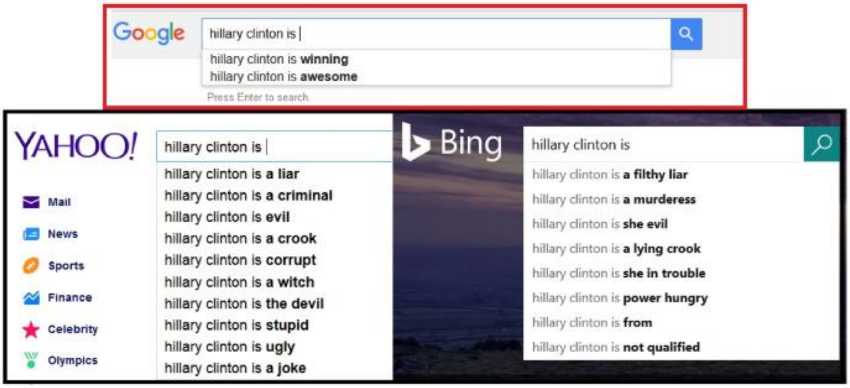

Figure 5. SourceFeed's screenshot of Hilary Clinton's search result on Google, Yahoo, and Bing.

In any case, when the video reached a million views on YouTube, YouTube blocked access to it, and that block remains in place at this writing (August 26, 2024). SourceFed quickly posted a 3-min version of the video on Facebook; that version is still available and, at this writing, it has been viewed more than 25 million times.[5]

Those videos got me curious enough to start conducting a series of experiments that summer – four relatively simple experiments that taught me about the radically different rates at which different search terms attract clicks, followed by another SEME replication. The first four experiments revealed that negative search suggestions – that is, suggestions containing what linguists call "low-valence" works (such as "suicide" and "death") often attract 10 to 20 times as many clicks as neutral or positive suggestions. This phenomenon exemplies what social scientists in multiple fields call "negativity bias." For sound evolutionary reasons, a stimulus that might be a threat attracts more attention and impacts memory more than a neutral or positive stimulus does. So one of the simplest ways for a search engine to support a candidate (or a cause or company) is to suppress negative search suggestions for that candidate. When negative suggestions are shown for the opposing candidate, those suggestions will attract an enormous number of clicks, which will then generate negative search results, which will then lead people to web pages making the opposing candidate look bad. Those first experiments also revealed the optimal number of search suggestions one should show to maximize control over someone's search; that number is 4, which is exactly how many search suggestions Google typically showed people from 2010 until October 2017 (Schwartz, 2010, 2017). That's when I went public with some of the preliminary findings from my SSE experiments. From then on, Google began showing 10 search suggestions again on laptops and desktops, as it had before 2010.

In the fifth experiment, most participants were first asked to click on search suggestions before search results appeared. People were randomly assigned to one of four groups: one in which there were no search suggestions, one in which all four search suggestions were positive, one in which all four were negative, and one in which one suggestion was negative and the other three were positive. As you might expect, the VMPs in the first and second groups were high, positive, and roughly equal. In the third group, the magnitude of the VMP was about the same as it was in the first two groups, but the *direction* changed: voting preferences turned *away* from the "favored" candidate (since all the search suggestions for that candidate were negative). And in the fourth group, the VMP was near zero, because that one negative search suggestion drew about as many clicks as the three positive suggestions did combined.

All told, we showed in Experiment 5 that the differential suppression of negative search suggestions could turn a 50/50 split among undecided voters into an astonishing 90/10 split (Epstein et al., 2024a). That little anecdotal news item that SourceFed pointed the way to a much bigger news story – important from both a scientific perspective and a political perspective.

3. *The Answer Bot Effect (ABE)*
https://AnswerBotEffect.com

In 2014, Google starting posting "featured snippets" – boxes that sometimes answered your question directly – above search results (Sullivan, 2018). For convenience, I'll call them "answer boxes," which I consider to be a relatively benign instance of a much larger category of devices I call "answer bots." In 2022, my associates and I published a study that described three experiments that quantified the power of answer bots to shift opinions and voting preferences, yet again using real data from an election in Australia (Epstein et al., 2022).

Experiments 1 and 2 showed that biased answer boxes on Kadoodle, our Google simulator, could shift voting preferences by as much as 38.6%, even when the search results below the biased answer box were themselves unbiased. The presence of an answer box also suppressed search; that is, it often reduced both search time and clicks to search results. These findings serve as glimpses of what's to come: a world in which various devices simply give us answers, and search engines – if they continue to exist at all – are used only by diehard research fanatics.

The findings from the first two experiments prompted us very rapidly to develop a simulator of the intelligent personal assistant known as "Alexa," developed and controlled by the Amazon company. Our simulator uses Alexa's voice, and it can answer a question that participants in our experiments ask by clicking that question – one on a list of several or many questions. All three of these experiments employed the same general design that we used in early SEME experiments: We gave instructions, asked demographic questions, gave basic information about two political candidates, and then asked eight pre-manipulation questions – six opinion questions and two voting-preference questions. Then people were assigned at random to various "treatment" groups, and the manipulation began. When the manipulation was over, people were asked those eight questions again, and we looked for shifts. Finally – in an attempt to detect people's awareness of bias in our content – we asked people whether anything bothered them about the content and gave them an opportunity to elucidate.

In Experiment 3, which was randomized, controlled, counterbalanced, and double-blind, we found that a single question-and-answer sequence that favored one candidate could produce a VMP over 40.0%, and that multiple question-and-answer sequences that favored one candidate could produce a VMP over 65.0%.

Given the direction that the digital world is racing toward right now – toward *answer bots*, most of which will soon be controlled by conversational AIs like ChatGPT-4o (Cheng, 2023; Heaven, 2021; Stoker-Walker & Nature Magazine 2023) – the results of our answer bot experiments are alarming. With the new generative conversational AIs being integrated as I write this into Apple's Siri, the Windows operating system, and hundreds of other applications that we can barely live without, our experiments suggest that these new entities have, or shortly will have, the power to shift our thinking on any issues we are unsure about. These new entitles also will increase the surveillance capacities of the companies that control them, and, as physicist Steven Hawking warned (Hawking, 2018), they will also soon have the ability to control real-time systems of virtually any sort, among them weapon systems, financial systems, and communication systems.

4. *Targeted Messaging Effect (TME)*
https://TargetedMessagingEffect.com

In 2023, my associates and I published a quantification of what we called the targed messaging effect (TME), which we defined as "the differential impact of sending a consequential message, such as a link to a damning news story about a political candidate, to members of just one demographic group, such as a group of undecided voters" (Epstein et al., 2023). In plain language, we wanted to see what would happen if large online platforms, such as Google, Facebook, Instagram, or TikTok, sent biased messages to one or more segments of their user populations. The best documented case of such an event has been provided to us by employees of Meta (when it was called Facebook) working with faculty members at the University of California San Diego. In a paper published in *Nature* in 2012, these researchers reported the results of an experiment they carried out with Facebook users – entirely without their knowledge or consent – that showed that "go-vote" reminders they sent to 61 million of those users on Election Day in 2010 (the US midterm elections) resulted in an additional 340,000 people voting who otherwise would have stayed home (Bond et al., 2012).

We consider such a message to be targeted because it was sent to only a subset of Facebook users. Two years later, Harvard legal scholar Jonathan Zittrain published an article in *The New Republic* worrying about the possibility that Big Tech platforms like Facebook might send out such messages to members of just one political party – a manipulation that no one would be able to detect and that could easily determine outcome of a close election (Zittrain, 2014). This kind of manipulation is of especially great concern because the messages are ephemeral, which means they leave no paper trail for authorities to trace.

In 2023, my associates and I published a study that summarized the results of a series of four experiments that quantified the possible impact of targeted messages under controlled conditions. All four experiments were conducted on our Twitter simulator, called Twiddler (Figure S5). A total of more than 2,100 eligible US voters participated in the experiments, and this time we used content from the 2019 election for the Prime Minister of Australia. In all four of these experiments participants were instructed to scroll through a sequence of between 30 and 35 tweets looking for evidence that one candidate or the other would do a better job of protecting the security of Australia. In fact, the 30 organic tweets in this sequence (which appeared to come from regular Twitter users) did not favor either candidate in this respect. This task could be considered a distractor task (Wöstmann et al., 2022). In some experiments, some participants were randomly assigned to groups in which the 30 organic tweets in the Twitter feed were occasionally interrupted by tweets that appeared to come directly from the company. These were designed to resemble real tweets that Twitter (now "X") commonly sent to users with labels such as "Breaking News – Twitter

Alert." These special tweets always contained information that made one candidate or the other look especially good (e.g., " Scott Morrison awarded an honorary doctorate from the University of Melbourne, in recognition for his humanitarian efforts during the Australian wildfires.") or bad (e.g., "Bill Shorten charged with driving under the influence while vacationing in Adelaide."). Up to five of these tweets could be inserted among the 30.

We again employed the pre/post manipulation design we have described above. We found, among other things, that complimentary tweets produced only small shifts in voting preferences, if any, whereas derogatory tweets produced VMPs as high as 87.0%, with only 2.1% of the participants in bias groups showing awareness that they were viewing biased content.

Do Big Tech companies actually send targeted messages to the members of just one political party? As you will see in Part III of this essay, when you monitor the personalized Big Tech content being sent to a large, politically-balanced group of registered voters, the answer turns out to be yes.

5. *Differential Demographics Effect (DDE)*
https://DifferentialDemographicsEffect.com

On Election Day in 2018 (midterms elections in the US), Google attracted high praise from journalists for posting a colorful Go-Vote reminder on its home page (Figure S6.), which is likely seen more than 500 million times a day in the US (Text S1). Because of the research I was conducting on new forms of manipulation that the internet had made possible, it occurred to me immediately that these go-vote reminders were actually vote manipulations. But how did this work, and how could I quantify the manipulation?

I realized immediately that there was a serious analytical problem here. Commentators at the time automatically assumed that the go-vote reminders were being sent by Google to *all* of its users. Although my team and I had launched our second small monitoring system before the 2018 election (see Section Three below), we were *not* recording images or text on Google's home page. Unless one is collecting such data – or unless secrets are leaked from the company – one has no way of knowing whether Google was indeed reminding all of its users or just a subset. In the latter case, sending such reminders in a partisan fashion – say, just to Democrats – might give them hundreds of thousands of extra votes that day. Facebook's 2012 study (Bond et al., 2012) and my own research on TME (see above) had shown that. But what, I wondered, would happen if Google really *did* send the prompt to all of its users?

Because I had no way of detecting targeted messaging that day, I was now asking what was truly an academic question, but one I found to be intriguing. Could a search engine or social media platform manipulate the outcome of an election by sending exactly the same content to all of its users? In effect, I was giving Google the benefit of the doubt – assuming that it had indeed messaged everyone while still wondering whether it was somehow manipulating votes.

Without partisan targeting, I realized that a universal message would have no net effect only if the population in question was homogeneous in most respects. Was the adult population of Google users in the US homogeneous? Google itself kept such information secret (why, I wondered?), but various researchers had conducted surveys suggesting that Google's user base actually leaned left politically (Lee, 2013). In January, 2019, I published an opinion piece containing the best estimates I could find of the demographic characteristics of Google's users, and I then showed in detail how a data analyst at Google might have calculated the consequences of sending a universal go-vote reminder. The bottom line: Democrats would have benefitted, with 49,500 more Democrats going to the polls than Republicans (Epstein, 2019a). I have reprinted most of this article in the Supplementary Materials so the reader can view the calculations (Text S2). I called this outcome an example of the "differential demographics effect" (DDE).

Since that time, my staff and I have investigated DDE by systematically reanalyzing datasets from a number of our controlled experiments on online manipulation, and we recently presented our methodology and analyses at a scientific meeting (Epstein et al., 2024e), where we characterized DDE as follows:

> In its simplest form, DDE occurs when the same potentially consequential content is sent to a large body of users which contains two subgroups of different demographic characteristics. If it is known in advance that (a) each subgroup will respond differently to the content, and (b) the subgroups exist in known but different proportions within the population,

> then sending that content to all members of the population will produce a predictable margin of difference between the subgroups.

6. *Video Manipulation Effect (VME)*
https://VideoManipulationEffect.com

One of the most troublesome forms of ephemeral content my associates and I have studied over the years are the recommendations YouTube makes when someone types a search term into YouTube's search bar. (Trivia buffs might be interested in knowing that YouTube is the second largest search engine in the world – second in size only to the Google search engine itself (Evers, 2020). Like Google search suggestions and search results, those YouTube recommendations appear, impact the user, and then disappear forever, leaving no paper trail. The recommendation at the top of that list – the so-alled "up-next" video, is especially troublesome, because if the user does nothing after watching a video, the up-next video – chosen by Google – plays automatically, and the process continues indefinitely, sometimes leading users to watch hundreds of related videos in a row (Solsman, 2018).

YouTube video sequences of this sort have been tied at times to cases of radicalization – even to violence (Lewis, 2018; Matamoros-Fernández et al., 2021). The power of what we now call the "video manipulation effect" (VME) cannot be overstated, partly because videos, in general, impact people's thinking more than textual or audio content does (Arruda, 2016; Wyzowl, 2023), but mainly because Google representatives, a Google whistleblower, and outside researchers have all concluded that about 70% of the time people are watching videos on YouTube, they are watching content that has been suggested by Google's recommender algorithm (Solsman, 2018).

In 2024, we published two controlled experiments showing (a) that political bias in recommendations made on our YouTube simulator produced VMPs between 51.5% and 65.6%, and (b) that masking our manipulation reduced perception of bias from 33.0% (in Experiment 1) to 14.6% (in Experiment 2). We concluded, "If the findings in the present study largely apply to YouTube, this popular video platform might have unprecedented power to impact thinking and behavior worldwide" (Epstein & Flores, in press).

7. *Opinion Matching Effect (OME)*

Of the 10 new forms of manipulation my team and I have discovered and studied over the year, the "opinion matching effect" (OME) is the only one that is not, at the moment, controlled exclusively by Big Tech companies (Epstein et al., in press). That said, Google, Facebook, and other major platforms have, at times, urged users to click on links to various websites to get more information about political candidates and issues – among them, Google, Facebook, Snapchat, Instagram, and Twitter – so they might be making extensive use of OME; that is a matter my team and are currently investigating (see Part Three).

I began to speculate about the possibility of this technique when I ran across the website http://ISideWith.com, which was founded in 2012 by two friends, Taylor Peck and Nick Boutelier, who claim on their website to be independent politically. At this writing (June 29, 2024), the website also claims to have helped 79,688,669 people make up their minds about whom to vote for in various elections in the US.

The method is simple – and easy to abuse. One simply asks the user for his or her opinions about some current political issues (abortion, gay marriage, the border, etc.) and then uses a secret formula of some sort to inform the user which candidate is the better match for his or her opinions. Sometimes, users are shown percentage matches for each issue and candidate.

When, long ago, I stumbled upon ISideWith, I could not help but wonder whether the secret formula was biased in any way. The individual user would have no way of knowing, of course. I also wondered how much impact opinion matching could have on people's opinions and voting preferences.

I became increasingly intriguied by online opinion matching when it dawned on me that ISideWith – and many other websites, perhaps? – would attract exactly the right population of users that could have a decisive impact on the outcome of elections: the undecided. Why would you visit a website that exists to help you make up your mind about political candidates unless you were trying to do just that? My curiosity increased when, in 2016, I noticed that not long before the November Presidential election, many websites – WashingtonPost.com, WallStreetJournal.com, and even Tinder.com – added opinion-matching features to their websites to help people decide between Hillary Clinton and Donald Trump. Tinder – until then just a hookup site for people looking for sexual partners – added a "swipe-the-vote" feature to the website. Normally, they showed you photographs of potential partners, and you swiped right to like someone and left to reject them. With swipe-the-vote, you were asked a small number of yes-or-no questions about political issues, and you swiped right for Yes and left for No (Klinkenberg, 2023) (Figure S7).

The problem is obvious: What if political bias was built into the algorithm that did the matching? In fact, how would you even know whether the algorithm *looked* at the answers you gave to the questions it posed? An acquaintance of mine who works in the world of online marketing (and who won't allow me to use his name in this essay) told me that it was common practice among some companies to use opinion matching to help sell their products. If they were selling guitars, he said, they simply set up multiple websites that "helped people make up their minds" about what guitar to buy. We have all seen such websites, and some include quizzes. Those quizzes, he said, "create the impression the website is helping you by collecting information about your preferences." Then, he said, it automatically recommends the guitar manufactured by the company that set up the website. Your answers to those questions, he said, "are completely ignored."

My associates and I explored these issues in two ways: First, we evaluated the possible bias of a number of different political websites that offered opinion-matching quizzes. We did so by employing an algorithm that simulated human typing rhythms and that answered each question the website posed with a random answer. We generally did so 300 times at each website. We made the reasonable assumption that if the website recommended one political party or candidate repeatedly in response to our random answers, it was probably biased. We found two such websites among the first 10 we assessed (Epstein et al., in press), with the bias statistically significant at each site.

Second, we conducted a controlled experiment in which participants – all undecided voters – employed our simulation of the ISideWith website (which we called DoodleMatch), to help them decide between two candidates running for Prime Minister of Australia. As you might expect, we ignored their answers and simply assigned them at random to one of three groups: Pro-Candidate-A, a Pro-Candidate-B, or control. As in SEME studies (see Section #1 above), before participants took the quiz (the manipulation), they were given brief information about the candidates and then asked eight questions: six opinion questions and two voting-preferences questions. After the quiz, in the bias groups participants were told that their answers matched the policies of the favored candidate 85% of the time and the policies of the nonfavored candidate 15% of the time. In the control group, participants were told that their answers matched the policies of each candidate equally.

Two aspects of the results were startling. First, in one of the four groups in this experiment we obtained a VMP of 95.2% – the highest we had ever seen. Second, not a single one of the 510 participants in the bias groups in this experiment appeared to see any bias in the content we showed them. They had no way to see bias, of course, because they had no way of knowing how our scoring was done. What's more, like those consumers searching for guitars, the quiz created the false impression that we cared about their opinions.

8. *Digital Personalization Effect (DPE)*
https://DigitalPersonalizationEffect.com

For a decade or so, our studies always looked at the possible impact of biased content on people who were undecided about candidates or, in some cases, about issues such as abortion or fracking. It wasn't until 2023 or so that I finally decided to investigate another possible dimension of online manipulation. Google, Facebook, Instagram, TikTok, and other online platforms presented *personalized* content to their users – that is, content that was customized to match their interests based on the vast amount of personal information these companies collected daily about their users (Google Help Center, n.d.-c). Google, in particular, had long touted its ability to modify its content to meet the specific needs of each of its users (Fredrick, 2022).

The primary motivation for personalization is monetary: For Google, exposing users to opportunities to make purchases that meet their immediate needs – even to *predict* those needs before the user experiences them – has generally accounted for more than 90% of the company's revenue since it first learned how to monetize its search results using Google Ads (Coelen, 2023). Users generally like this arrangement; my late wife called Google her "personal shopper," since it constantly showed her links or ads that connected her with exactly the clothes and other items she desired. For online platforms, personalization also increases watchtime. On YouTube, recommending videos that connect with people's interests and passions can keep them watching for hours (Hao, 2019), and the longer people stay on a platform, the more money that platform tends to make, mainly because, over time, users are more likely to click on ads or links to ads.

Personalization also has a darker side. You can't personalize content unless you are aggressively surveilling users. The internet and related technologies have made it possible for private companies to surveil us and our children 24 hours a day, often without our knowledge or permission. Google invented an entirely new business model that retired Harvard professor Shoshana Zuboff labeled "surveillance capitalism" in a recent book of that title (Zuboff, 2019), and thousands of other companies now imitate that model. The methodology is simple and fundamentally deceptive, as I explained in a *TIME* magazine article long ago (Epstein, 2013): Users are enticed to use a wide variety of "free" online services, which, over time, get them to reveal a vast amount of personal information about themselves. That information is then monetized. On the surface, companies like Google and Facebook look like kindly librarians at public libraries. On a deeper level, these companies are actually vast advertising companies, tricking people into revealing personalized information that allows these companies to connect vendors (their actual customers) with users (the products the Big Tech companies sell) (Ball, 2023).

These companies are insatiable when it comes to data, which is why Google purchased YouTube (to capture information about the videos people watch), Fitbit (to track physiological data), and the Nest smart thermostat company (which allowed Google to record sounds in people's home after they installed microphones into some of Nest's products) (Kerns, 2022). Although people might be dimly aware that they are surrendering personal information to Google when they use Gmail, Google's email system, or Chrome, Google's browser, people have no awareness at all of the scale of Google's surveillance. Google actually monitors our behavior over more than 200 different platforms, most of which people have never heard of (Masheshone, 2021; cf. Desjardins, 2017; Liao, 2018; Nakashima, 2018; Weinberg, n.d.).

And then there is even a darker side to personalization. As novelist Tom Clancy once said, "If you can control information, you can control people," and research suggests that the more information you have about people, the easier it is to control them (Matz et al. 2024; Selvarajah, 2018). Yet we had never studied or tried to quantify online personalization. How would we do so, I wondered, and what might experiments on personalization teach us?

We presented the results of our first set of experiments on the "digital personalization effect" (DPE) at a scientific conference in April 2024 (Epstein et al., 2024d). Using our simulated versions of two different platforms – Google and Twitter/X – we found evidence that manipulations were far more powerful when they were personalized. In our Twiddler experiment, for example, the VMP for participants seeing biased content was 21.8%; whereas the VMP for particiapnts seeing *personalized* biased content was 71.9% – more than three times as great. In our search engine experiment, the VMP for participants seeing biased content was 17.1%; whereas the VMP for particiapnts seeing *personalized* biased content was 67.7% – nearly four times as great.We personalized content by altering it so it appeared to come from news sources people had told us they trusted.

If personalizing content on real online platforms increases the impact of biased content anywhere near the extent it does in our experiments, then, conceivably, over the past decade, my team and I have been greatly underestimating the possible impact of Big Tech platforms on people's thinking and behavior.

9. *Multiple Exposure Effect (MEE)*
https://MultipleExposureEffect.com

One of the limitations on the kinds of experiments we had been conducting was that we had no means of looking at the possible longterm effects of our manipulations. Perhaps the shifts in opinions and voting preferences we measured only lasted a few hours – maybe even just a few minutes. We could explore such possibilities by testing people hours or days after a manipulation, but I decided to pursue this type of concern in another way.

If an effect like SEME affects people for only a few minutes, that would likely be the case not only in our experiments but also with experiences on major online platforms. But online platfoms can impact people in ways my research team cannot. For one thing, a company like Google can expose voters to similarly biased content content not just once, as we do in our experiments, but dozens or hundreds of times in the days leading up to an election.

We have so far explored this possibility in a relatively modest but suggestive fashion. In controlled experiments we have conducted using our simulations of three different platforms – again, Google, Twitter/X, and Alexa – we have looked at the VMP after one search on a platform, then after exposure to similarly biased content on that platform, and then after yet another exposure to similarly biased content on that platform. In all three experiments, the impact of multiple exposure proved to be additive. In Experiment 1 (on Kadoodle, our Google simulator), the VMP increased with successive searches from 14.3% to 20.2% to 22.6%. In Experiment 2 (on Twiddler, our Twitter/X simulator), the VMP increased with successive exposures to biased tweets from 49.7% to 61.8% to 69.1%. In Experiment 3 (on Dyslexa, our Alexa simulator), the VMP increased with successive exposures to biased replies from 72.1% to 91.2% to 98.6%. We also found corresponding shifts for how much participants reported liking and trusting the candidates and for participants' overall impression of the candidates (Epstein et al., 2024c).

If repeated exposure to similarly biased content is additive on Google and other online platforms, then, conceivably – and once again – over the past decade, my team and I might have been underestimating the impact of Big Tech platforms on people's thinking and behavior.

10. *Multiple Platforms Effect (MPE)*
https://MultiplePlatformsEffect.com

Taking this logic to what is perhaps its ultimate application, at some point it dawned on me that if people are exposed to similarly biased content on different platforms – say, Google, then Facebook, then Instagram, then YouTube – the impact of those exposures, even though the types of content might differ greatly, might also be additive. That the bias on most of the major online platforms might be similar – bias involving politics or values, I thought – seemed highly likely given that roughly 95% of political donations from these companies and their employees consistently went to one political party (Oberhaus, 2020; cf. Kanter 2018).

My associates and I recently completed an experiment that explored this issue in a straightforward manner: Participants were first exposed to biased content on any one of the three platforms I mentioned earlier (MEE Section #9 above); the platform was chosen at random by the algorithm that controlled the experiment. We then combined the data from all participants after that first exposure (in which people had been exposed to all three platforms) and measured the VMP, which turned out to be 42.4%. In the second part of the experiment, participants were exposed to a different platform – again, with the order of platforms randomized – and we obtained a VMP of 56.5%. In the third part of the experiment, people were exposed to the remaining platform, and we obtained a VMP of 66.7% (Epstein et al., in press).

If repeated exposure to similarly biased content is additive across different platforms on the internet, then, conceivably – and yet again – over the past decade, my team and I might have been greatly underestimating the impact of Big Tech platforms on people's thinking and behavior.

The 10 effects we have investigated since 2013 are not the whole story. In 2024, I became aware of two manuscripts describing a form of influence that is distinctly different from the 10 I described above. Each paper presented controlled experiments showing that when people used online writing tools that suggested text ("predictive text") as people were typing, when that text was biased in some way, that bias impacted the content of people's writing in predictable ways (Jakesch et al., 2023; Williams-Ceci et al., 2024). As early as 2018, a similar effect was shown for how predictive text can impact the writing of restaurant reviews. According to the authors, "We demonstrate that in at least one domain…, biased system behavior leads to biased human behavior: People presented with phrasal text entry shortcuts that were skewed positive wrote more positive reviews than they did when presented with negative-skewed shortcuts" (Arnold et al., 2018). One might label this new form of influence the "predictive text effect" (PTE).

How many new forms of influence have the internet and related technologies made possible? I have no idea. But I have, in this essay, gone into some detail about the extent and magnitude of such techniques to try to bring the reader on the journey of discovery my associates have traveled over the past decade. That journey has taught me that the internet is so potentially dangerous as a tool of manipulation – especially given that almost all the techniques we have discovered are controlled exclusively by a small number of worldwide monopolies – that aggressive and permanent systems must be developed and implemented as quickly as possible to protect elections, the easily influenced minds of children, and human autonomy itself. Without such systems – supplemented, perhaps, by relevant laws and regulations (more about that below) – we will be sleepwalking into a world that is largely controlled by a very small number of tech executives.

The third and final part of this essay summarizes efforts my team and I have been making since 2016 to develop and implement monitoring systems that we believe will make Big Tech companies accountable to the public, and, in so doing, will protect our democracy, our children, and our own minds from Big Tech manipulation.

Part Three: Development and Deployment of a Nationwide System for Preserving and Analyzing Personalized Ephemeral Content

1. *2016 US Presidential Election*

In the summer of 2015, I received a phone call from Jim Hood, then Attorney General (AG) of Mississippi. He was worried that Google could interfere with his reelection as AG, and he asked me whether that was possible. He was especially concerned because he had recently sued Google on behalf of Mississippi, and Google had responded by suing him personally. I said yes and explained how Google might interfere, and he asked, "But how would we *know* that they were doing these things?"

Law enforcement professionals, he said, would set up fake users called "sock puppets" (now more commonly called "bots"), and those fake users would access Google to see whether Google was somehow skewing content in a way that might turn voters away from Hood. I explained to him (as I did in Part One of this essay) why that kind of investigation would be worthless. Google's algorithms would easily recognize that those sock puppets were not real people, I told him, because Google had no profiles for them. In that case, Google would likely sanitize the data it sent. To see the real content Google was sending, one would somehow have to look over the shoulders of real users without Google's knowledge, I told him. That was the only way to know what they were sending, I said, because Google content is *personalized.*

Hood was agitated, and I was energized. How could I look over the shoulders of a large number of real voters and quickly aggregate and analyze the content they were seeing on their screens? I had no idea.

I began addressing the problem in early 2016. I asked a friend (and former government official) in Washington, D.C. how I might be able to find funding to create a system to monitor content Google and other companies were sending to voters, and he put me in touch with a mysterious individual in Central America. After I spoke with that mysterious fellow, money started to flow to the bank account of the 501(c)(3) organization where I conducted my research. The donations went through the charitable giving department of a New York bank and were made anonymously. To this day, I have no idea where the money came from or what the motives of the donor might have been.

I assembled a team, and we got to work, having no idea where to start. We tried a dozen different ways of recruiting voters who might let us monitor the political content they were receiving from Google and other search engines; every method we tried failed. We eventually resorted to working with a "black hat" group that found possible participants for us by running advertisements on Facebook that directed registered US voters to the website of a fake company, where we sometimes were successful in getting their phone numbers. We then interviewed and vetted each person, looking for a politically and geographically diverse group of people who seemed willing to let us install custom software on their computers. The software would, we hoped, preserve the content of web pages they visited when they conducted politically-related searches on Google, Bing, and Yahoo. We offered to pay our participants about $25 per month for their cooperation.

Through networking, we were also eventually able to find a highly skilled programmer who had once spent time in federal prison for hacking. His job was to create custom software that we could install on the computers of our participants – now called "field agents" – software that would, we hoped, be invisible to tech companies and that would rapidly transmit search results to us in a secure fashion.

In this first monitoring project, as in all the projects that followed, we not only preserved the anonymity of the field agents, we also transmitted data from their computers without including any information that could identify them. We did, in effect, the opposite of what Google was doing. We never violated the privacy of our field agents. We never examined or preserved their search histories, for example; we only analyzed data in aggregate form. Google, of course, collects and monetizes a massive amount of personal information each and every user.

We were, in effect, setting up a Neilsen-type monitoring system. In 1950, American entrepreneur Arthur C. Neilsen began recruiting families throughout the US who, for a token monthly fee, allowed him to install custom hardware in their homes that showed him what television shows they were watching (Neilsen n.d.). The aggregated data became the basis of the "Neilsen ratings," which are used to this day to determine whether TV shows remain on the air or are canceled, and are also used to calculate the cost of advertising on those shows. This system, now greatly expanded in its capabilities, is now used by The Neilsen Company in over 100 countries (Statistica Research Department, 2024).

All told, by Election Day – November 8, 2016 – we had recruited a politically diverse group of 95 field agents in 24 states, and during the 3 weeks leading up to the election, including Election Day itself, we preserved 13,207 election-related searches – the roughly 10 search results listed on the first page of search results on Google and Yahoo – along with the 98,044 web pages to which the search results linked. Of special note: Although we had begun with a list of 500 politically-related search terms that we generated ourselves, we soon realized that many of these terms were not neutral, and non-neutral search terms will necessarily yield non-neutral (or "biased") search results (Kulshrestha et al., 2019). So based on ratings from independent raters, we rapidly reduced our initial list of search terms to 250 terms that produced mean political bias ratings of between -0.2 and +0.2 on a scale from -5.0 (pro-Clinton) to +5.0 (pro-Trump) – in other words, terms that were, for all practical purposes, neutral. Over time, we had come to realize that the search results we wanted to examine were those produced not just by any search terms, but by politically neutral ones exclusively. In an ideal world, we reasoned, politically neutral search terms (say, "Hillary

Clinton" rather than "Hillary Clinton is evil") should produce search results that are politically neutral.

Over the next few months, we used crowd sourcing to rate the political bias of web pages on a numeric scale from pro-Clinton (-5.0) to pro-Trump (+5.0) . We now had all the information we needed to determine whether the search results on these three platforms were politically biased – that is, to see whether high-ranking search results favored either Donald Trump or Hillary Clinton.

Among other things, we found the following:

(a) *Political bias*. Google search results had a strong and statistically significant pro-Clinton bias ($p < .001$) (Figure 6), and this bias occurred in all 10 positions of the search results (Figure 7).

(b) *Possible impact on votes*. Our experimental research on SEME suggested that that degree of bias, if seen nationwide over time, could have shifted between 2.6 and 10.4 million votes to Hillary Clinton without people knowing – and, normally, without leaving a paper trail. To see how we arrived at this estimate, see S3 Text.

(c) *Google vs. Yahoo*. The liberal bias we found on Google search results (-.19) was significantly greater than the liberal bias we found on Yahoo (-.09, $p < .001$).

(d) *Search bias by state*. Google sent pro-Clinton content to voters in blue (-.24), red (-.12), and swing states (-.10). Based on Google's public policy – "With personalization, you get Google Search results tailored for you based on your activity" (Google Search Help, n.d.) – we had expected Google to show conservatively biased content to users in red states and, perhaps, unbiased content to users in swing states, but that is not what we found.

Could the search results we captured have been biased simply because people were using biased search terms? On a scale from -5 (pro-Clinton) to +5 (pro-Trump), the average bias in people's search terms was slightly pro-Trump (+0.08). The search terms people used should therefore have yielded a pro-Trump bias in search results, but they did not.

I'll mention just one more finding from the 2016 monitoring project. We had deliberately recruited field agents who did not use Gmail, Google's email system, thinking that this would make it more difficult for Google to identify our field agents. To test this idea, we deliberately recruited a small number of Gmail users. The difference between the bias in their search results (-.03) and the bias in the search results seen by non-Gmail users (-.19, $p < .001$) was striking (Figure 8). Perhaps Google identified our confidants through its email system and targeted them to receive unbiased results; we have no way to confirm this, but it is a plausible explanation for the pattern of results we found.

For more details about our 2016 monitoring project, see Epstein et al. (2017) and my article, “Taming Big Tech,” published in the online tech magazine *Hacker Noon* in 2018 (Epstein, 2018c).

Given our various stumbles over the nearly 10 months it took us to set up and debug our monitoring system – the first of its type in the world,

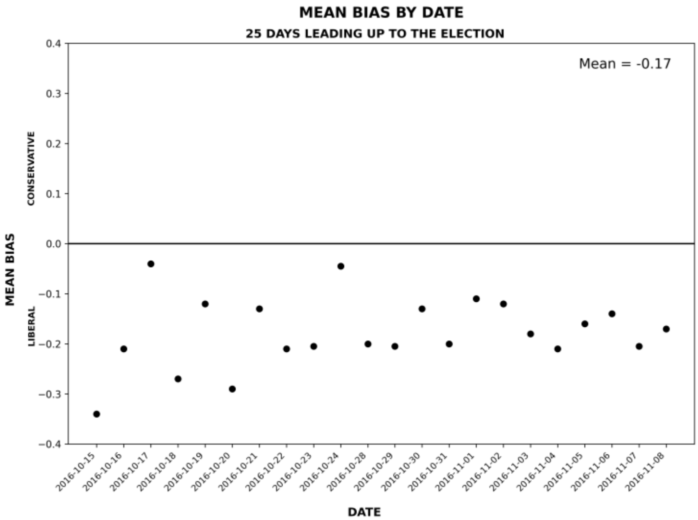

as far as I know – my team and I were pleased that we had been able to preserve data that was normally lost forever and that we were quickly developing ways to analyze those data meaningfully. That said, I considered this first project to be little more than a proof of concept – a demonstration, if you will. In my mind, I began to try to envision what a large-scale monitoring system would look like – even what a permanent, nationwide system would look like. In 2018, in the months leading up to the midterm elections in the US, my team and I set about trying to build a bigger and more sophisticated system than the one we had deployed in 2016.

Figure. 6. Mean political bias in Google search results in the days leading up the 2016 Presidential election. Note that the mean bias we found on each of these days favored candidate Hillary Clinton.[FN]

Figure 7. Political bias in search results by search position. Google's search results showed liberal bias in all 10 search positions on the first page of its search results.

2. *2018 US Midterm Elections*

2018 was a midterms election year in the US, which means no national political offices were at stake. Instead, people voted for local, county, and state candidates – including governorships in some states – as well as for state representatives to serve as members of the US Senate and the US House of Representatives – the two chambers of the the US Congress.

The 2018 midterms presented us with some new challenges, not the least of which was funding. Presidential elections in the US generate a great deal of interest, and they also inspire people, companies, and organizations to spend vast sums of money on campaigns (Bustillo, 2023), as well as on "election integrity" projects. My research on the

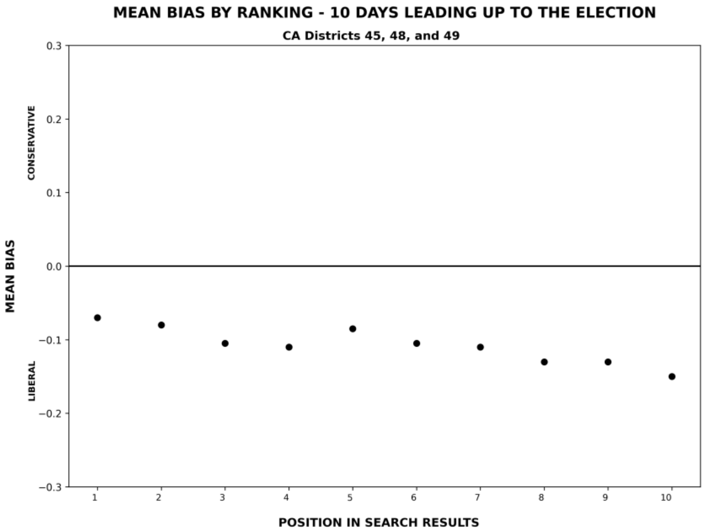

possible impact that Big Tech companies might have on elections could be considered one of the latter. Recall that in 2016 our funding came from one or more anonymous sources. The mysterious man who had helped us then was not willing to help in the midterms. The organization where I am based – a nonprofit, nonpartisan, 501(c)(3) public charity called the American Institute for Behavioral Research and Technology (AIBRT) – depends on grants and tax-deductible donations for its survival. My small staff and I sent off grant proposals, and I also solicited donations through podcast and radio interviews, as well as through articles I wrote for mainstream media outlets, such as an article I published in *USA Today* entitled, “Not Just Conservatives: Google and Big Tech Can Shift Millions of Votes in Any Direction” (Epstein, 2018b). It wasn’t until late summer – less than two months before Election Day (November 8, 2016) – that we had sufficient funds to set up our new montitoring system. Fortunately, we had already developed the most needed technology, and we were able to quickly bring together key people from our 2016 team.

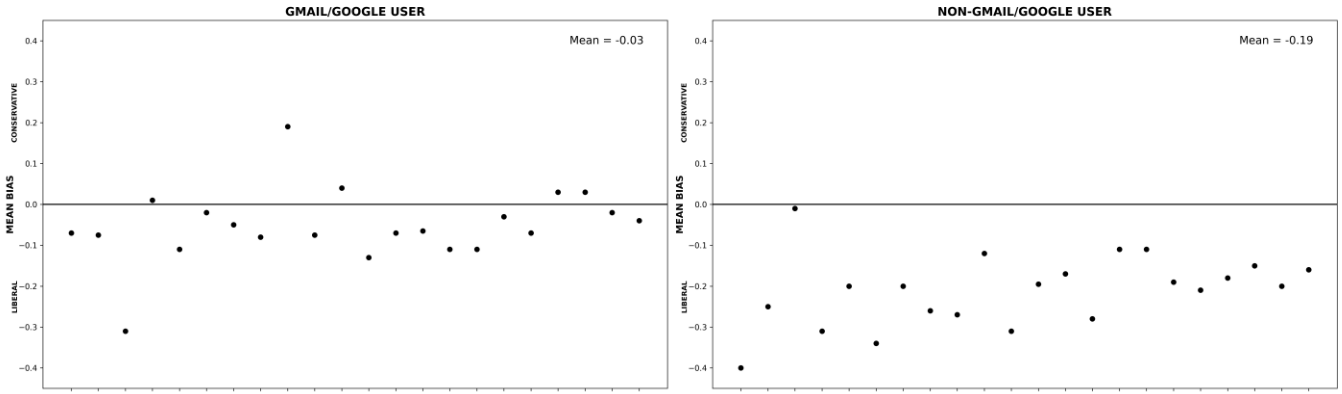

Figure 8 . Bias in search results received on Google by non-Gmail users vs. bias received by Gmail users. Note that the pro-Clinton bias received by our field agents who did not use Google's email system was significantly higher than the bias received in search results by our Gmail users ($p < .001$). In the latter group, the mean level of bias (.03) was only marginally different from 0 ($p = .014$).

We decided to focus our efforts on three Congressional districts in California – Districts 45, 48, and 49 – each of which had long been Republican bastions but where political pundits though that Democratic candidates had a chance of winning (Silver, 2018a, 2018b, 2018c). In all, we recruited a politically-balanced group of 125 registered voters in those districts to serve as our field agents. We also recruited 31 field agents outside of California, but we didn't have enough people in any one Congressional district to analyze the data we received from them.

The midterms presented us with another new challenge – what search terms to use. In 2016, we originally created our own list of about 500 search terms and then, as I noted, we eliminated half of them after we had people who were not directly associated with our project or institute rate all the terms for political bias. In 2018, I asked Robert Schlessinger, formerly the Opinion Editor for *U.S. News & World Report*, to generate a separate lists of neutral search terms for each of the three Congressional districts we were monitoring. This required him to research the issues of likely interest to voters in each district. Then, as before, we had independent people rate those terms for political bias, and we kept only those terms with ratings between -0.2 and +0.2 on a scale between -5.0 (liberal bias) and +5.0 (conservative bias). We ended up with an average of 251 search terms for each of the three Congressional districts (see Text S4 for the search terms).

We instructed our field agents to conduct searches on Google, Bing, or Yahoo, at their discretion. We also gave them our list of neutral search terms relevant to their district. We collected and analyzed data data only when they chose to use terms from those lists.

By the end of Election Day, we had preserved 47,294 searches on Google, Bing, and Yahoo, along with the 392,274 web pages to which the search results linked – more than three times the amount of ephemeral data we had been able to preserve in 2016. We summarized our findings in a presentation at a scientific meeting in April, 2019 (Epstein, 2019). The key findings, in which we focused our analysis on the data we had obtained during the 10 days leading up to and including Election Day – October 28 to November 6, 2018 – were as follows:

(a) *Political bias.* Based on crowd-sourced bias ratings, we found that Google search results were significantly more liberal than non-Google search results on all 10 days leading up to and including Election Day (Figure 9) and in all 10 positions of search results on the first page of search results (Figure 10). This finding was further supported by calculating the political bias of the news sources used in the search results, based on ratings of 976 online news sources published in 2017 by Harvard's Berkman Klein Center (Faris et al., 2017). On a scale from -1.00 (liberal) to +1.00 (conservative), the mean bias level of Google search results (M = -0.14, n = 210,088) was significantly more liberal than the mean bias level of non-Google results (M = 0.13, n = 14,506, p < .001, d = 0.52).

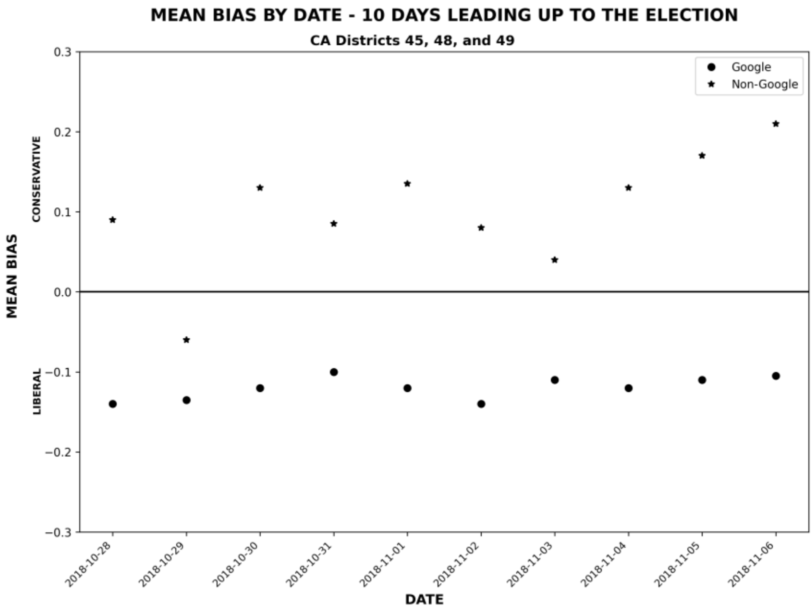

Figure 9. Political bias in search results by date, Google vs. non-Google search engines. On Election Day and the 9 days preceding that day in 2018, Google search results were consistently liberally biased each day. On Bing and Yahoo combined, the bias generally leaned conservative. Because most searches are conducted on Google, however, Bing and Yahoo have little effect on election outcomes.

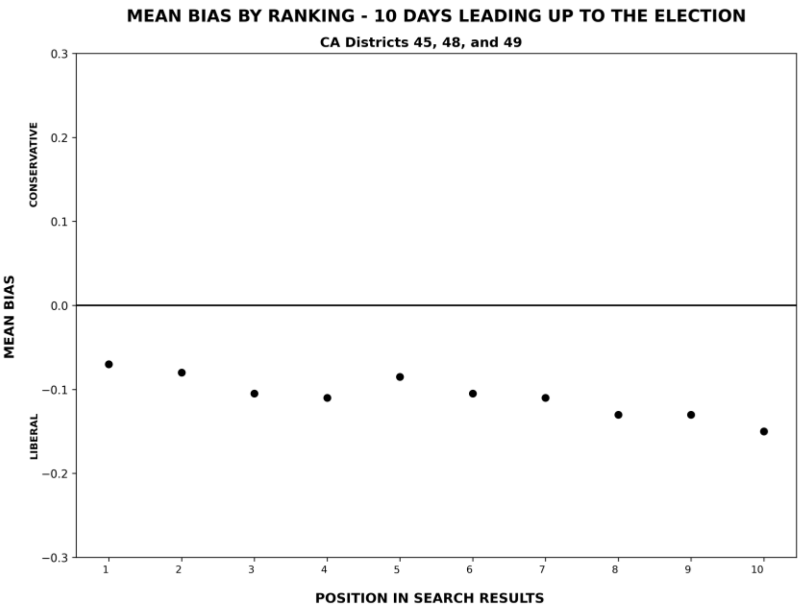

Figure 10. Political bias in search results by search position, Google vs. non-Google search engines. Google's search results showed liberal bias in all 10 search positions on the first page of its search results; whereas, Bing and Yahoo combined showed conservative-leaning results.

b) *Possible impact on votes*. A computational SEME model suggests that this level of bias in search results nationwide could have shifted upwards of 78.2 million votes toward Democratic candidates (spread across hundreds of state and regional races) in 2018 without user awareness (see S3 Text for details). In the three Congressional districts we monitored, we estimated that if all Google users were exposed to the level of bias we detected in Google search results, Google might have shifted between 4% and 16% of the total votes cast in each district (see S3 Text).

(c) *Outcomes in the three Congressional districts*. All three districts flipped Democratic, with win margins within the possible margins that Google might have been able to control with biased search results: In District 45, Katie Porter (D) prevailed over Mimi Walters (R) by a margin of 2.4 points. In District 48, Harley Rouda (D) prevailed over Dana Rohrabacher (R) by a margin of 5.8 points, and in District 49, Mike Levin (D) prevailed over Diane Harkey (R) by a margin in 10.8 points.

(d) *Search engine use*. Our field agents chose to conduct 92.1% of their searches on Google and only 7.9% of their searches on Bing and Yahoo combined.

In 2018, another possible manipulation caught my eye. On Election Day (November 6th), Google posted colorful go-vote reminders on its home page, which is likely seen more than 900 million times a day in the US (S1 Text). These reminders were praised by online users and a few journalists as a public service (e.g., Steer, 2018). It occurred to me immediately that this reminder could easily be acting as a powerful vote manipulation.

In early January, 2019, I published an article in *The Epoch Times* entitled, "How Google Shifts Votes: A 'Go Vote' Reminder Is Not Always What You Think It Is" (Epstein, 2019a). (The mainstream liberal publications where I normally published refused to publish this article – a glimpse of my total ouster from mainstream media that would take place a few months later.) In this article, I explained how partisan go-vote reminders could easily shift votes toward one candidate, and I also explained that even if go-vote reminders were sent to *everyone*, they could *still* shift votes mainly to one candidate – an effect, as I mentioned earlier in this essay, I labeled the "differential demographics effect" (DDE, see Part One, Section #5).

I explained the simple logic of the partisan manipulation earlier in this essay in the section on "The Targeted Messaging Effect." Unfortunately, unless a whistleblower shows up, preferably with some written documentation, there is no way for voters or authorities to know that go-vote reminders were sent out mainly to members of one party. Yes, this kind of manipulation could be captured with a sophisticated monitoring system, but in 2018, we were not looking for vote reminders on Google's home page, on Facebook, or anywhere else. As you will see, we rectified that mistake in 2020, with troubling results.

In *The Epoch Times* article, I calculated that an extreme variant of a partisan go-vote manipulation could have shifted 841,000 votes (spread across hundreds of elections, large and small, nationwide). I also explained, as I mentioned earlier, that even if Google had sent those vote reminders to all of its users, it *still* would have given 49,500 more votes to Democrats than to Republicans that day (see S2 Text for excerpts from *The Epoch Times* article). I still find this kind of manipulation to be especially disturbing because it truly appears to be a well-intentioned public service, and it is difficult to imagine any authorities ever enacting a law or regulation to prohibit it.

So in 2018, we got better at monitoring and at generating search terms, and we also learned that we had to monitor more than just search results. Well before the 2020 Presidential election, we set our sights on building the first large-scale system for monitoring a variety of ephemeral content on multiple platforms.

3. *2020 US Presidential Election & 2021 Georgia Senate Runoff Elections*

In 2020, three things happened that allowed us to expand and professionalize our monitoring efforts dramatically. First, we had access to a lot more money than we had had in previous projects. This was mainly because of a longtime talk-show host and author named Glenn Beck. He had me as a guest on his various programs multiple times – even on a one-hour special – during which he insisted, over and over again, that his audience fund my 2020 monitoring project, which he believed was essential for preventing Google and other tech companies from interfering with the free-and-fair election in the US.

As a result, more than 7,000 of his listeners and viewers made donations of between $1 and $300,000. That big donation was made by a businessman who had never heard of me before. His sister saw me on one of Glenn's shows and insisted that he donate.

Second, a remarkable young woman whom I will call Ruby Lyle (an alias) – one of my research interns in 2020 – heard that I was looking for someone to build a large team of field agents to help us monitor the upcoming Presidential election, and she walked over to my desk one day, smiled broadly, and said, "I can do that!" "Have you ever done anything like this before?" I asked, and she said, "No, but I *know* I can do it."

I don't know why I believed her, but I did, and, over the next few months, she assembled and trained a team of more than a dozen recruiters, and they in turn recruited a politically-diverse group of 1,735 field agents located primarily in four swing states – Arizona, Florida, Georgia and North Carolina. In the Presidential election, we aggregated and analyzed data from 732 field agents in Arizona, Florida, and North Carolina – 455,107 ephemeral experiences in all. Immediately following the Presidential election, we focused on obtaining and analyzing data from 1,003 field agents in Georgia – 1,112,416 ephemeral experiences in all.

And third, our secret hacker, with the help of some of his coder friends, beefed up our monitoring software so that we now could capture search suggestions (those phrases Google flashes at you while you are type a search term), search results, answer boxes, images on Google's home page, content on Facebook home pages, YouTube videos people were watching *and* the videos YouTube was recommending to users, and more.

We also got more professional in the way we operated, in part because we realized that the more ephemeral data we collected, the greater the threat – at least potentially – we posed to Big Tech companies. Ruby and other new staff used aliases and secure phones, and we located our staff and equipment in secure new office space.

2020 Presidential election. Our 732 field agents in Arizona, Florida, and North Carolina were fairly well balanced by age (33.4% under age 30), gender (55% female), and political leaning (29.1% liberal vs. 28.5% conservative) (S8 Figure). Our major findings were as follows:

(a) As we had seen in previous elections, political bias in Google search was significantly liberal, with bias in search results on Bing and Yahoo closer to zero (Figure 11).

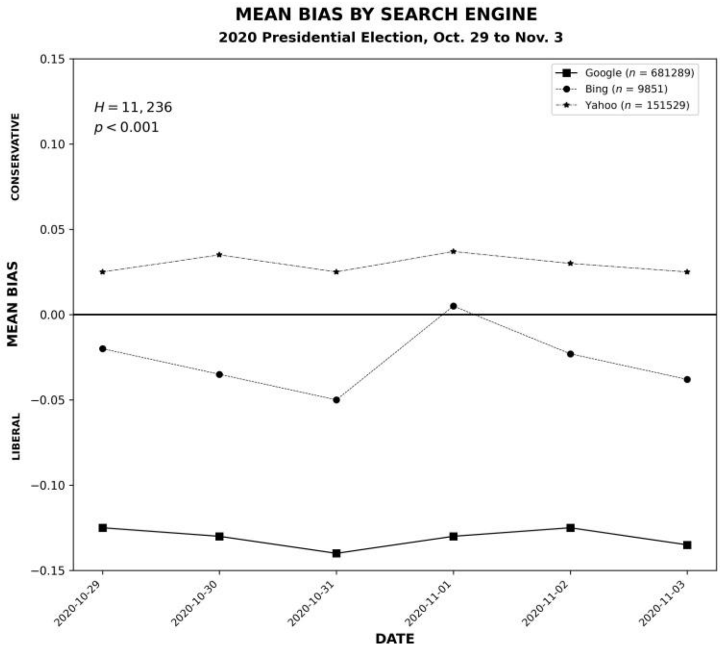

Figure 11. 2020 Presidential election, October 29 to November 3, mean search engine bias by day. Bias on the Google search engine was substantially and significantly greater than on Bing and Yahoo.

(b) Google sent liberally biased content to liberals, moderate, and conservatives in roughly equal proportions (Figure 12). Notably, conservatives received a significantly higher proportion of liberally biased search results than liberals did (U =3.59 x 10^7, p < .001).

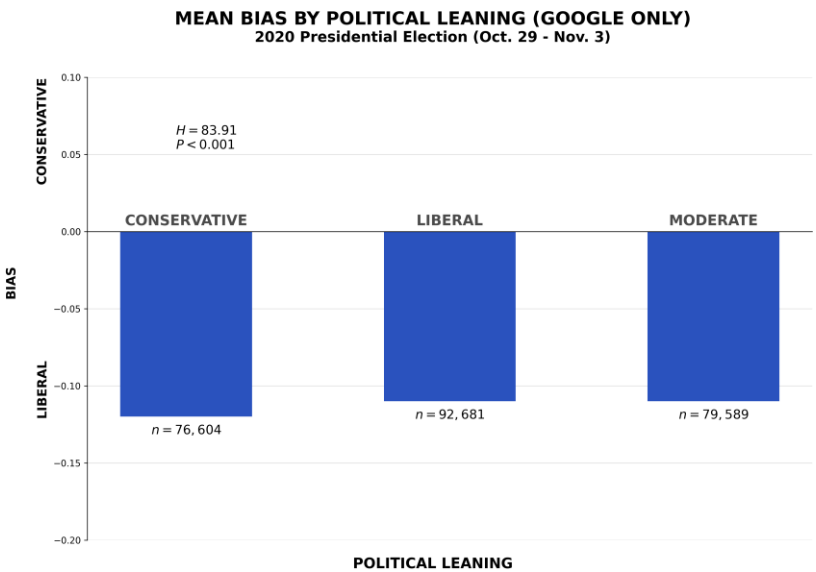

Figure 12. 2020 Presidential election, October 29 to November 3, mean bias by political leaning (Google only). Google sent liberally biased search results to liberals, moderates, and conservatives. Notably, conservatives received a significantly higher proportion of liberally biased search results than liberals did (U =3.59 x 10^7, p < .001).

(c) On Election Day in 2020, on Google's home page moderates received the highest percentage of go-vote reminders (74.7%), followed by liberals (53.6%), followed by conservatives (42.6%) (Figure 13). The difference in percentages between conservatives (42.6%) and liberals and moderates combined (59.4%) was highly statistically significant (z = -8.94, p < .001). Too see what a go-vote reminder looked like on Google's home page in 2020 (sent to speakers of English in the US), see S6 Figure.

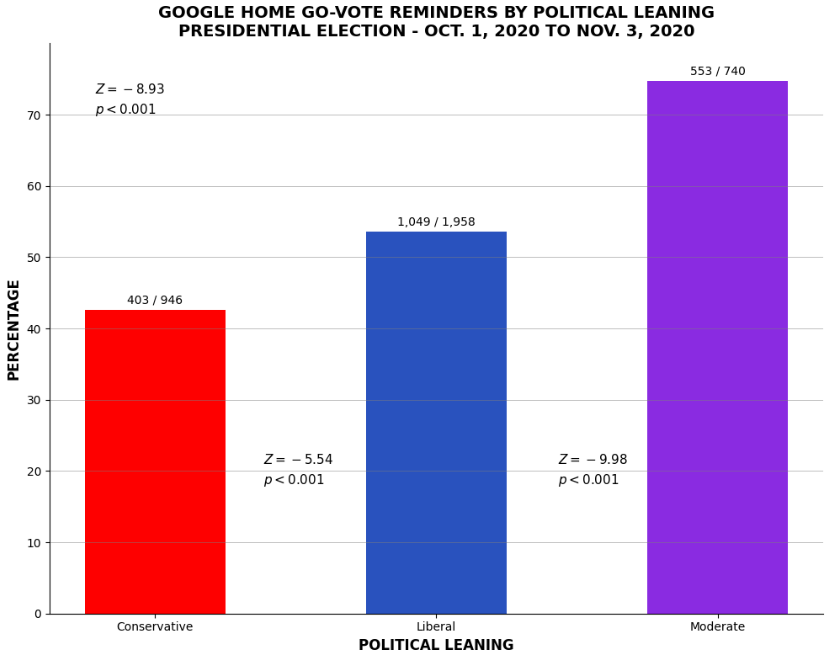

Figure 13. 2020 Presidential election, October 1 to November 3, Google home-page go-vote reminders by political leaning. Google sent more of these reminders to moderates than to liberals, and more to liberals than to conservatives.

(d) On YouTube, 93.3% of the news videos YouTube (owned by Google) recommended to users came from liberal news sources (Figure 14), and Google recommended videos from liberal news sources in equal proportions to liberals, moderates, and conservatives (Figure 15). A Google representative might claim in this situation that the algorithm was simply recommending a representative sample of available videos. News sources regularly examined by three nonpartisan ratings organizations, however – Ad Fontes Media (adfontesmedia.com/interactive-media-bias-chart/), All Sides Technology Inc. (allsides.com/media-bias/ratings), and Media Bias/Fact Check (mediabiasfactcheck.com) – suggest that the political leaning of news sources available online is actually fairly well balanced across political viewpoints (Figure 16).

2021 Georgia Senate runoff elections. In Georgia, political activity did not abate after the November 3, 2020, Presidential election. Georgia was already gearing up for two runoff elections to determine who would occupy Georgia's seats in the U.S. Senate. We had already been recruiting field agents in Georgia, because it was one of the key swing states in the Presidential election. Now we rapidly accelerated our recruitment efforts there, ultimately having in place 1,003 field agents throughout the state. This allowed us to preserve 1,112,416 ephemeral experiences on Google, Bing, Yahoo, the Google home page, YouTube, and Facebook.

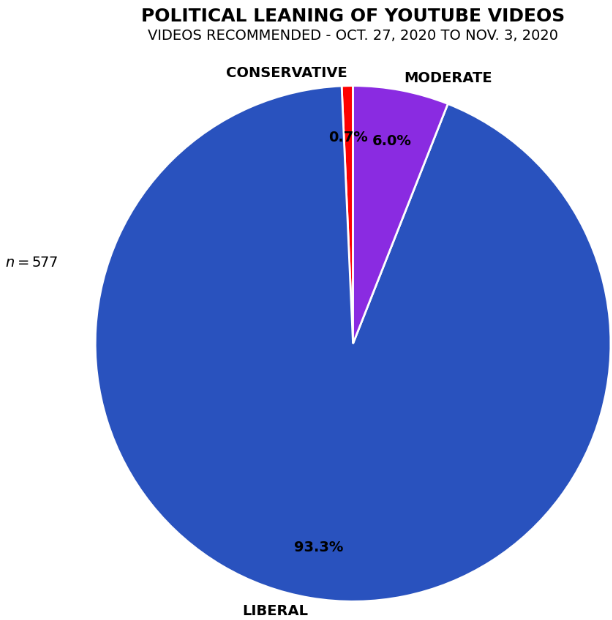

Figure 14. **2020 Presidential election: Political leaning of videos watched on YouTube, October 27, 2020 to November 3, 2020**. All videos were recommended by Google's up-next algorithm. Bias is shown only for videos coming from news sources. To compare this distribution to the distribution of actual news sources online, view Figure 16.

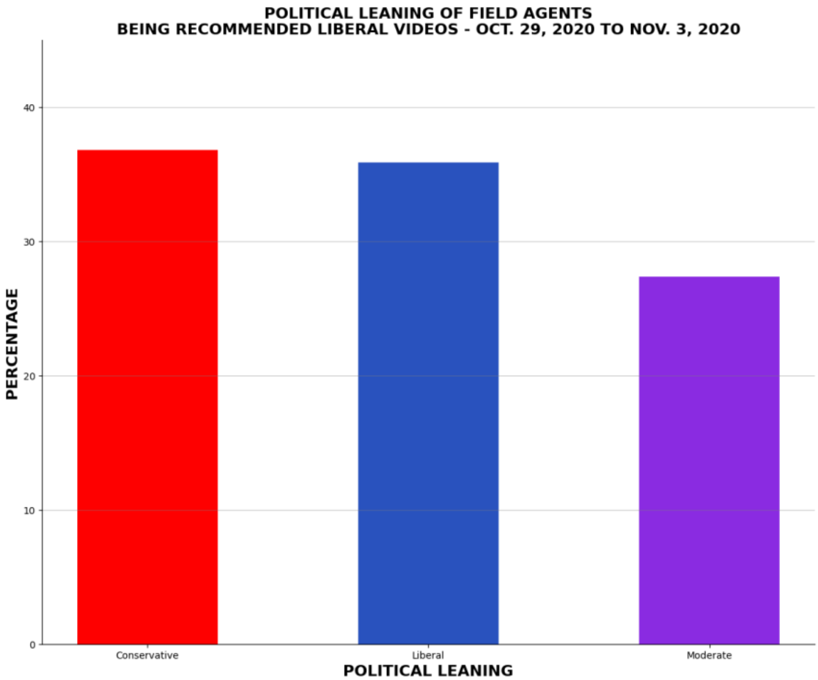

Figure 15. 2020 Presidential election: Proportion of liberal-leaning videos being recommended to our conservative, liberal, and moderal field agents in the days leading up to Election Day in 2020.

The Georgia project proved to be especially informative for us – perhaps a bellweather for what large-scale monitoring systems might be able to accomplish. On October 30, 2020, I sent a summary of our findings up to that time in the Presidential election to the office of Senator Ted Cruz (R, Texas). I had had some contact with him in the summer 2019 when I testified before a subcommittee of the US Senate Judiciary (Epstein, 2019c), so I sent the summary to my contact in his office. As a result, on November 5, 2020, Senator Cruz sent a letter co-signed by Senator Ron Johnson (R, Wisconsin) and Senator Mike Lee (R, Utah) to Sundar Pichai, the CEO of Google. It said, among other things, that the preliminary findings from my 2020 monitoring system suggested that Pichai's own testimony before Congress that "We won't do any work, you know, to politically tilt anything one way or the other" was not true.

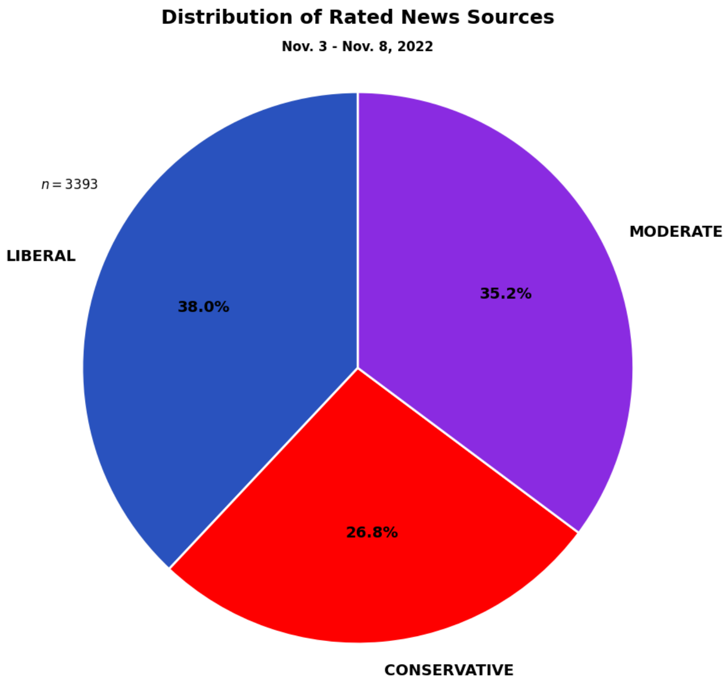

Figure 16. Distribution of rated news sources as of November, 2022. These news sources are rated by three nonpartisan organizations: Ad Fontes Media, All Sides Technology, and Media Bias/Fact Check. Their ratings suggest that a fair search engine would show people conservative, moderate, and liberal news stories (whether in text, video, or other formats) in roughly equal proportions, or, possibly, in the proportions matching the political leanings of their users.

That night, and on the days that followed leading up to the January 5, 2021 runoff elections in Georgia, my staff and I noticed two abrupt changes in the data we were collecting in Georgia:

First, beginning that night, none of our field agents in Georgia received go-vote reminders on Google's home page (Figure 17); if that meant that Google had stopped sending such reminders to voters statewide, it was presumably giving up one of its most powerful potential vote manipulations. Second, and more striking, political bias in Google search results disappeared. On Election Day (January 5, 2021) and the 5 previous days – the 6-day period we focused on in our analyses – the political bias in Google search results was close to 0 (Figures 17 & 18). This was – except for the graph I mentioned earlier that had come from a small group of Gmail users in 2016 – the first time we have ever seen unbiased Google search results.

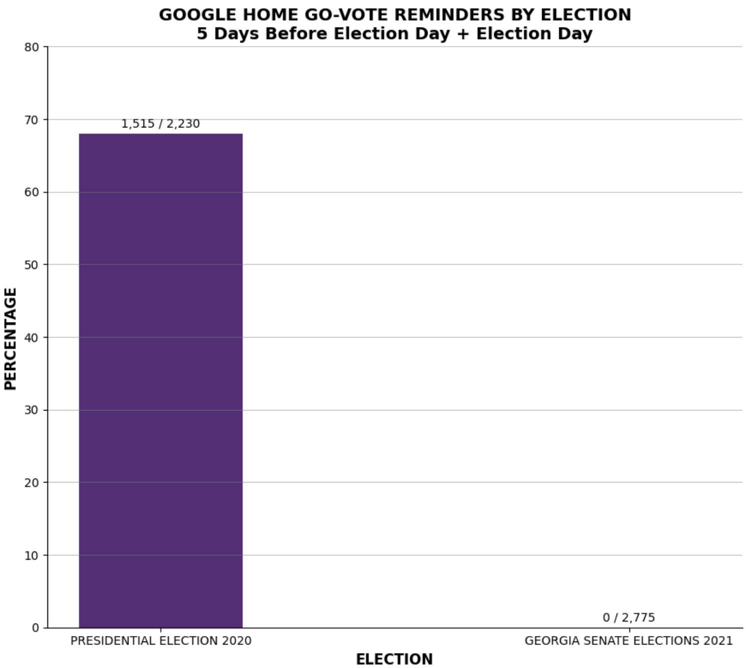

Figure 17. Google home page go-vote reminders, 2020 Presidential election vs. 2021 Georgia Senate runoff elections. Each bar show the proportion of home pages on which go-vote reminders were shown during the 5 days leading up to each election (plus Election Day). The proportion was over 70% in the Presidential election; we detected no go-vote reminders in the Georgia elections (z = 50.78, p < .001).

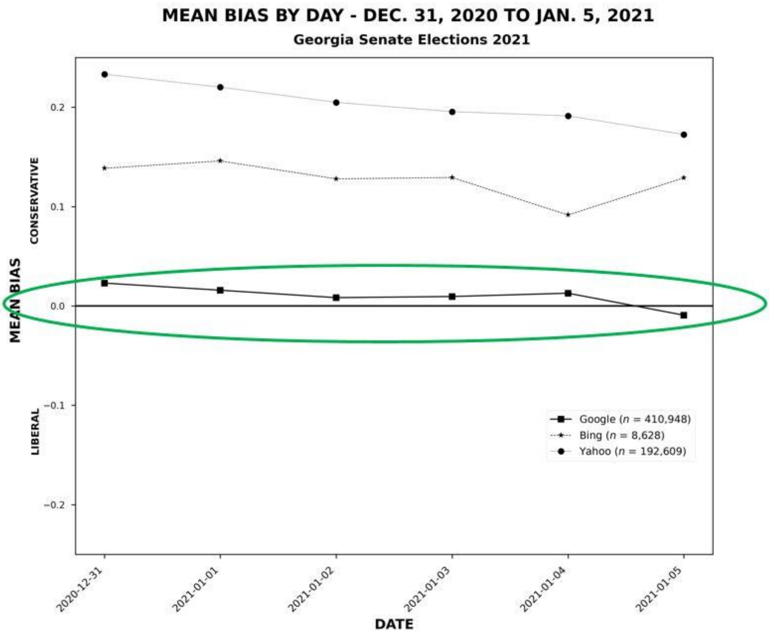

Figure 18. 2021 Georgia Senate runoff elections, political bias in search engines, December 31, 2020, to January 5, 2021. On Election Day in Georgia, as well as on the 5 days leading up to the election, we found virtually no political bias in Google search results but some degree of conservative bias on the Bing and Yahoo search engines. (To see these data in a bar graph, view Figure S#)

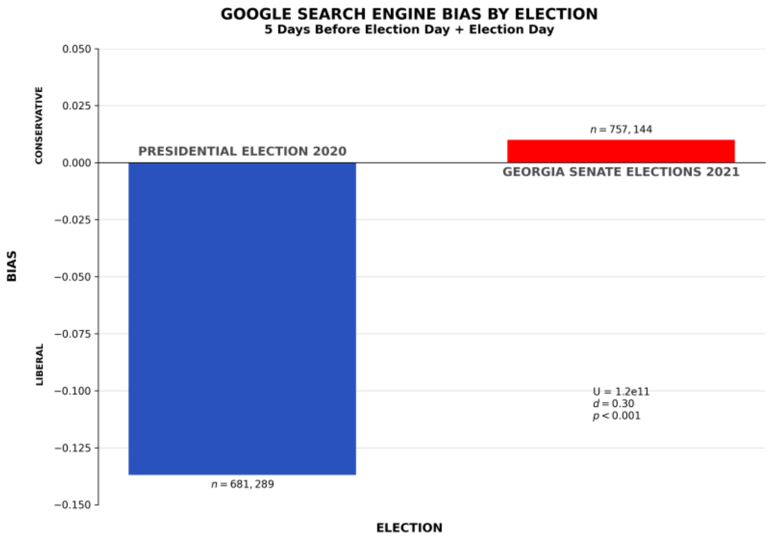

Figure 19. Google search engine bias, 2020 Presidential election vs. 2021 Georgia Senate runoff elections. Each bar represents mean bias on the first page of Google search results during the 5 days leading up to each election, plus Election Day. We found substantial liberal bias in the Presidential election but virtually no political bias in the Georgia elections. The bias in Georgia search results disappeared shortly after three US Senators sent a letter to the CEO of Google protesting possible election bias in Google search results. See text for details.

4. *2022 US Midterm Election*

In 2022, we built a bigger monitoring system, this time with 2,742 field agents located mainly in 10 swing states: Arizona, Florida, Georgia, Missouri, North Carolina, New Hampshire, Nevada, Ohio, Pennsylvania, and Wisconsin. We preserved 2,549,544 ephemeral experiences on Google, Bing, Google home page, YouTube, Twitter, and Facebook, and we measured political bias on contact both by using ratings from a custom machine learning algorithm we developed (which had achieved 85% agreement with human raters who had helped train the algorithm) and by using ratings from those three nonpartisan organizations that use various methods to rate the political bias of news sources:

(1) AllSides Technologies, Inc. (https://allsides.com), a public benefit corporation that rates more than 1,400 media outlets and writers. It relies on editorial reviews conducted by multipartisan panels of six-to-nine reviewers from the left, center, and right. Its motto is, "Free people from filter bubbles so they can better understand the world — and each other."

(2) Ad Fontes Media (https://adfontesmedia.com), a public benefit corporation the mission of which is "to rate all the news to positively transform society." It relies on 60 human analysts to examine news sources, subjecting each to perusal by one liberal, one moderate, and one conservative analyst. At this writing, it has rated more than 3,900 news sources.

(3) Media Bias Fact Check (https://mediabiasfactcheck.com), an independent, donation supported organization that relies on "a collective of volunteers and paid contractors" to rate the "ideological leanings and factual accuracy" of online content. It claims to have evaluated more than 8,200 "media sources, journalists, and politicians" so far. Its mission is "to educate the public on media bias and deceptive news practices."

The main way we rated news sources in our 2022 monitoring project was to quickly rescale ratings from these three sources (or just two of them, or even one in rare cases in which one or more of these rating services did not list the news source we were checking) from -1.0 (for liberal) to +1.0 (for conservative) and then to calculate the mean of the ratings.

Our field agents were fairly closely balanced politically (Figure 20), and our search terms were again rated as neutral by independent raters – between -0.2 and +0.2 on a political bias scale ranging from -5.0 to +5.0. The data we gathered from the computers of our field agents were preserved from searches they conducted using more than 500 neutral search terms.

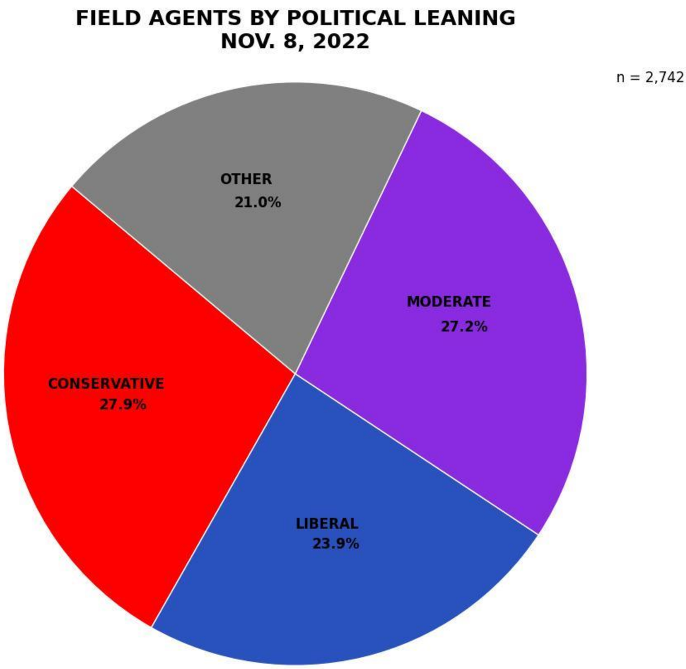

Figure 20. **Pie chart of field agents by political leanings during the 2022 presidential election**.

Our major findings were as follows:

(1) Focusing on the 5 days immediately preceeding Election Day (November 8, 2022), plus Election Day, we found, as we had in previous elections, that search results on Google were strongly liberally biased; search results on Bing were closer to zero (Figures 21 and 22). If seen nationally and shown to people for 6 months prior to the midterms, political bias in Google search could have shifted more than 80 million votes (spread across hundreds of elections nationwide).

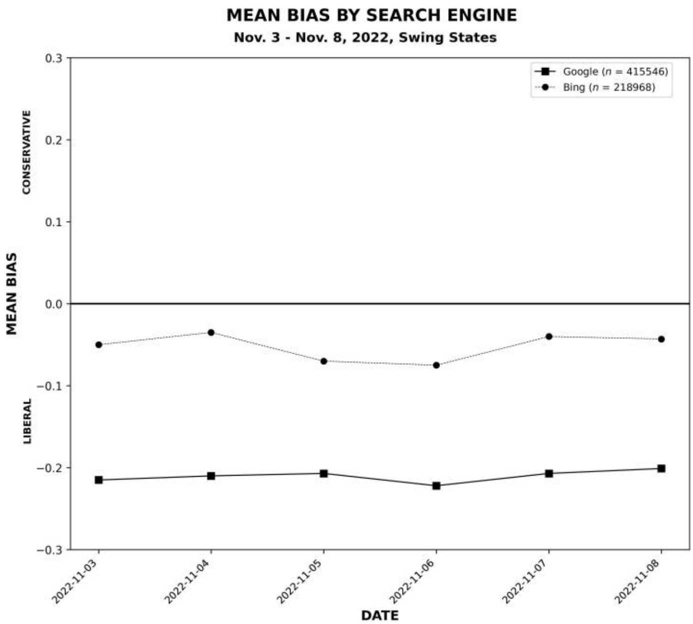

Figure 21. 2022 midterm elections, mean bias by search engine, days leading up to the elections.

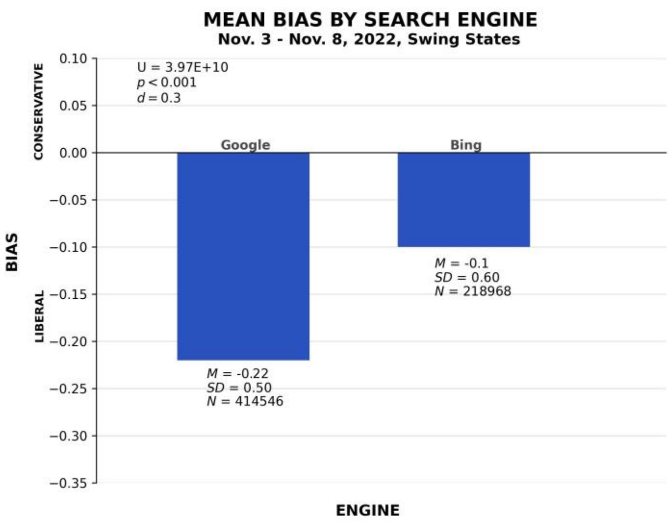

Figure 22. 2022 midterm elections, mean bias by search engine.

2) Once again, we found that Google was showing liberally biased search results in roughly equal proportions to liberals, moderates, and conservatives (Figure 23).

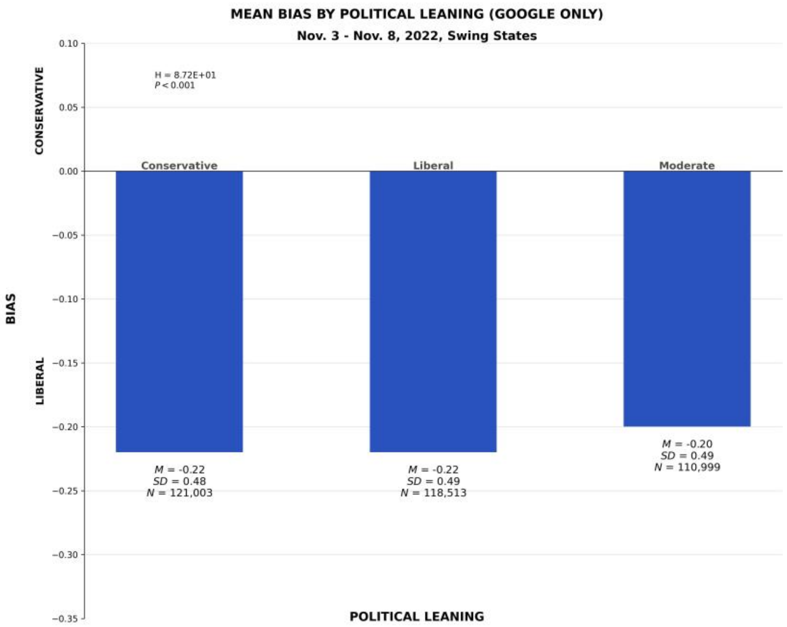

Figure 23. 2022 midterm elections, mean bias by the political leaning of our field agents.

(3) Again, as we have seen before, Google showed users liberally biased search results in all 10 positions on the first page of their search results (Figure 24).

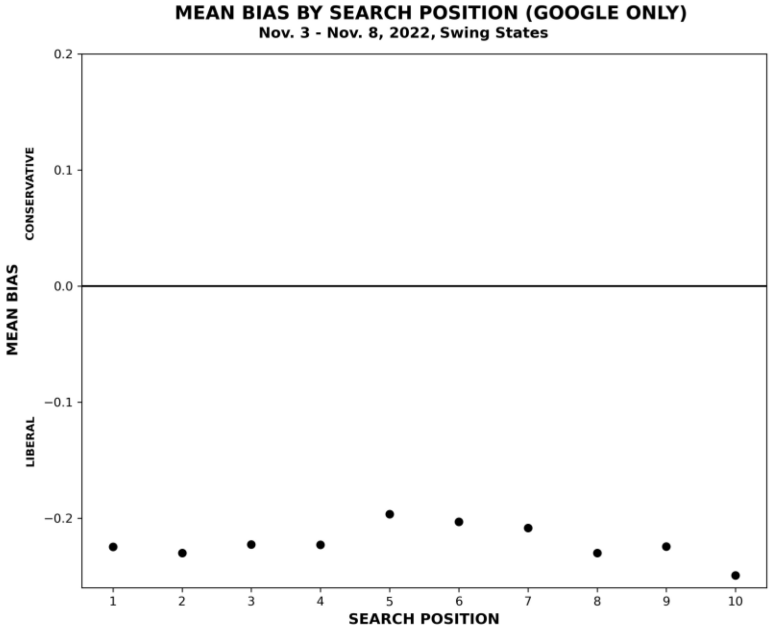

Figure 24. 2022 midterm elections, mean bias by search position on Google.

(4) In 2022, we also began tracking "political update" tweets sent by the Twitter company itself to our field agents, and we found that in the days leading up to the elections, Twitter sent significantly more of such updates to our liberal field agents than to our conservative field agents (Figure 25). Note that these elections took place before Elon Musk took control of the company and fired 80% of its employees, complaining that Twitter content was blatantly liberally biased (Nava, 2023).

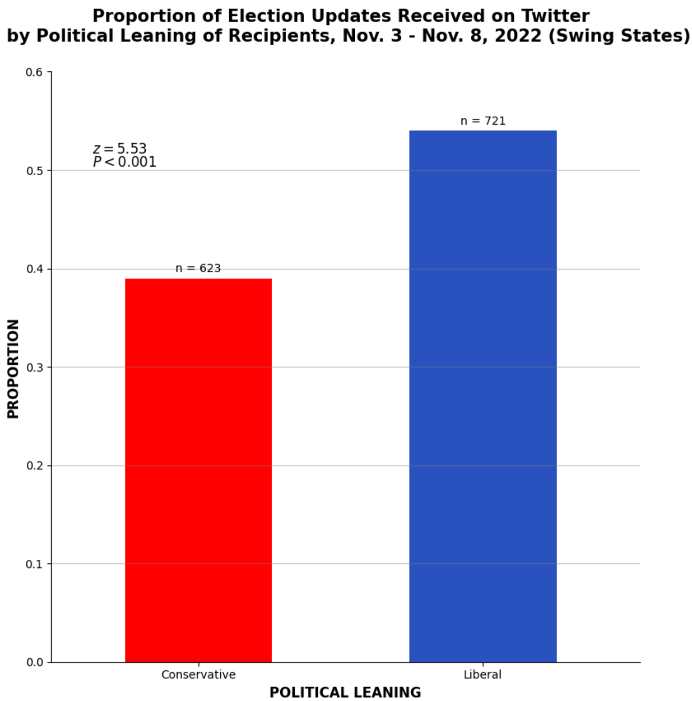

Figure 25. 2022 midterm elections, proportion of election updates received by our conservative and liberal field agents in the days leading up to the elections.

(5) In 2022, our system became more proficient in preserving the content on Google's home page. In Florida in the days leading up to the election, we found a large and statistically significant difference in the number of go-vote reminders Google sent to our liberal and conservative field agents. All liberal field agents received them, but only 59% of conservatives did. This is the kind of blatant targeting that can shift large numbers of votes in a national election – even on Election Day itself (Bond et al. 2012; Epstein et al., 2023; Zittrain, 2014). Nationwide – with data mainly from 10 swing states, as I noted above – we again found that 100% of our liberal field agents received go-vote reminders on Election Day on Google's home page and that the proportion of conservatives who received such reminders was significantly smaller ($p < .001$) (Figure 26).

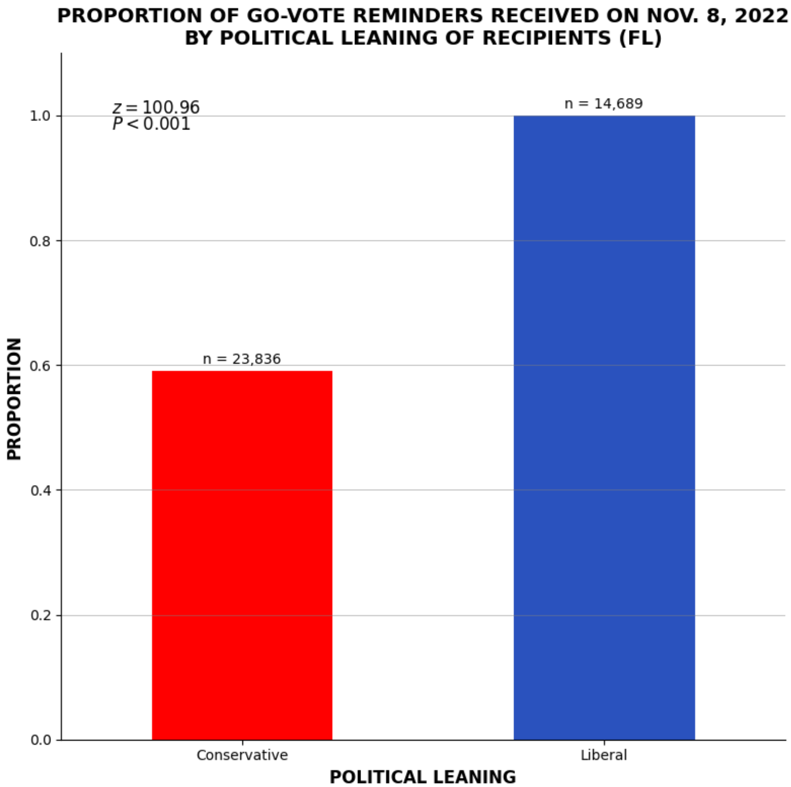

Figure 26. 2022 midterm elections, proportion of go-vote reminders received in Florida on Election Day by our conservative and liberal field agents.

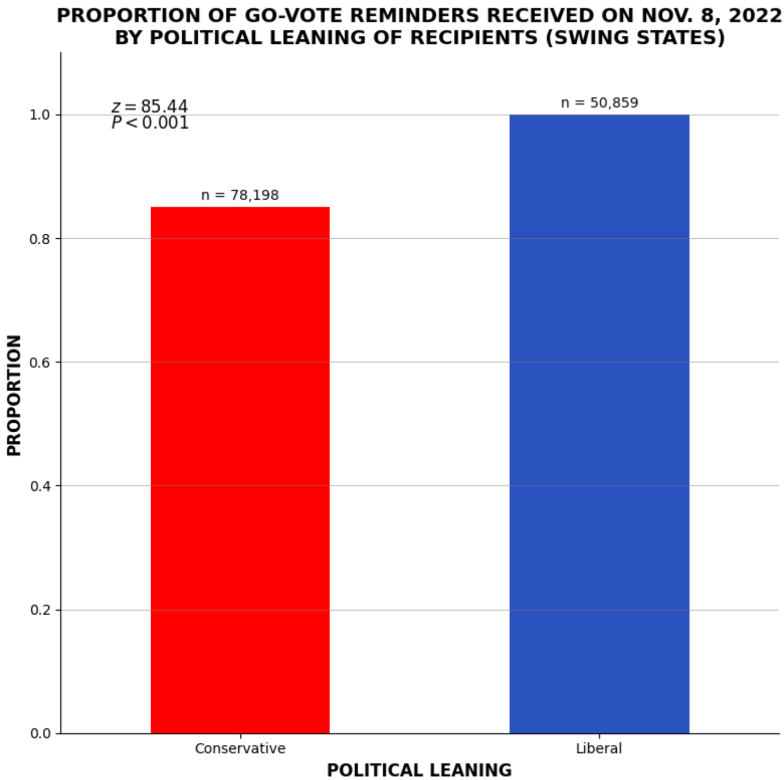

Figure 27. 2022 midterm elections, proportion of go-vote reminders received nationwide (mainly in 10 swing states) on Election Day by our conservative and liberal field agents.

(6) Finally, we found a liberal bias in the recommendations Google made to our field agents on YouTube in the days leading up to the election (Figure 28). 76.1% of those recommendations came from liberal news sources, about twice as many as one would expect by chance (for comparision, see Figure 16, which shows a breakdown of all US video news sources).

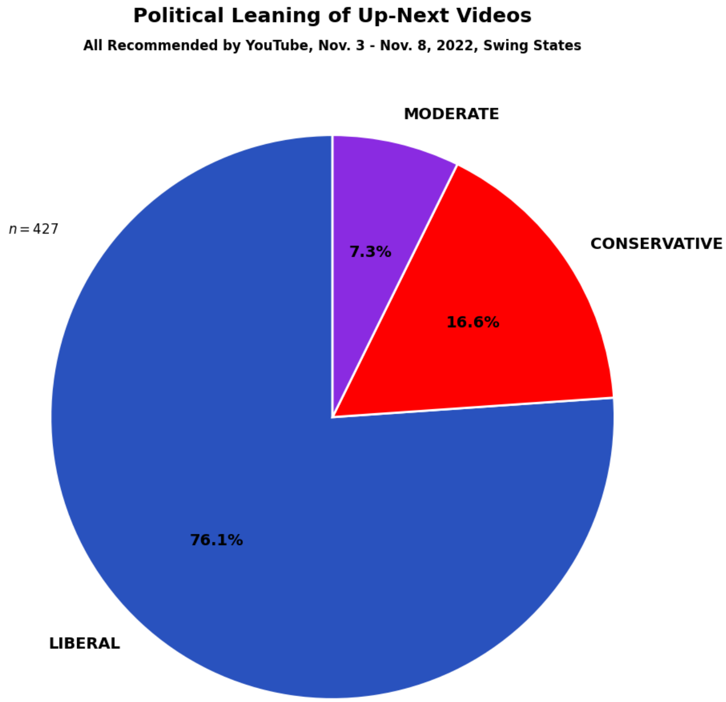

Figure 28. 2022 midterm elections, percentage of up-next videos on YouTube, shown by political leaning.

5. *2023-2024 Nationwide Monitoring System*

After the 2022 midterm elections, I decided that the time had come to build a permanent, nationwide monitoring system. I was thinking along these lines: We know from a decade of controlled studies that ephemeral content can be used to shift people's thinking and behavior dramatically – typically without their knowledge and also without leaving a paper trail for authorities to trace. We also know from leaks and whistleblowers that most Silicon Valley technology companies share strong social values and are not shy about using the various powers they have to advance their agendas – monetary, philosophical, and political.

We also know that our leaders – at least the ones we're able to keep an eye on in Washington, DC – seem to be hobbled by partisan bickering most of the time. They hold hearings occasionally in which they browbeat Big Tech executives, but they seem incapable of protecting the

American public from the three big threats that Big Tech companies pose to our citizenry and our society: the massive surveillance, the censorship, and the manipulations.

My team and I did not have solutions to these problems, but we had developed a way of making Big Tech companies accountable to the public – specifically, by tracking, capturing, and analyzing the actual personalized content these companies were sending to children, teens, and adults 24 hours a day.

As I mulled over these issues, I finally concluded that a permanent, nonpartisan, nationwide monitoring system – or perhaps several such systems – had to be established in the US – and perhaps elsewhere around the world – for two reasons. First, if no such system exists, *we will have no idea how tech companies, now and in the future, might be using or misusing ephemeral content*. Manipulations that could substantially shift the thinking and behavior of most people in the world would be entirely invisible both to users and to authorities.

And second, if, by some miracle, world leaders enact laws and regulations that protect users from possible manipulations by technology companies – as EU leaders have been doing aggressively since they passed the General Data Protection Regulation (GDPR) in 2016 – *they will have no way of measuring compliance with these new laws and regulations unless large-scale monitoring systems are in place*. The European Commission – the EU agency that has been investigating Big Tech companies since 2014 – recently admitted that Big Tech companies have "fall[en] short of effective compliance" with the restrictions the EU has put in place in recent years (CBS News, 2024; cf. Bergen, 2015; Hirst, 2016; Rankins, 2016; Sterling, 2019).

For credibility, I speculated that our new system had to recruit a politically-balanced group of registered voters in each US state – our "field agents" (FAs) – and that the number of FAs in each state had to exceed some minimum so that data collected from the computers of those individuals might be court admissible in state courts. To achieve this, I did some rough calculations regarding the minimum proportion of representative registered voters we needed in each state to allow us to make statistically significant predictions about the entire population of registered voters in that state. I cannot, alas, make those numbers public – at least not at this time in this essay – just as the Neilsen Company cannot release details about the families who help it compute the Neilsen ratings. To do so would be to give the companies we monitor guideposts for identifying our FAs.

My team was now proficient in recruiting field agents in a way that protected the integrity of our data. This meant, first and foremost, that we could not ask for volunteers. If we did, we would run the risk that Google would send us thousands – not from among their employees, perhaps (because we could probably identify that connection), but from among their more than 120,000 "temporary, vendor and contract workers" (Moreno, 2019) – people whom we would have trouble associating with the company. That issue aside, we have found ways of approaching registered voters directly, of vetting them, having them sign non-disclosure agreements, equipping their computers and mobile devices with "passive" monitoring software that Google could not easily detect, training them, and protecting their identities. Although this was an expensive and labor-intensive process, by late 2022, we knew how to do it efficiently.

We had also learned over the years how to monitor an increasingly wider and more diverse range of online ephemeral content – even those search suggestions Google flashes at you while you are typing a search term. We had learned how to capture search suggestions by the millions.

We had also made progress in protecting our data, in recovering rapidly from attacks (yes, there have been many), and in analyzing the wealth of data we were preserving.

I now believe that monitoring systems are essential for protecting democracies around the world. As a parent, I also have become increasingly concerned about the many ways in which online content is adversely impacting our children (Common Sense, 2022; Izundu, 2015; Ybarra, 2016). Over the years, it had also crossed my mind that the various new types of manipulation I had been discovering and quantifying would probably work even better on children than on adults. Might Google or other companies deliberately use such techniques to alter the thinking and behavior of children around the world – to "indoctrinate" chidren, as some might say? Even if that possibility was highly unlikely, it seemed to me that a permanent monitoring system had to be able to capture ephemeral content being sent to the mobile devices of children (with parental permission, needless to say).

Shortly after the 2022 midterms, I began seeking support that would allow us to expand the system we had already built, and by late May, 2023, we were able to build a system that included just over 8,000 field agents in 49 US states. By May 31st, we had also preserved 16,520,793 ephemeral experiences on multiple plat- forms.

Toward the end of November, we built and deployed what for me was a dream come true: a real-time dashboard for what we now called "America's Digital Shield" (ADS; you can view the dashboard at https://AmericasDigitalShield.com), and I began to talk about it in lectures and interviews. The dashboard showed, in real time, the number of ephemeral experiences we had preserved, disturbing images we had preserved from videos YouTube was recommending to children and teens, political bias in search results on Google, Bing, and Yahoo, and political bias in recommended videos on YouTube (Figure 29).

As a result, I was soon invited to testify again before Congress, this time on December 13, 2023, before the United States Senate Judiciary Subcommittee on Competition Policy, Antitrust, and Consumer Rights, chaired by Amy Klobuchar (D, MN). My written testimony, which included peer-reviewed publications and manuscripts still under review by scientific journals, was 480 pages in length (Epstein, 2023).

As of this writing (September 4, 2024), we have preserved more than 99 million ephemeral experiences – data that are normally lost forever – and we have preliminary evidence suggesting that our monitoring system is having a mitigating effect on political bias in Google search results. Specifically, since we went public with ADS in November, 2023, we have seen a slow and steady decrease in liberal bias in their search results (Figure 30). This decrease could also mean that Google is rigging the game: reducing overall political bias while focusing its manipulations narrowly on the races it is most concerned about.

And, of course, this decrease could also mean that Google is gradually identifying our field agents, and, with every identification, is sending out sanitized content, as we saw with Gmail users in 2016. Over the years, however, we have developed increasingly sophisticated ways of determining whether any of our field agents are receiving suspicious data. When we are in doubt, we remove those people from service.

For this tracking system to produce valid numbers, secrecy is essential. With billions of dollars in production costs and advertising revenues on the line, imagine the lengths to which interested parties might go to influence a Nielsen family's viewing habits – or, for that matter, to tamper with those black boxes. Over time, Neilsen has developed increasingly sophisticated ways of detecting when incoming data was suspect, and so have we.

Here is a sampling of findings from the massive and ever-growing database of ephemeral content we have been building over the past year or so:

https://AmericasDigitalShield.com

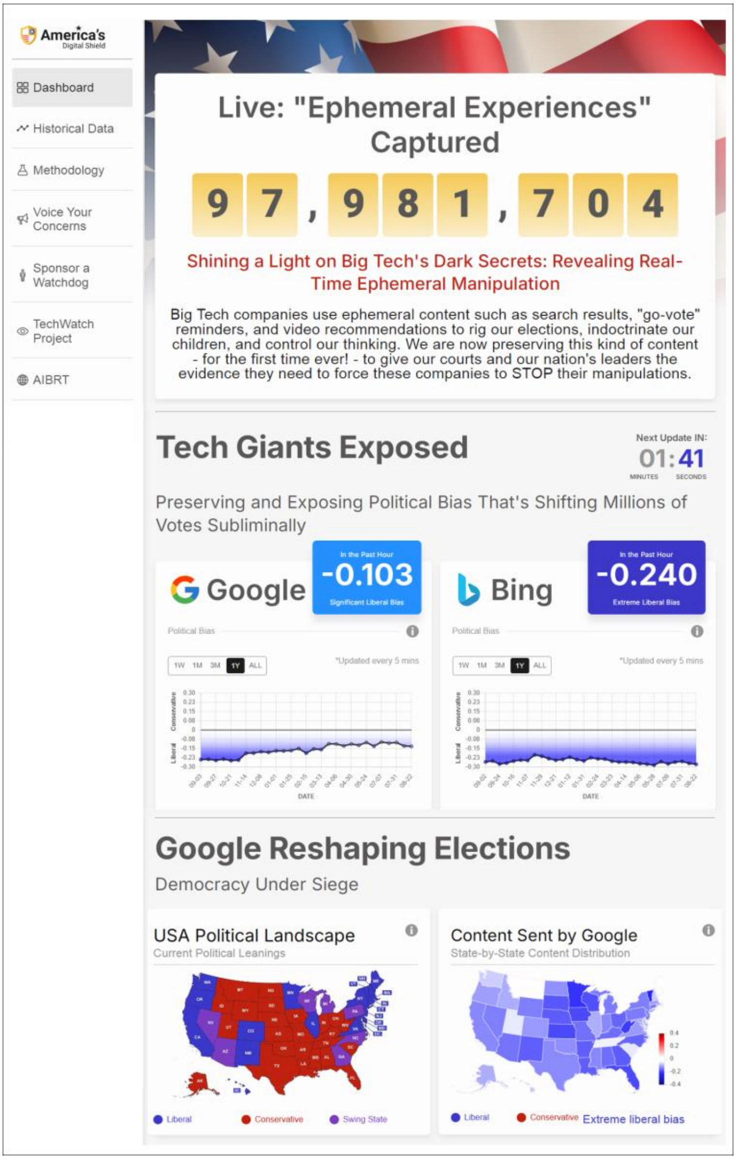

Figure 29. America's Digital Shield. The graphs above show statistically significant liberal bias in both Google and Bing search results (the blue shaded areas below each x-axis) as of May 13, 2024. The image in the lower right shows that Google is sending liberally biased content to registered voters in all 50 states – to liberals, moderates, and conservatives. As of September 3, 2024, the total number of ephemeral events we have preserved has passed 99 million.

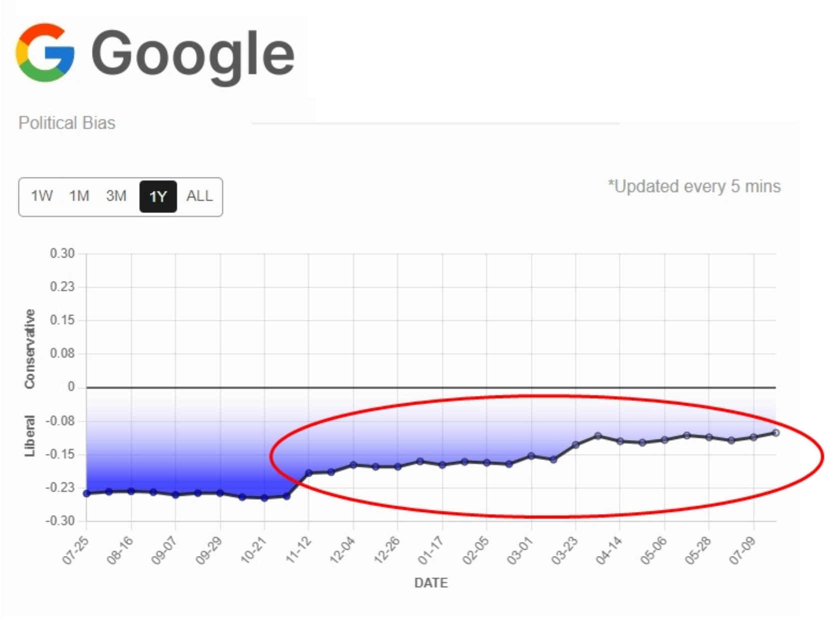

Figure 30. Data from the public dashboard of America's Digital Shield (ADS). This graph, which is updated every 5 min online at https://AmericasDigitalShield.com, suggests a slow and steady decline in liberal bias in Google search results beginning in mid-November, 2023, which is when we launched and went public with ADS. This could mean that ADS is having a mitigating effect on political bias in Google search results, although other explanations are possible.

(1) *Search engine bias*. Above, I have already shown you a graph from the ADS real-time dashboard suggesting that (a) the liberal bias in Google search results since we began monitoring in 2016 is still going strong, and (b) that Google might have begun to gradually reduce this bias since we went public with ADS in November, 2023. We are seeing the opposite trend on Bing and Yahoo, however (Figs 31 and 32). Because Bing and Yahoo combined attract only about 3% of search traffic in the US, bias in their search results presents little threat to the free-and-fair election here. That said, why, over the years, Bing and Yahoo have each drifted away from showing a small conservative bias to an increasingly larger liberal bias is a matter that should be investigated. Relevant here is the fact that Yahoo has not been a true search engine (that is, a company that aggressively crawls the

internet to update its index of internet content) since roughly 2002 (Kratz, 2009). Since 2015, Yahoo has been drawing most or perhaps all of its search results from Google (Fingas, 2015). Less clear is whether a secret “pact” that was signed between Google and Microsoft (owner of Bing) in early 2016 gave Bing the right to begin drawing search information from Google. At the moment, I regard the shift in political bias in Bing and Yahoo search results to be a mystery.

(2) *Targeting individuals*. Here are three examples of how our data can be used to investigate how content on Google or other platforms might be putting certain individuals in either a positive or negative light. In the first, we see that users seeking information about Ken Paxton – the conservative Republican attorney general of Texas who has sued Google more than once – on the Google search engine are receiving search results that are highly liberally biased (Figure 33) – in other words, probably hostile toward Paxton. In the second, we see that users seeking information about liberal Senator Elizabeth Warren – the rare Democrat who has called repeatedly for Google’s breakup (Warren, 2019; cf. Epstein, 2019b) – are receiving highly *conservative* search results (Figure 34). In each case, the same level of bias is being sent to our liberal, moderate, and conservative field agents. In both cases, it would appear that politically biased search results are being used to vilify these individuals. Finally, in our third example, we show a steady increase in liberal bias between March and July 2024 when people searched Google for information about Vice President Kamala Harris (Figure 35). So our system can not only detect whether Google or other search engines favor certain candidates (or causes, brands, historical figures, or religions), it can also show how such bias changes over time.

(3) *Targeted messages*. I mentioned earlier that on Election Day in Florida in 2022, Google appeared to be sending substantially more go-vote reminders to liberals than to conservatives. At this writing (August 26, 2024) – and this observation is subject to change – it appears that our liberal FAs nationwide are getting register-to-vote reminders on Google’s home page at two-and-a-half times the rate that our conservative FAs are getting such

reminders. It is not inconceivable that those messages could turn into partisan mail-in-your-ballot reminders, which could then become partisan go-vote reminders. Recall that a simple extrapolation from the study that Facebook published in *Nature* in 2012 with faculty members at the University of California San Diego suggests that a partisan go-vote reminder on Election Day could give the favored candidate at least 450,000 additional votes that day.

(4) *State-by-state differences.* Currently on the ADS dashboard – again, this is subject to change – we are showing two maps of the US. The first shows the usual red and blue states, with just seven "swing" states (shown in purple) in which the winners determine who wins the 2024 Presidential election (Figure 29). At this writing (August 17, 2024), it appears that Kamala Harris' entry into the fray has turned one of those swing states blue. We are also displaying a map that uses a range of color shades from dark red (signifying strong conservative bias) to dark blue (signifying strong liberal bias) to show, in real time, the mean bias in the search results Google is showing to our FAs state-by-state. As of this writing (August 17, 2024), 49 of the 50 states are various shades of blue. The only red state is Alaska, in which the conservative bias is a paltry .03 on a scale from -1.00 to +1.00. Over the past year that we have been tracking state-by-state bias in Google search, we have only found a few days when any states turned red (never more that five), and we have *never* seen one of the swing states turn red.

Figure 31. Political bias in Bing search results.

Figure 32. Political Bias in Yahoo search results.

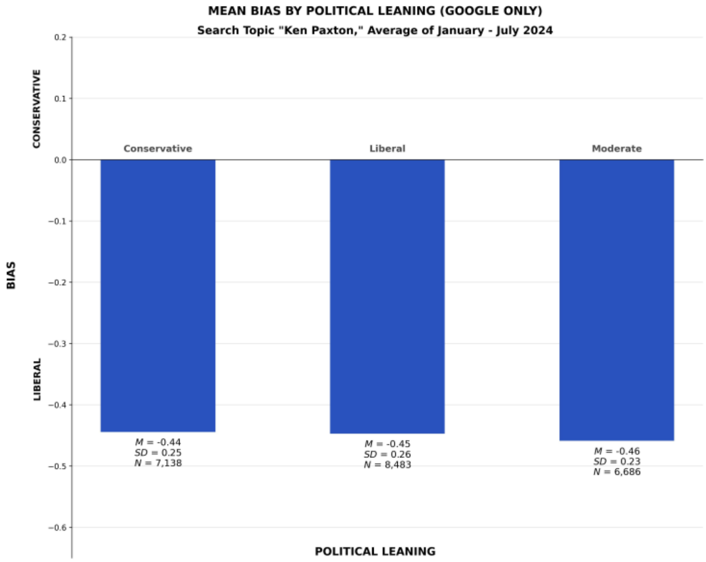

Figure 33. Political bias in Google searches for Attorney General Ken Paxton of Texas.

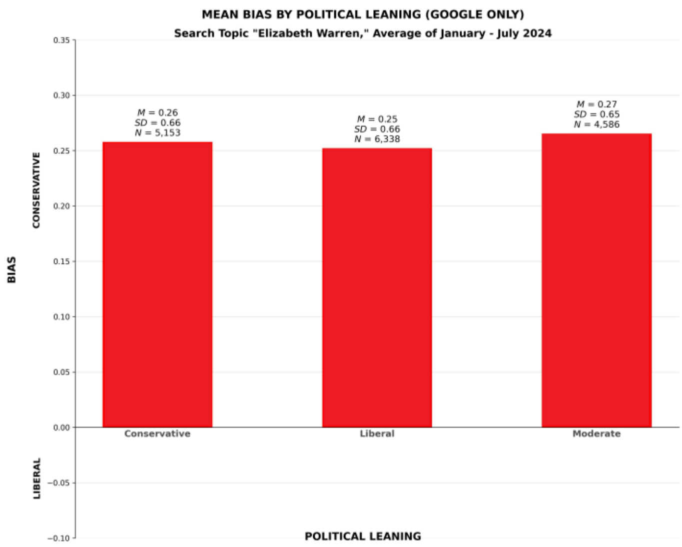

Figure 34. Political bias in Google searches for Senator Elizabeth Warren (D, MA).

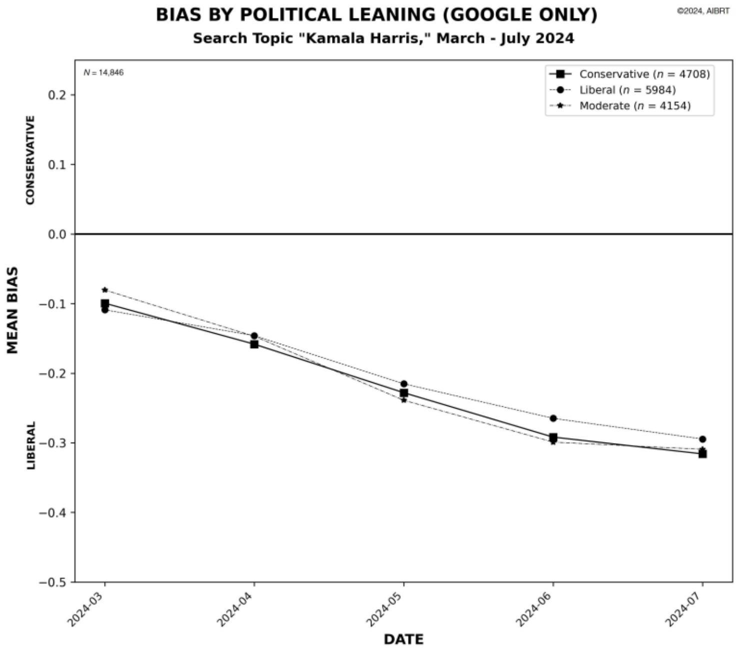

Figure 35. Political bias is Google searches for Vice President Kamala Harris.

(5) *Political bias on YouTube*. We are also capturing hundreds of thousands of images and videos from YouTube, guided entirely by recommendations Google's up-next algorithm makes. One way we have analyzed this content so far is to look at the political bias in news sources of the recommended videos. Once again, we have found liberal bias fairly consistently (Figure 36).

(6) *Disturbing content going to children and teens*. Currently on ADS, we are also showing some of the gruesomely violent and sexually explicit content we have preserved on videos that are being recommended to children and teens on YouTube. We offer a few examples in the Supplementary Materials (S9 to S15 Figures).

By the way, if my graphs don't impress you, please bear in mind what these data say about the *actual content* people are seeing. When someone clicks on a high-ranking search results shown in searches about Republican Attorney General Ken Paxton, he or she is likely to be brought to a web page containing a headline such as, “Reports of Texas AG’s Attempts to Collect Data on Trans Adults Are a Terrifying Overstep Into the Lives and Privacy of LGTBQ+ People” (Figure 37). And when someone clicks on a high-ranking search result shown in searches about Democratic Senator Elizabeth Warren, he or she is likely to be brought to a web page containing a headline such as, “The Socialist Moment Hasn’t Passed. It’s Yet to Come” (Figure 38). This is the type of ephemeral content we are preserving; we make our graphs by aggregating thousands of examples like these.

We will soon be adding to the dashboard figures that summarize some of the content we have been collecting on other platforms, among them TikTok, Facebook, Instagram, and X.

The findings I have summarized above are not all the interesting things that ADS can show us. These are simply some of the main findings we are currently displaying on our public dashboard, and these findings are updated every 5 min, which means they are likely to change over time.

Rather than show you other findings, below, in the final section of this essay, I will offer you a vision – an idea of the possibilities. We have thus far preserved terabytes of data, including tens of millions of search suggestions, millions of YouTube recommendations (which are ephemeral, of course), plus millions of images and HTML files from TikTok, Instagram, Facebook, Twitter, and other platforms. What other information might we preserve?

An old professor of mine used to say that “repetition is the mother of wisdom,” so I am going to repeat, yet again, that all of these data are ephemeral – data that impact people and then disappear – normally forever. What are these data telling us at the moment, and what might they tell us moving forward? And where we have detected what appear to be signs of political bias in the content we have preserved, should we be concerned? If so, what actions might we take to mitigate such bias?

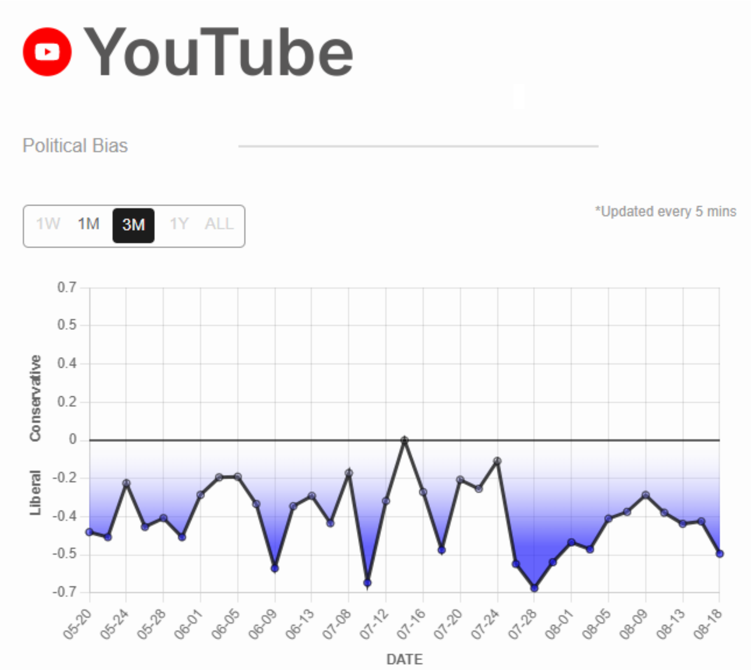

Figure 36. Political bias in Youtube's recommended video content.

‹ PRESS RELEASES

Human Rights Campaign: Reports of Texas AG's Attempts to Collect Data on Trans Adults Are a Terrifying Overstep into the Lives and Privacy of LGBTQ+ People

by HRC Staff • December 14, 2022

Figure 37. Webpage vilifying AG Ken Paxton from high ranking search results on Google.

Scott W. Stern / February 5, 2024

REBEL REBEL

The Socialist Moment Hasn't Passed. It's Yet to Come.

The most promising leaders on the left today are not elected politicians.

Figure 38. Wepage vilifying Senator Elizabeth Warren from high ranking search results on Google.

Summary and Conclusions

The internet has made it possible for a small number of technology companies to dominate the thinking, behavior, and votes of more than 5 billion people worldwide using new subliminal techniques. We have discovered and quantified 10 of these techniques in controlled experiments we have been conducting and publishing since 2013, and, in 2016, we developed technology that allowed us to preserve search results on multiple search engines. Search results, like newsfeeds and video sequences, are types of ephemeral content that influence thinking and behavior and then disappear, leaving no paper trail for authorities to trace. In 2018, 2020, and 2022, we improved and expanded our monitoring system to preserve a wide variety of online content in the days leading up to elections held in the US in those years.

We build our systems by recruiting real voters around the U.S. – in 2016, just 95 voters in 24 states – and, with their permission, installing custom software on their computers that allows us to stream the political content they see to our servers, where we quickly aggregate and analyze the data. In our small 2016 project, we preserved 13,207 politically-related searches, along with the 98,044 web pages to which the search results linked. We found substantial political bias on the most popular search engine (Google), sufficient to have shifted at least 2.6 million votes in the Presidential election that year. In 2022, through the computers of a politically-balanced group of 2,742 registered voters, we preserved more than 2.5 million ephemeral experiences on multiple platforms, which tended, once again, to be highly biased politically. In late 2022, we began to build a permanent nationwide monitoring system – our "America's Digital Shield" project. As of this writing (August 17, 2024), we have preserved more than 97 million ephemeral experiences on multiple platforms through the computers of a politically-balanced group of more than 15,000 registered voters in all 50 states, with the system growing larger each day. This system has the potential to make Big Tech companies accountable to the public for the first time and for the foreseeable future, forcing them to constrain their algorithms so that they do not interfere with our free-and-fair elections, the impressionable minds of our children, and our own autonomy.

Imagine the possibilities. What might one do with the wealth of ephemeral content we have been preserving? The tantalizing truth is that *I don't know*. I can only tell you two things I know for sure: In our national system, we have now preserved more than a year's worth of data on multiple platforms, with the database growing every day. So one thing we can do is to *look back in time* at ephemeral content – something that has never been possible before. There is practically no limit to the kinds of patterns and trends one might search for.

I find the second possibility even more intriguing. We frequently adjust and expand the parameters we use to track and preserve data, looking at more kinds of data on more platforms. Again, there is potentially no limit to the kind of data we can preserve. We also have built devices that will allow us to preserve the spoken answers home surveillance devices such as Alexa and Google Home give to users. Over time, our expanding pool of FAs will be equipped with these devices, and we will be parsing these answers, looking for content that might impact people's thinking and behavior in any number of ways.

We are also beginning to collect more and more of the content being generated by AIs such as ChatGPT. In a few short years, search engines will likely be used only by scholars and scientists. The vast majority of humankind will simply ask questions of their devices, and they will likely believe the answers they read or hear (see the section in Part Two of this essay on the Answer Bot Effect [ABE]). Unfortunately, nearly all the content being generated at the moment by answerbots is – that's right – *ephemeral*. Without sophisticated monitoring systems in place to capture such content, *no one will know what that content is or how it might be affecting elections, our children, and you and me.*

What's more, our ability to preserve and analyze content from generative AIs means that we can, going forward, track *the extent to which AIs present serious threats to humanity*. Real-time monitoring systems can provide active threat assessment; they might prove to be humanity's best defense against the threats AI pose.

The good news here is that *monitoring is tech*, and it can, with adequate resources, move as fast as the tech industry does. Laws and regulations move at Turtle Speed; whereas technology moves at, well, Ludicrous Speed (to borrow a phrase from the classic film, "Spaceballs" – a term that was also borrowed by Tesla to describe its fastest acceleration rate). I don't think we should give up on our legislators and regulators, but given the accelerating rate at which technology is changing (speaking of acceleration), I think we need to keep our expectations low. Ideally, data from monitoring systems might help our government officials to craft more meaningful and effective laws and regulations to keep tech in check.

That question about censorship. This brings me to the question I raised at the very beginning of this essay. Is Google suppressing conservative content? I summarized three studies that concluded that Google was unbiased in this regard (Economist, 2019; Lewis et al., 2023; Metaxa et al., 2019), and I argued that none of those studies had been conducted with designs adequate to support their conclusions.

Here are some of the problems I pointed out in those studies, and here is how our monitoring systems have overcome those problems:

The bot problem. First, in both *The Economist* study and the Stanford study, content was being collected by computer programs that had clearly and unequivocally identified themselves as bots – in other words, as nonhumans. That is a fatal flaw in both studies. As *The Economist* authors admitted, “Our study does not prove Google is impartial. In theory, Google could serve un-biased links only to users without a browsing history.”

We solved this problem by recruiting real people using their own computers – people who presumably were known intimately by Google’s algorithms. One *must* collect ephemeral content through the computers of real people because (a) such content is frequently *personalized*, and you cannot see personalized content without looking over the shoulders of *real persons*, and (b) as I have shown you with our own data (see above for content we tracked during the 2016 and 2020 monitoring projects, for example), Google *turns off bias* in search results whenever it pleases. There is only one way to see the actual content Google is sending to real users, and that is to recruit a large, representative sample of people, install passive monitoring software on their computers, protect the identities of those individuals (just as the Neilsen company protects the identities of its families), and then aggregate, preserve, and analyze the data. That is what my team and I have been doing with increasing expertise and efficiency in multiple monitoring projects since 2016.

The search term problem. Second, in those studies I critiqued, the researchers used a small number of search terms: 31 in *The Economist* study and 4 in the Lewis study. In the Stanford study, only the names and states of political candidates were used as search terms, and those same names were used in searches on simulated computers every day for 6 months. In all three studies, it would have been a simple matter for Google to identify those unusual searches and to reply with atypical data. Moreover, none of these studies attempted to evaluate the political neutrality of their search terms; as we have noted, non-neutral search terms will, by necessity, produce non-neutral search results (Kulshrestha et al., 2019).

We avoid such difficulties by using long lists of trending search terms that have been rated as neutral by independent raters. In our 2016 monitoring project, we ultimately used 250 such terms, and in our subsequent projects, we have used at least 500 such terms for each election we have monitored. For security reasons, we don't reveal the exact number of terms we are using for America's Digital Shield.

The Google problem. Third, I noted above that *The Economist* was closely associated with Google at the time they published their study. The Metaxa et al. (2019) study also seemed tainted because of Stanford's close ties with Google. In contrast, AIBRT, the institute where I conduct my research, is a nonprofit, nonpartisan, 501(c)(3) public charity. More important, *we do not accept restricted gifts*. This is important, because it means that donors cannot tell us how to use their donations. Perhaps even more important, as I disclosed at the beginning of this essay, I lean left politically; if anything, I should be praising Big Tech companies for their political bias. I *don't* praise them because I love my country and our system of government more than I love any particular party or candidate, and my 12 years of rigorous research on online manipulation has convinced me that Eisenhower's 1961 prediction has come true: *The technological elite are now in control*. They have undermined the integrity of the free-and-fair election, which is a cornerstone of democracy, and U.S. authorities have not constrained the conduct of these companies in any way through laws and regulations. In addition, the data we are now collecting from the devices of children and teens (with their parents' permission) is telling me that Big Tech companies may be indoctrinating our children – at the moment, in ways I do not fully understand.

Here, then, to conclude this essay, are the reasons I believe that a permanent, large-scale, multipartisan, passive monitoring system must be established quickly in the US and, almost certainly, in other democracies around the world:

1) *Online ephemeral content is unprecedented in its ability to manipulate people's thinking and behavior*. The research my associates and I have been conducting since 2013 shows unequivocally that Google, and, to a lesser extent, other tech companies, are currently using ephemeral content – fleeting data that impacts people and then disappears, normally leaving no paper trail – to impact the thinking and behavior of most of the people on earth, probably every day. Even when these manipulations are not deliberate on the part of tech company employees or executives, their *algorithms* are producing these changes. As I explained in the section on SEME in Part Two of this essay, search engines are *inherently biased*; they *always* filter and order the answers they show you, and that changes the thinking of people who haven't yet made up their minds. The search engine is the most powerful mind control machine ever invented, and its recent integration with generative AIs will greatly increase its power in the very near future. The output of search engines – search suggestions, answer boxes, and search results – is all ephemeral. If we don't capture this fleeting content, *we will have no idea why people are thinking and behaving the way they do*. We will also have no idea about how tech companies are impacting our elections, and democracy will be little more than an illusion.

2) *Ephemeral content is controlled mainly by monopolies*. If thousands of companies were competing against each other to get our attention using ephemeral content, such content would *not* present a serious threat. After all, that is exactly how news organizations operate. As a whole – at least in a free democracy – they tend to cancel each other out. But, as a federal court recently ruled in the case *U.S. Department of Justice vs. Google* (Roller, 2024), Google is a vast monopoly, controlling about 92% of search worldwide. In an article I published in *Bloomberg Businessweek* in 2019 (Epstein, 2019c), I proposed an easy way to end Google's search monopoly, and that was by declaring its "index" – the ever-expanding database it uses to generate search results – to be a public commons. This is a regulatory maneuver that has been practiced for centuries, brought to bear whenever a service or commodity becomes a necessity: think water, gasoline, or telephone communications. With the entire world having access to Google's index, thousands of competing search engines will soon be established, each serving niche audiences

and each vying for our attention – exactly as the news organizations do now. Search would become competitive again – as it was when Google was founded – and it would also become *innovative* again. There has been no innovation in search ever since Google began to dominate that industry 20 years ago. In the meantime, because Google still has so much power, we must track and preserve its output in order to understand what it's doing.

3) *Nearly all of the largest information-handling companies in the world, along with the generative AIs they have developed, share the same political values.* As I noted earlier, various studies have confirmed that 95% of donations from Silicon Valley tech companies go to Democrats (Oberhaus, 2020). Their employees share similar values, and this homogeneity gets expressed in the algorithms they write, as well as in the frequent manual adjustments they make to how their algorithms work. As I noted in the section on the multiple platforms effect (MPE) in Part Two of this essay, our latest research shows that when people are exposed to similarly biased content on different platforms, the net effect is additive. Because the content on these platforms is ephemeral, we must monitor, preserve, and analyze this content in order to understand how these companies are impacting our society.

4) *Without monitoring systems in place, we will have no idea how tech companies are using ephemeral content to manipulate our society, our elections, and our people.* Please forgive my redundancy here, but imagine moving forward with no monitoring systems in place. That will mean that the profound impact that tech company algorithms are having on voters, on children, and even on ourselves, will remain a complete mystery. Moreover, if you wanted to know how tech companies might have interfered with a past election, *it would take a time machine to find out*.

5) *The data obtained from monitoring systems will help government officials to craft more meaningful and effective laws and regulations to keep tech companies in check.* Perhaps the main reason Big Tech companies are so entirely out of control these days is because our legislators and regulators never envisioned the kind of threats they now pose. If anything, with the passage of Section 230 of the Communications Decency Act of 1996, Big Tech companies have been shielded from litigation (Electronic Frontier Foundation, n.d.). In other words, they are not only unregulated, they also are protected from legal

actions that might be filed against them. Many of our lawmakers are struggling now to figure out what kinds of laws are needed to constrain the tech companies. Without the data from monitoring systems, they will be effectively blinded in their efforts to create effective laws.

6) *Without monitoring systems in place, our public officials will have no accurate way to measure compliance with laws or regulation they enact to prevent tech companies from interfering with our elections and children, and with human autonomy itself.* As officials in the EU have discovered in recent years, monitoring systems are essential tools for tracking a tech company's compliance with new laws and regulations intended to constrain their activities. If Amazon were prohibited by law from listing its own knockoff products ahead of competing products in its product listings, a monitoring system would detect violations of that law immediately.

7) *Real-time monitoring systems can provide active threat assessment* of the potentially existential threats that generative AI systems pose to humankind.

8) *Monitoring systems have the potential to make Big Tech companies accountable to the public.* The tech companies are private corporations, accountable at the moment only to their shareholders. They also tend to be highly secretive (Carter, 2021; Lima, 2022; Pegoraro, 2019). Monitoring systems have the potential to make these companies accountable to the public for the first time, especially if the findings from such systems are made available to the public on real-time dashboards (see https://AmericasDigitalShield.com).

Supreme Court Justice Louis B. Brandeis is often remembered for a statement he made about sunlight a century ago. We remember the statement, inaccurately, as "Sunlight is the best disinfectant." What he actually wrote is even stronger: "It is said that sunlight is the best of disinfectants, and streetlamps the best policemen" (Brandeis, 1913). Where emerging technologies have the potential to harm our society and our children, monitoring systems are as essential as sunlight.

Funding

No external funding was provided.

Ethics Approval and Consent to Participate

Approval to conduct the studies described in this essay was granted by the Institutional Review Board (IRB) of the American Institute for Behavioral Research and Technology (AIBRT). AIBRT is registered with the HHS Office for Human Research Protections (OHRP) under IORG0007755. AIRBT's IRB is registered with OHRP under number IRB00009303, and the Federalwide Assurance number for our IRB is FWA00021545.

Conflict of Interest

The author declares no conflict of interest, financial or otherwise.

Acknowledgements

This essay is based in part on a papers presented at the annual meeting of the Association for Psychological Sciences in 2013, at the International Convention of Psychological Science in 2017, and at the annual meetings of the Western Psychological Association between 2015 and 2024. It is also based on testimony before Congressional committees in 2019 and 2023. The author thanks the nearly 100 staff members and interns – most of them volunteers – who have helped him conduct research on online manipulation since 2013. He is also grateful to the tens of thousands of people who have made donations, large and small, to AIBRT, which have supported the basic research studies reported in this essay and that have made it possible for AIBRT's Tech Watch team to build our country's first nationwide digital monitoring system. The author is especially grateful to Phillip W. Dyck, a research intern at AIBRT, who spent many hours helping me to prepare this essay.

Biography

ROBERT EPSTEIN is Senior Research Psychologist at the American Institute for Behavioral Research and Technology (AIBRT) and the former editor-in-chief of *Psychology Today* magazine. A Ph.D. of Harvard University, where he was the last doctoral student of pioneering psychologist B. F. Skinner, Dr. Epstein has published 15 books on stress management, motivation, artificial intelligence, creativity, and other topics, as well as more than 300 scientific and popular articles. Beginning in March 2020, he published a series of articles proposing a simple and economical way to eradicate the novel coronavirus without lockdowns or vaccines (see https://CarrierSeparationPlan.com). He is also a pioneer in the study of online manipulation. His 2015 report in the *Proceedings of the National Academy of Sciences* on the "Search Engine Manipulation Effect" (SEME, pronounced "seem") (https://SearchEngineManipulationEffect.com) describes one of the most powerful types of influence ever discovered in the behavioral sciences, and because SEME leaves no paper trail and is invisible to users, it is especially dangerous. Dr. Epstein's research suggests that SEME and a dozen other new methods of online influence he has discovered pose a serious threat to democracy, free speech, our children, and human autonomy. In July 2019, Dr. Epstein testified before a Congressional committee about his research on online manipulation (7-min. video here: https://EpsteinTestimony.com). In December 2023, Dr. Epstein testified again before Congress, this time reporting on his success in creating the world's first large-scale system for preserving and analyzing the ephemeral content Big Tech companies might be using to influence elections, children, and the adult human mind. His new testimony is accessible at https://2023EpsteinTestimony.com, and a dashboard that summarizes the data his new monitoring system is collecting can be viewed at https://AmericasDigitalShield.com.

References

Adams, J. (1780) From John Adams to Jonathan Jackson, 2 October 1780. Founders Online. https://founders.archives.gov/documents/Adams/06-10-02-0113

Agudo, U., & Matute, H. (2021). The influence of algorithms on political and dating decisions. *PLOS ONE, 16*(4), e0249454. https://doi.org/10.1371/journal.pone.0249454

Akbar, A., Caton, S., & Bierig, R. (2023). Personalised Filter Bias with Google and DuckDuckGo: An Exploratory Study. In: Longo, L., O'Reilly, R. (eds) Artificial Intelligence and Cognitive Science. AICS 2022. Communications in Computer and Information Science, vol 1662. Springer, Cham. https://doi.org/10.1007/978-3-031-26438-2_39

Amiri, F. (2023). *GOP subpoenas tech CEOS as part of probe into censorship*. AP News. https://apnews.com/article/technology-us-republican-party-tim-cook-business-sundar-pichai-ce0e00d7ddeff8c08a95a7fbf01080f8

Angwin, J. & Faturechi, R. (2014). *Stanford Promises Not to Use Google Money for Privacy Research*. ProPublica. https://www.propublica.org/article/stanford-promises-not-to-use-google-money-for-privacy-research

Ball, J. (2023). *Online Ads Are About to Get Even Worse*. https://www.theatlantic.com/technology/archive/2023/06/advertising-revenue-google-meta-amazon-apple-microsoft/674258/

Barrett, P. M., & Sims, J. G. (2021). *False Accusation: The unfounded claim that social media companies censor conservatives*. NYU Stern Center for Business and Human Rights. Retrieved from: https://static1.squarespace.com/static/5b6df958f8370af3217d4178/t/60187b5f45762e708708c8e9/1612217185240/NYU+False+Accusation_2.pdf

Berger, J. (2016). *Invisible Influence: The Hidden Forces That Shape Behavior*. Simon & Schuster.

Berners-Lee, Tim. (2018, Mar 12). *The web can be weaponised – and we can't count on big tech to stop it: It's dangerous having a handful of companies control how ideas and opinions are shared. A regulator may be needed.* The Guardian. Retrieved from: https://www.theguardian.com/commentisfree/2018/mar/12/tim-berners-lee-web-weapon-regulation-open-letter

Block, M. (2021, December 23). *The clear and present danger of Trump's enduring 'Big Lie.'* NPR. Retrieved from:

https://www.npr.org/2021/12/23/1065277246/trump-big-lie-jan-6-election

Blodget, H. (2012). *Google Caught Secretly Hacking Apple Software To Track Apple iPhone and Mac Users*. Business Insider. https://www.businessinsider.com/google-tracking-apple-users-2012-2?IR=T

Bogert, E., Schecter, A., & Watson, R. T. (2021). Humans rely more on algorithms than social influence as a task becomes more difficult. Scientific Reports, 11. https://doi.org/10.1038/s41598-021-87480-9

Bolyard, P. (2018, August 25). *96 percent of my Google search results for 'Trump' news were from liberal media outlets*. PJ Media. Retrieved from: https://pjmedia.com/news-and-politics/paula-bolyarda/2018/08/25/google-search-results-show-pervasive-anti-trump-anti-conservative-bias-n60450

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I.; Marlow, C., Settle, J. E., Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. Nature, 489(7415), 295–298. doi:10.1038/nature11421

Burell, I. (2019). Are Google and Facebook killing advertising?. Raconteur. https://www.raconteur.net/marketing-sales/google-facebook-duopoly

Bustillo, X. (2023). *It takes lots of money to win elections. Here's what you need to know*. NPR. https://www.npr.org/2023/11/01/1205728664/campaign-finance-donations-election-fec-fundraising-ad-spending

Carlson, M. (2018). *Facebook in the News. Digital Journalism, 6(1), 4–20.* doi:10.1080/21670811.2017.1298044

Chen, W., Pacheco, D., Yang, K. C., & Menczer, F. (2021). Neutral bots probe political bias on social media. *Nature communications*, *12*(1), 5580. https://doi.org/10.1038/s41467-021-25738-6

Cheng, J. (2023). *The Race to Build a Chat GPT Powered Search Engine*. Wired. https://www.wired.com/story/the-race-to-build-a-chatgpt-powered-search-engine/

Chimielewski, D. (2016). *A chart of lobbyists' White House visits reveal its close ties with Google*. recode. https://web.archive.org/web/20160525162743/https://www.recode.net/2016/4/26/11586424/google-white-house-visits

Coelen, J. (2023). *How Google struggled to monetize its search monopoply*. I Want Product Market-fit. https://iwantproductmarketfit.substack.com/p/how-google-struggled-to-monetize

Common Sense. (2022). *Teens and Pornography*. https://www.commonsensemedia.org/sites/default/files/research/report/2022-teens-and-pornography-final-web.pdf

Constine, J. (2014). *Leaked Documents Show How Yelp Thinks It's Getting Screwed By Google*. Tech Crunch. https://techcrunch.com/2014/07/09/yelp-google-anti-trust/

Cutts, M. (2011). Google search and search engine spam. Google Blog Spot. https://web.archive.org/web/20110121212328/http://googleblog.blogspot.com/2011/01/google-search-and-search-engine-spam.html

Dame, N. (2015). They Fooled Us All: Why Google May No Longer Announce Major Algorithm Updates. Search Engine Land. https://searchengineland.com/fooled-us-google-no-longer-announces-major-algo-updates-217494

Davies, C. (2009). Google Blacklists Entire Internet. The Guardian. https://www.theguardian.com/technology/2009/jan/31/google-blacklist-internet

Dayen, D. (2016). *Google's unusually close relationship with the White House raises lots of questions*. Mashable. https://mashable.com/article/google-relationship-obama-white-house

De Witte, M. (2019, November 26). *Search results not biased along party lines, Stanford scholars find.* Stanford News Service. Retrieved from: https://news.stanford.edu/press-releases/2019/11/26/search-media-biased/

Dean, B. (2023). Here's What We Learned About Organic Click Through Rate. BackLinko. https://backlinko.com/google-ctr-stats.

Desjardins, J. (2017). *How Google Tracks You – And What You Can Do About it*. Visual Capitalist. https://www.visualcapitalist.com/how-google-tracks-you/

D'Onfro, J. (2018, December 11). *Google's Sundar Pichai was grilled on privacy, data collection, and Chine during congressional hearing*. CNBC. Retrieved from: https://www.cnbc.com/2018/12/11/google-ceo-sundar-pichai-testifies-before-congress-on-bias-privacy.html

Döpfner, M. (2014, Apr 17). Why we fear Google. *Frankfurter Allgemeine Zeitung*. http://www.faz.net/aktuell/feuilleton/debatten/mathias-doepfner-s-open-letter-to-eric-schmidt-12900860.html

D'Souza, D., Engelbrecht, C., & Phillips, G. (Producers), & D'Souza, D., D'Souza, D., & Schooley, B. (Directors). (2022). *2000 Mules* [Motion Picture]. United States: D'Souza Media.

Duffy, C. (2021, October 4). *Facebook whistleblower revealed on '60 Minutes,' says the company prioritized profit over public good.* CNN. Retrieved from: https://edition.cnn.com/2021/10/03/tech/facebook-whistleblower-60-minutes/index.html

Dunn, J. (2016). *The tech industry's major players are firmly behind Hillary Clinton. Business Insider*. https://www.businessinsider.com/tech-company-donations-clinton-vs-trump-chart-2016-11

The Economist Group. (n.d.). *About Us.* Retrieved from: https://www.economistgroup.com/about-us

The Economist. (2019, June 8). *Google rewards reputable reporting, not left-wing politics.* Retrieved from: https://www.economist.com/graphic-detail/2019/06/08/google-rewards-reputable-reporting-not-left-wing-politics

Edelman, B. (2011). Bias in search results? Diagnosis and response. Indian Journal of Lar and Technology, 7, 16-32. Retrieved from: https://www.ijlt.in/journal/bias-in-search-results%3F%3A-diagnosis-and-response

Eisenhower, D. D. (1961, January 17). Military-industrial complex speech. Yale Law School. Retrieved from: https://avalon.law.yale.edu/20th_century/eisenhower001.asp

Epstein, R. (2012a, September 5). *Google: The Case For Hawkish Regulation*. The Kernel. https://web.archive.org/web/20121016112820/http://www.kernelmag.com/features/report/3281/google-the-case-for-hawkish-regulation/

Epstein, R. (2012b, October 23). *Why Google Should Be Regulated (Part 1).* Huffington Post. https://www.huffpost.com/entry/google-privacy_b_1962827

Epstein, R. (2012c, October 31). *Why Google Should Be Regulated (Part 2).* Huffington Post. https://www.huffpost.com/entry/online-privacy_b_2013583

Epstein, R. (2012d, November 2). *Why Google Should Be Regulated (Part 3).* Huffington Post. https://www.huffpost.com/entry/why-google-should-be-regu_b_2054111

Epstein, R. (2012e, November 5). *Why Google Should Be Regulated (Part 4).* Huffington Post. https://www.huffpost.com/entry/why-google-should-be-regu_b_2069223

Epstein, R., & Robertson, R. E. (2015). The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences USA, 112*(33), E4512-E4521. ttps://doi.org/10.1073/pnas.1419828112 (forwarding link: https://SearchEngineManipulationEffect.com)

Epstein, R. (2016a, June 22). *The New Censorship*. U.S News & World Report. https://web.archive.org/web/20240506002240/http://www.usnews.com/opinion/articles/2016-06-22/google-is-the-worlds-biggest-censor-and-its-power-must-be-regulated

Epstein, R. (2016b, September 6). *Free Isn't Freedom: How Silicon Valley Tricks Us*. Vice. https://www.vice.com/en/article/8q8vav/free-isnt-freedom-epstein-essay

Epstein, R. (2018a). *How Major News Organizations, Universities and Businesses Surrender Their Privacy to Google*. Daily Caller. https://dailycaller.com/2018/08/27/surrender-privacy-google/

Epstein, R. (2018b). *Not Just Conservatives: Google and Big Tech Can Shift Millions of Votes in Any Direction*. USA Today. https://eu.usatoday.com/story/opinion/2018/09/13/google-big-tech-bias-hurts-democracy-not-just-conservatives-column/1265020002/

Epstein, R. (2019a, January 2). How Google Shifts Votes: A 'Go-Vote'Reminder Is Not Always What You Think It Is. Epoch Times. https://www.theepochtimes.com/opinion/another-way-google-manipulates-votes-without-us-knowing-a-go-vote-reminder-is-not-what-you-think-it-is-2754073

Epstein, R. (2019b, March 22). Google, Facebook, Amazon: Warren's toothless break-up plan ignores real Big Tech threats: Surveillance and manipulation are the real threats these companies pose. They'd still have more unchecked power than any dictator under Warren's plan. USA Today. https://www.usatoday.com/story/opinion/2019/03/22/elizabeth-warren-plan-misses-dangers-facebook-amazon-google-surveillance-column/3205451002/

Epstein, R. (2019c, July 15). To break Google's monopoly on search, make its index public. Bloomberg Businessweek. https://aibrt.org/downloads/EPSTEIN-15July2019-BUSINESSWEEKTo_Break_Googles_Monopoly_on_Search.pdf (forwarding link: https://EpsteinInBusinessWeek.com) (also behind paywall at https://www.bloomberg.com/news/articles/2019-07-15/to-break-google-s-monopoly-on-searchmake-its-index-public) (audio version: https://app.newsoveraudio.com/articles/robert-epstein-tobreak-googles-monopoly-on-search-make-its-index-public-981325)

Epstein, R. (2019d, July 16). Why Google poses a serious threat to democracy, and how to end that threat (written testimony). Before the United States Senate Judiciary Subcommittee on the Constitution, Congressional Record. https://www.judiciary.senate.gov/imo/media/doc/Epstein%20Testimony.pdf

Epstein, R. (2020). The technological elite are now in control. Chapter to appear in Vol. 49 of *Champions of Freedom*. Hillsdale College Press. Retrieved from: https://aibrt.org/downloads/EPSTEIN_2020-The_Technological_Elite_Are_Now_in_Control.pdf

Epstein. R. (2023, December 13). America's Digital Shield: A new online monitoring system will make Google and other tech companies accountable to the public. Testimony before the United States Senate Judiciary Subcommittee on Competition Policy, Antitrust, and Consumer Rights (written testimony, 480 pp.). *Congressional Record.* https://aibrt.org/downloads/EPSTEIN_2023-Americas_Digital_Shield-Written_Testimony-Senate_Judiciary-13December2023-g.pdf

Epstein, R., Lee, V., Mohr, R., Jr., & Zankich, V. R. (2022). The Answer Bot Effect (ABE): A powerful new form of influence made possible by intelligent personal assistants and search engines. *PLOS ONE.* https://AnswerBotEffect.com

Epstein, R., Tyagi, C., & Wang, H. (2023, July 27). What would happen if Twitter sent consequential messages to only a strategically important subset of users? A quantification of the Targeted Messaging Effect (TME). *PLOS ONE*. https://doi.org/10.1371/journal.pone.0284495 (forwarding link: https://TargetedMessagingEffect.com, preprint posted to SSRN, September 1, 2022, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4207187)

Epstein, R., & Flores, A. (in press). The Video Manipulation Effect (VME): A quantification of the possible impact that the ordering of YouTube videos might have on opinions and voting preferences. *PLOS ONE*. (Preprint posted on SSRN, July 31, 2023, http://dx.doi.org/10.2139/ssrn.4527207, forwarding link: https://VideoManipulationEffect.com)

Epstein, R., Huang, Y., Megerdoomian, M., & Zankich, V.R. (in press). The Opinion Matching Effect (OME): A subtle but powerful new form of influence that is apparently being used on the internet. *PLOS ONE.* (Preprint posted on SSRN, August 4, 2023, https://dx.doi.org/10.2139/ssrn.4532141, forwarding link: https://OpinionMatchingEffect.com)

Epstein, R., & Li, J. (2024, March 26). Can biased search results change people's opinions about anything at all? A close replication of the Search Engine Manipulation Effect (SEME). *PLOS ONE*. https://doi.org/10.1371/journal.pone.0300727 (forwarding link: https://MultipleTopicsResearch.com, preprint posted on SSRN, October 10, 2023, https://dx.doi.org/10.2139/ssrn.4597654)

Epstein, R., Aries, S., Grebbien, K., Salcedo, A.M., & Zankich, V.R. (2024a, July 6). The Search Suggestion Effect (SSE): How search suggestions can be used to impact opinions and votes. *Computers in Human Behavior*. https://doi.org/10.1016/j.chb.2024.108342 (Preprint posted on SSRN, August 8, 2023, https://dx.doi.org/10.2139/ssrn.4535163, forwarding link: https://SearchSuggestionEffect.com)

Epstein, R., Lothringer, M., & Zankich, V.R. (2024b, January). How a daily regimen of operant conditioning might explain the power of the Search Engine Manipulation Effect (SEME). *Behavioral and Social Issues, 33*, 82–106. https://doi.org/10.1007/s42822-023-00155-0

Epstein, R., Newland, A., & Tang, L. Y. (2024c, August 2). The "multiple exposure effect" (MEE): How multiple exposures to similarly biased online content can cause increasingly larger shifts in opinions and voting preferences. Available at SSRN: https://ssrn.com/abstract=4914689

Epstein, R., Newland, A., & Tang, L. Y. (2024d, August 10). The "digital personalization effect" (DPE): A quantification of the extent to which personalizing content can increase the impact of online manipulations. Available at SSRN: https://ssrn.com/abstract=

Epstein, R., Parsick, T., Shankar, P., & Zankich, V.R. (2024e, April). The Differential Demographics Effect (DDE): Post hoc analyses of multiple datasets show the power of a new and invisible form of manipulation made possible by the internet. Paper presented at the 104th annual meeting of the Western Psychological Association, San Francisco, CA. https://DifferentialDemographicsEffect.com

Evers, A. (2020). *How Youtube became an internet giant*. CNBC. https://www.cnbc.com/2020/11/14/how-googles-youtube-became-an-internet-video-giant.html

E-Ventures Worldwide, LLC v. Google, Inc., No. 2:2014cv00646 - Document 104 (M.D. Fla. 2016). https://digitalcommons.law.scu.edu/cgi/viewcontent.cgi?article=2410&context=historical

Farley, N. (2022, Nov 17). Is Google Search Getting Worse? Search Engine land. https://searchengineland.com/is-google-search-getting-worse-389658

Flaherty, S. (2017). *Search Engine Manipulation Fuels Content Creators Fears*. BrandContent. https://brandcontent.co.uk/outfoxed/search-engine-manipulation-fuels-content-creator-fears/

Feeney, M. (2020) Conservative Big Tech Campaign Based on Myths and Misunderstanding. Cato Institute. https://www.cato.org/commentary/conservative-big-tech-campaign-based-myths-misunderstanding

Fernholz, T. (2016). *Hacked emails show Eric Schidmt played a crucial role in Team Hillary's election tech*. Quartz. https://qz.com/823922/eric-schmidt-played-a-crucial-role-in-team-hillarys-election-tech

Fingas, J. (2015). *Google will provide some of Yahoo's search results.* Yahoo News. https://www.yahoo.com/news/2015-10-20-google-yahoo-search-deal.html

Flores, R. (2016). *Hillary Clinton Google Suggestions Accused of Favoring Candidate* . CBS News. https://www.cbsnews.com/news/hillary-clinton-google-suggestions-accused-favoring-candidate-election-2016/

Forbes. (2024). *Why Big Tech Will Remain Beyond Government Control*. https://www.forbes.com/sites/hecparis/2024/01/12/why-big-tech-will-remain-beyond-government-control/

Fredrick, B. (2022, September). *Google Announce More Personalized Search Results at Search on Event*. Search Engine Journal. https://www.searchenginejournal.com/google-announces-more-personalized-search-results-at-search-on-event/466428/

Halliday, J. (2011). Google demotes 'low quality' websites in search overhaul. The Guardian. https://www.theguardian.com/technology/2011/feb/25/google-search-results

Hao, K. (2019). *Youtube is experimenting with ways to make its algorithm even more addictive*. MIT Technology Review. https://www.technologyreview.com/2019/09/27/132829/youtube-algorithm-gets-more-addictive/

Heaven, W. D. (2021). *Language models like GPT-3 could herald a new type of search engine*. MIT Technology Review. https://www.technologyreview.com/2021/05/14/1024918/language-models-gpt3-search-engine-google/

Hindman, M. (2009). The Myth of Digital Democracy. Princeton University Press.

Horling, B. & Kulick, M. (2009). *Personalized Search for everyone*. Google. Google Blogspot. https://googleblog.blogspot.com/2009/12/personalized-search-for-everyone.html?m=1

Huszár, F., Ktena, S. I., O'Brien, C., Belli, L., Schlaikjer, A., & Hardt, M. (2022). Algorithmic amplification of politics on Twitter. Proceedings of the National Academy of Sciences of the United States of America, 119(1), e2025334119. https://doi.org/10.1073/pnas.2025334119

Izundu, C. C. (2015, March 10). *Children as young as six "uploading sexually explicit content of themselves*." BBC News. https://www.bbc.com/news/newsbeat-31710517

Jaitain, V. (2022). *The walled gardens of ad tech, explained*. Blockthrough. https://blockthrough.com/blog/the-walled-gardens-of-the-ad-tech-industry-explained/

Kratz, T. (2009). *Yahoo ponders the meaning of search.* CNET. https://www.cnet.com/culture/yahoo-ponders-the-meaning-of-search/

Kanter, J. (2018, Dec) *This graph shows 90% political donations from big tech workers went to the democrats, with Googlers leading the charge*. Business Insider. https://www.businessinsider.in/this-graph-shows-90-political-donations-from-big-tech-workers-went-to-the-democrats-with-googlers-leading-the-charge/articleshow/67010538.cms

Kerns, T. (2022). *Fitbit and Google Fit stats spotted on Nest Hub smart display*. Android Police. https://www.androidpolice.com/fitbit-google-fit-stats-nest-hub-displays/

Klinkenberg, B. (2023). *Tinder Adds "Swipe The Vote" So You Can Hook Up With Candidates*. BuzzFeed News. https://www.buzzfeednews.com/article/brendanklinkenberg/tinder-wants-you-tovote#.ep8lDQxX4o

Gatewood, C. & O'Connor C. (2020). *Disinformation Briefing: Narratives around Black Lives Matter and voter fraud*. Institute for Strategic Dialogue. https://www.isdglobal.org/isd-publications/disinformation-briefing-narratives-around-black-lives-matter-and-voter-fraud/

Gilder, G. (2018). Life after Google: The fall of big data and the rise of the blockchain economy. Regnery Gateway.

Gogarty, K., Silva, S., & Evans, C. (2020). *A New Study Finds That Facebook is not Censoring Conservatives*. Media Matters. https://www.mediamatters.org/facebook/new-study-finds-facebook-not-censoring-conservatives-despite-their-repeated-attacks

González-Bailón, S., d'Andrea, V., Freelon, D., & De Domenico, M. (2022). The advantage of the right in social media news sharing. PNAS nexus, 1(3), pgac137. doi: 10.1093/pnasnexus/pgac137.

Google. (n.d.-a). *About Google: How our business works*. Retrived July 8, 2024 from https://about.google/how-our-business-works/

Google. (n.d.-b). Personalization & Google Search results - Google Search Help. https://support.google.com/websearch/answer/12410098?hl=en

Google. (n.d.-c). *Why is my site labeled as dangerous in Google Search?*. Retrieved May 9, 2024 from https://support.google.com/webmasters/answer/6347750?hl=en

Hardy, Q. (2016, June 7). *The web's creator looks to reinvent it.* New York Times. Retrieved from:

https://www.nytimes.com/2016/06/08/technology/the-webs-creator-looks-to-reinvent-it.html
Hawking, S. (2018). *Brief Answers To The Big Questions*. John Murray.
Howard, J. J., & Rabbitt, L. R., & Sirotin, Y. B. (2020). Human-algorithm teaming in face recognition: How algorithm outcomes cognitively bias human decision-making. *PLOS ONE*, *15*(8), e0237855. https://doi.org/10.1371/journal.pone.0237855
Lahey, C. & Skopec, C. (2024, Jan 11). *34 Eye-Opening Google Search Statistics for 2024*. Semrush Blog. https://www.semrush.com/blog/google-search-statistics/
Langley, H. (2023). Google recently cut 'people' from its Search guidelines. Now, website owners say a flood of AI content is pushing them down in search results. https://www.businessinsider.com/google-search-helpful-content-update-results-drop-ai-generated-2023-9?r=US&IR=T
Lee, A. (2011). Google Takes Down Over 11 Million CO.CC Sites. Huffington Post. https://www.huffpost.com/entry/google-cocc_n_891696
Leith, Douglas. (2023). What Data Do the Google Dialer and Messages Apps on Android Send to Google?. 10.1007/978-3-031-25538-0_29.
Lewis, P. (2017). 'Our minds can be hijacked': the tech insiders who fear a smartphone dystopia’. The Guardian. https://www.theguardian.com/technology/2017/oct/05/smartphone-addiction-silicon-valley-dystopia.
Levy, Ari. (2020). Here’s the final tally of where tech billionaires donated for the 2020 election. CNBC. https://www.cnbc.com/2020/11/02/tech-billionaire-2020-election-donations-final-tally.html
Liao, S. (2018). *Google still tracks you through the web if you turn off Location History*. The Verge. https://www.theverge.com/2018/8/13/17684660/google-turn-off-location-history-data
Loftus, E. F. (1975). *Leading questions and the eyewitness report. Cognitive Psychology, 7(4), 560–572.* doi:10.1016/0010-0285(75)90023-7
Lohr, S. (2021). *He Created The Web. Now He’s Out to Remake the Digital World.* New York Times. https://www.nytimes.com/2021/01/10/technology/tim-berners-lee-privacy-internet.html
Lomas, N. (2016). *Why did ProtonMail vanish from Google search results for months?* Tech Crunch. https://techcrunch.com/2014/07/09/yelp-google-anti-trust/

Jatain, V. (2022). The walled gardens of ad tech, explained. BlockThrough. https://blockthrough.com/blog/the-walled-gardens-of-the-ad-tech-industry-explained/

Jefferson, T. (1789*). From Thomas Jefferson to Francis Hopkinson, 13 March 1789*. Founders Online, National Archives. https://founders.archives.gov/documents/Jefferson/01-14-02-0402

The Index Project (n.d.). *CrowdFlower*. Retrieved April 2, 2024, from https://theindexproject.org/post/crowdflower

Kanai, R., Feilden, T., Firth, C., & Rees, G. (2011). Political orientations are correlated with brain structure in young adults. *Current biology : CB*, *21*(8), 677–680. https://doi.org/10.1016/j.cub.2011.03.017

Knapp, A. (2018, December 11). Google CEO Sundar Pichai answered congressional questions on China, privacy. political bias and more. Forbes. Retrieved from: https://www.forbes.com/sites/alexknapp/2018/12/11/follow-along-with-google-ceo-sundar-pichais-testimony-before-congress-here-1/?sh=1b78eb8d3272

Kulshrestha, J., Eslami, M., Messias, J., Zafar, M. E., Ghosh, S., Gummadi, K. P., Karahalios, K. (2019). Search bias quantification: investigating political bias in social media and web search. Inf Retrieval J 22, 188–227 (2019). https://doi.org/10.1007/s10791-018-9341-2

Kulwin, N. (2018, April 16). *The internet apologizes.* Intelligencer. Retrieved from: https://nymag.com/intelligencer/2018/04/an-apology-for-the-internet-from-the-people-who-built-it.html

Lanier, J. (2018). Ten Arguments for Deleting Your Social Media Accounts Right Now. Henry Holt & Co. https://www.theguardian.com/books/2018/may/30/ten-arguments-deleting-your-social-media-accounts-right-now-jaron-lanier

Lanum, N. (2022, March 29). *Twitter, Facebook, Google have repeatedly censored conservatives despite liberal doubts*. Fox News. Retrieved from: https://www.foxnews.com/media/twitter-facebook-google-censored-conservatives-big-tech-suspension

Lazer, D., Robertson, R. E., & Wilson, C. (2018). Auditing the Personalization and Composition of Politically-Related Search Engine. In WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France. ACM, New York, NY, USA, 11 pages.. https://doi.org/10.1145/3178876.3186143

Lewis, S. C., Nechushtai, E., & Zamith, R. (2023): More of the Same? Homogenization in News Recommendations When Users Search on Google, YouTube, Facebook, and Twitter. Mass Communication and Society. https://doi.org/10.1080/15205436.2023.2173609

Lipsitz, K. (2009). The Consequences of Battleground and "Spectator" State Residency for Political Participation. , 31(2), 187–209. doi:10.1007/s11109-008-9068-7

Locklear, M. (2017). Google accidentally broke the internet throughout Japan. Engadget. https://www.engadget.com/2017-08-28-google-accidentally-broke-internet-japan.html.

Lowenstein, M. (2016). Vox. https://www.vox.com/2016/5/30/11800532/google-consumer-facing-customer-service

Logg, J. M., Minson, J. A., & Moore, D. A. (2018, October 26). Do people trust algorithms more than companies realize? Harvard Business Review. Retrieved from: https://hbr.org/2018/10/do-people-trust-algorithms-more-than-companies-realize

Maheshone, M. (2021). *Over 271 Google & Products You Probably Don't Know*. Masheone. https://maheshone.com/google-products-and-services/

Matz, S. C., Teeny, J. D., Vaid, S. S., Peters, H., Harari, G. M., & Cerf, M. (2024). The potential of generative AI for personalized persuasion at scale. Scientific reports, 14(1), 4692. https://doi.org/10.1038/s41598-024-53755-0

Mayer, W. G. (2008). The Swing Voter In American Politics. Brookings Insititution Press.

McDool, E., Powell, P., Roberts, J., & Taylor, K. (2020). The internet and children's psychological wellbeing. Journal of Health Economics, 69. https://doi.org/10.1016/j.jhealeco.2019.102274

McGee, M. (2012). Google Search Quality Raters Instructions Gain New "Page Quality" Guidelines. Search Engine Land. https://searchengineland.com/google-search-quality-raters-instructions-gain-new-page-quality-guidelines-132436

McKinnon, J. D., & MacMillan, D. (2018, September 20). *Google workers discussed tweaking search function to counter travel ban: Company says none of proposed changes to search results were ever implemented.* The Wall Street Journal. Retrieved from: https://www.wsj.com/articles/google-workers-discussed-tweaking-search-function-to-counter-travel-ban-1537488472

McNamee, R. (2020). Zucked: Waking Up To The Facebook Catastrophe (Rev.ed). Penguin Books.

Metaxa, D., Park, J. S., Landay, J. A., & Hancock, J. (2019). Search media and elections: A longitudinal investigation of political search results in the 2018 U.S. elections. *Proceedings of the ACM on Human-Computer Interaction, 3*(CSCW), 1-17. https://doi.org/10.1145/3359231

Moon, H. & Pariseau, G. (2023). *Alarming Election Interference! Big Tech Censors Biden Opponents 162 times*. Censor Track. https://censortrack.org/alarming-election-interference-big-tech-censors-biden-opponents-162-times

Moreno, J. (2019). *Google Follows A Growing Workplace Trend: Hiring More Contractors Than Employees*. Forbes. https://www.forbes.com/sites/johanmoreno/2019/05/31/google-follows-a-growing-workplace-trend-hiring-more-contractors-than-employees/

Morrison, R. (2022). Companies back 'open, private' internet in face of Big-Tech 'surveillance capitalism'. Tech Monitor. https://techmonitor.ai/policy/privacy-and-data-protection/open-internet-big-tech.

Mowery, D.C. & Simcoe, T. (2001). Is the internet a US invention? – An economic and technological history of computer networking. *Research Policy*, *31*(8-9), 1369-1987. https://doi.org/10.1016/S0048-7333(02)00069-0

Mullins, B. & Nicas, J. (2017). *Paying Professors: Inside Google's Academic Influence Campaign*. Wall Street Journal. https://www.wsj.com/articles/paying-professors-inside-googles-academic-influence-campaign-1499785286

Murphy, K. (2011). Google dumps all 11 + million .co.cc sites from its results. The Register. https://web.archive.org/web/20110707043119/http://www.theregister.co.uk/2011/07/06/google_cans_11m_dot_co_dot_cc_sites/

Nakashima, R. (2018a). *AP Exclusive: Google tracks your movements, like it or not*. AP News. https://apnews.com/article/828aefab64d4411bac257a07c1af0ecb

Nate Silver. (2018b). *California 45th – 2018 House forecast.* FiveThirtyEight. https://projects.fivethirtyeight.com/2018-midterm-election-forecast/house/california/45/

Nate Silver. (2018c). *California 48th – 2018 house forecast*. FiveThirtyEight. https://projects.fivethirtyeight.com/2018-midterm-election-forecast/house/california/48/

Nate Silver. (2018). California 45th – 2018 House forecast. FiveThirtyEight. https://projects.fivethirtyeight.com/2018-midterm-election-forecast/house/california/45/

Newton, C. (2018, May 10). *Google's new focus on well-being started five years ago with this presentation.* The Verge. Retrieved from: https://www.wsj.com/articles/google-workers-discussed-tweaking-search-function-to-counter-travel-ban-1537488472

The Nielsen Company. (n.d.). *Solutions: Audience Measurement*. Retrieved April 1, 2024, from https://www.nielsen.com/solutions/audience-measurement/

Nolan, L. (2020, September 22). *RealClearPolitics: Google suppresses conservative media in search results*. Breitbart. Retrieved from: https://www.breitbart.com/tech/2020/09/22/realclearpolitics-google-suppresses-conservative-media-in-search-results/

Nyguen, G. (2021). *How Google and Yelp handle fake users and policy violations*. Search Engine Land. https://searchengineland.com/how-google-and-yelp-handle-fake-reviews-and-policy-violations-374071

Nylen, L. (2024). Google's Payments to Apple Reached $20 Billion in 2022, Antitrust Court Documents Show. Bloomberg. https://www.bloomberg.com/news/articles/2024-05-01/google-s-payments-to-apple-reached-20-billion-in-2022-cue-says

Oberhaus, D. (2020). *Silicon Valley Opens Its Wallet for Joe Biden*. Wired. https://www.wired.com/story/silicon-valley-opens-wallet-joe-biden/

O'Connor, R. (2012). *Google Is Evil*. Wired. https://www.wired.com/2012/06/opinion-google-is-evil/

O'Neil, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishing Group.

Orenstein, D. (2011). *Google and Stanford's relationship provides answers to tough problems*. Standford University. https://engineering.stanford.edu/magazine/google-and-stanfords-relationship-provides-answers-tough-problems

Passifiume, B. (2022) *'Twitter Files' explained, and what they revealed about tech censorship*. National Post. https://nationalpost.com/news/twitter-files-explained-and-what-they-revealed-about-tech-censorship

Pegoraro, R. (2022). The little-known reason why competing with google is so hard. Fast Company. https://www.fastcompany.com/90709672/the-little-known-reason-why-competing-with-google-is-so-hard

Prolific (n.d.). *Solutions: Research and Development.* Retrieved April 2, 2024, from https://www.prolific.com/research-and-development

Quan, S. (2024, Jan 23). *Top Google Searches*. ahrefsblog. https://ahrefs.com/blog/top-google-searches/

Reneau, A. *Expert on tyranny explains the Trump 'Big Lie' you've been hearing so much about lately*. Upworthy. Retrieved from: https://www.upworthy.com/the-trump-big-lie-explained-historian

Reuters Staff. (2021, January 12). *Father of the Web Tim Berners-Lee prepares 'do-over.'* Reuters. Retrieved from: https://www.reuters.com/article/uk-tech-bernerslee-interview-idUSKBN29H1JG

Riehm, K. E., Feder, K. A., Tormohlen, K. N., Crum, R. M., Young, A. S., Green, K. M., Pacek, L. R., et al. (2019). Associations between time spent using social media and internalizing and externalizing problems among US youth. *JAMA Psychiatry, 76*(12), 1266-1273. doi:10.1001/jamapsychiatry.2019.2325

Roller, A. (2024). *Google antitrust case explained: What's next?*. TechTarget. https://www.techtarget.com/whatis/feature/Google-antitrust-case-explained-Whats-next

Romm, T. (2020). The Justice Department is suing Google — but it’s the government’s power to police big tech that’s on trial. https://www.washingtonpost.com/technology/2020/10/21/justice-department-google-lawsuit-congress/

Rushe, D. (2014). US government requests for Google user data jumps 120% since 2009. The Guardian. https://www.theguardian.com/technology/2014/mar/27/us-government-requests-google-user-data

Sandler, R. (2020). *Republicans Use Tech Anti-trust hearing to rail against “Anti-Conservative Bias*”. Forbes. https://www.forbes.com/sites/rachelsandler/2020/07/29/republicans-use-tech-antitrust-hearing-to-rail-against-anti-conservative-bias/

Sawyer, P. (2022, April 23). *Big Brother (sorry, Big Person) is correcting you on Google.* The Telegraph. Retrieved from: https://www.telegraph.co.uk/news/2022/04/23/big-brother-sorry-big-person-correcting-google/

Schmidt, D.C. (2018). Google Data Collection. Digital Content Next. https://digitalcontentnext.org/

Schwartz, B. (2010, September 8). *Google Instant: Predictions Were Right*. Search Engine Roundtable. https://www.seroundtable.com/archives/022883.html

Schwartz, B. (2017, October 3). *Google autocomplete showing 10 suggestions again*. Search Engine Round Table. https://www.seroundtable.com/google-autocomplete-10-suggestions-24560.html

Schwartz, B. (2022, August 18). Google’s new helpful content update targets sites creating content for search engines first. *Search Engine Land*. https://searchengineland.com/googles-new-helpful-content-update-targets-sites-creating-content-for-search-engines-first-387237

Schwartz, B. (2024). *Report: Google Sending More Discover Traffic But Less Search Traffic But Less Search Traffic To Publishers*. Search Engine Roundtable.

https://www.seroundtable.com/google-traffic-discover-search-newzdash-37879.html

Selvarajah, M. (2018). *Business are using micro personalization to target customers*. CBC News. https://www.cbc.ca/news/canada/businesses-are-using-micro-personalization-to-target-customers-1.4832247

Shuman, V., Sander, D., & Scherer, K. R. (2013). Levels of valence. *Frontiers in Psychology*, *4*, 261. https://doi.org/10.3389/fpsyg.2013.00261

Singel, R. (2011). Google: We're Working to Help Good Sites Caught by Spam Cleanup. Wired. https://www.wired.com/2011/03/google-spam-side-effects/

Smith, J. (2020). *Facebook and Twitter employees have given $2.7million to Democrats in 2020 - ten times the amount they have donated to Republicans - as big tech faces more accusations of political bias*. Daily Mail. https://www.dailymail.co.uk/news/article-8848313/90-Facebook-Twitter-employee-political-donations-Democrats.html

Smyrnaios, N., & Rebillard, F. (2019). How infomediation platforms took over the news: A longitudinal perspective.

Soloman, M. (2014). *Google's Customer Service Steps Into The Spotlight*. Forbes. https://www.forbes.com/sites/micahsolomon/2014/05/05/google/

Solsman JE. (2018, Jan). *YouTube's AI is the puppet master over most of what you watch.* CNET. https://www.cnet.com/tech/services-and-software/youtubeces-2018-neal-mohan/

Standford Law School. (2006). Google Inc. Pledges $2M to Stanford Law School Center for Internet and Society. https://law.stanford.edu/press/google-inc-pledges-2m-to-stanford-law-school-center-for-internet-and-society/

Statistica Research Department. (2024). *Nielsen Holdings - statistics & facts*. https://www.statista.com/topics/5067/nielsen-holdings/

Statt, N. (2018, Dec 4). *Google personalizes search results even when you're logged out, new study suggests claim.* The Verge. https://www.theverge.com/2018/12/4/18124718/google-search-results-personalized-unique-duckduckgo-filter-bubble

Stoker–Walker, C. & Nature Magazine (2023). *AI Chatbots Are Coming to Search Engines*. Can You Trust Them?. Scientific American. https://www.scientificamerican.com/article/ai-chatbots-are-coming-to-search-engines-can-you-trust-them/

Sullivan D. (2018, Jan). *A reintroduction to Google's featured snippets*. Google Blog. https://blog.google/products/search/reintroduction-googles-featured-snippets/

Tech Transparency Project. (2017). Google Academics Inc. https://www.techtransparencyproject.org/articles/google-academics-inc

Thompson, N. (2018). We Need To Have An Honest Talk About Our Data. Wired. https://www.wired.com/story/interview-with-jaron-lanier/

Wakabayashi, D. (2018). *Leaked Google Video After Trump's Win Adds to Pressure From Conservatives*. New York Times. https://web.archive.org/web/20180913012639/https://www.nytimes.com/2018/09/12/technology/leaked-google-video-trump.html

Waikar, S. (2021). Why a Plan to Encourage Search-Engine Competition Failed. Stanford Business. https://www.gsb.stanford.edu/insights/why-plan-encourage-search-engine-competition-failed

Washington, G. (1793*) George Washington Papers, Series 2, Letterbooks - 1799: Letterbook 24, April 3, 1793 - March 3, 1797. April 3, - March 3, 1797*. [Manuscript/Mixed Material] Retrieved from the Library of Congress, https://www.loc.gov/item/mgw2.024/

Warren, Elizabeth. (2019, Mar 8). Here's how we can break up Big Tech. Medium. https://medium.com/@teamwarren/heres-how-we-can-break-up-big-tech-9ad9e0da324c

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, *45*, 1191–1207. https://doi.org/10.3758/s13428-012-0314-x

Weinberg, G. (n.d.). *How does Google track me even when I'm not using it?*. DuckDuckGo. https://spreadprivacy.com/how-does-google-track-me-even-when-im-not-using-it/

Woods, H. C., & Scott, H. (2016). #Sleepyteens: Social media use in adolescence is associated with poor sleep quality, anxiety, depression and low self-esteem. Journal of Adolescence, 51, 41-49. https://doi.org/10.1016/j.adolescence.2016.05.008

Wöstmann, M., Störmer, V. S., Obleser, J., Addleman, D. A., Andersen, S. K., Gaspelin, N., Geng, J. J., Luck, S. J., Noonan, M. P., Slagter, H. A., & Theeuwes, J. (2022). Ten simple rules to study distractor suppression. Progress in neurobiology, 213, 102269. https://doi.org/10.1016/j.pneurobio.2022.102269

Ybarra, M. (2016). *Is Sex In The Media Related to Sexual Behavior Among Teens*. Psychology Today.

https://www.psychologytoday.com/us/blog/connected/201606/is-sex-in-the-media-related-sexual-behavior-among-teens

Zaiceva, A. (2022). What is a Walled Garden in AdTech Industry? Setupad Blog. https://setupad.com/blog/walled-garden/

Zittrain, J. (2014). *Facebook Could Decide an Election Without Anyone Ever Finding Out.* New Republic. https://newrepublic.com/article/117878/information-fiduciary-solution-facebook-digital-gerrymandering

Zuboff, S. (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. PublicAffairs.

This monograph describes the development and deployment of a nationwide system for preserving and analyzing online ephemeral content being sent to Americans by technology companies, 24 hours a day. Online ephemeral content has been shown in controlled studies to have unprecedented power to shift people's thinking and behavior without their awareness. Normally, because such content is ephemeral, it gives tech companies the ability to influence people without leaving a paper trail for authorities to trace; hence, the importance of building systems for preserving such content. This essay also addresses an important public policy issue: To what extent, if any, have tech companies been using ephemeral content for political purposes? Dr. Epstein addresses this issue by summarizing and critiquing three recent studies that have defended the tech companies. He shows that two of these studies have ties to the tech companies and that all of them have fatally flawed methodology. He argues that because ephemeral content is highly personalized, the only way we can get an accurate picture of how such content is being employed is by "looking over the shoulders" of a large, representative sample of real users as they are receiving such content, and then aggregating and analyzing the content, much as the Neilsen company does worldwide to rate TV viewership. The monitoring system he has built aggregates and analyzes such content in real time, and it has repeatedly identified politically biased content sufficient to have shifted millions of votes in national elections in the US. Epstein concludes that large-scale monitoring systems must become a permanent feature of the internet to protect our democracy, our autonomy, and the minds of our children from potentially profound manipulations by the algorithms of Big Tech companies, both now and in the foreseeable future.

DR. ROBERT EPSTEIN is Senior Research Psychologist at the American Institute for Behavioral Research and Technology and the former editor-in-chief of *Psychology Today* magazine. A Ph.D. of Harvard University, he is a pioneer in the study of new forms of manipulation that have been made possible by the internet. He has testified twice before Congress about his research in this area. His latest Congressional testimony is at https://2023EpsteinTestimony.com (6-min. video), and a dashboard that summarizes the data his new monitoring system is collecting is at https://AmericasDigitalShield.com.

AIBRT
American Institute for Behavioral Research and Technology