ARIZONA HOUSE OF REPRESENTATIVES
Fifty-sixth Legislature - First Regular Session

HOUSE AD HOC COMMITTEE ON OVERSIGHT, ACCOUNTABILITY AND BIG TECH

Report of Interim Meeting
Tuesday, September 5, 2023
House Hearing Room 3 -- 10:00 A.M.

MINUTES RECEIVED
CHIEF CLERK'S OFFICE
9-6-2023

Convened 10:04 A.M.
Recessed 12:12 P.M.
Reconvened 1:10 P.M.
Adjourned 3:01 P.M.

Members Present	**Members Absent**
Representative Kolodin, Chairman	
Representative Aguilar	
Representative Carter	

Agenda
Original Agenda – Attachment 1

Request to Speak
Report – Attachment 2

Committee Attendance
Report – Attachment 3

Presentations

Name	**Organization**	**Attachments** (Handouts)
The Impact of Big Tech's Election Interference and How We Can Stop It	Dr. Robert Epstein	4
Free Speech Rights and Government Infuence over Social Media Platforms	James Kerwin	5

Astrickland
Amanda Strickland, Committee Secretary
September 6, 2023

(Original attachments on file in the Office of the Chief Clerk; video archives available at http://www.azleg.gov)

Interim agendas can be obtained via the Internet at http://www.azleg.gov /Interim-Committees

ARIZONA HOUSE OF REPRESENTATIVES

Convened - 10:04am
Recessed - 12:12pm

INTERIM MEETING NOTICE
OPEN TO THE PUBLIC

HOUSE AD HOC COMMITTEE ON OVERSIGHT, ACCOUNTABILITY AND BIG TECH

Reconvened - 1:10pm
Adjourned - 3:01pm

Date: Tuesday, September 5, 2023

Time: 10:00 A.M.

Place: HHR 3

Members of the public may access a livestream of the committee here:
https://www.azleg.gov/videoplayer/?clientID=6361162879&eventID=2023091000

AGENDA

1. Call to Order
2. Roll Call
3. Introduction of Committee Members
4. Opening Remarks
5. Presentations:
   - The Impact of Big Tech's Election Interference and How We Can Stop It—Dr. Robert Epstein
   - Free Speech Rights and Government Influence over Social Media Platforms—James Kerwin
6. Public Testimony
7. Closing Remarks
8. Adjournment

Members:

Representative Alexander Kolodin, Chair
Representative Cesar Aguilar
Representative Neal Carter

08/29/2023
RA

People with disabilities may request reasonable accommodations such as interpreters, alternative formats, or assistance with physical accessibility. If you require accommodations, please contact the Chief Clerk's Office at (602) 926-3032 or through Arizona Relay Service 7-1-1.

Attachment 1

PLEASE COMPLETE THIS FORM FOR THE PUBLIC RECORD

HOUSE OF REPRESENTATIVES

Please PRINT Clearly

Committee on ______ Bill Number ______

Date Sept. 5, 2023 ❑ Support ❑ Oppose ❑ Neutral

Name Robert S. Hathorne Need to Speak? ❑ Yes ❑ No

Representing Myself Are you a registered lobbyist? No

Complete Address 10295 E. Rising Sun Dr 85262

E-mail Address THEHUTCHS1@icloud.com Phone Number 602-579-0939

Comments: ______

FIVE-MINUTE SPEAKING LIMIT

Attachment 2

PLEASE COMPLETE THIS FORM FOR THE PUBLIC RECORD

HOUSE OF REPRESENTATIVES

Please PRINT Clearly

Committee on BIG Tech
Bill Number N/A
Date 9/5/23
☒ Support ☐ Oppose ☐ Neutral
Name Merissa Hamilton
Need to Speak? ☒ Yes ☐ No
Representing Self
Are you a registered lobbyist? No
Complete Address Phoenix
E-mail Address merissa@merissahamilton.com
Phone Number
Comments:

*****FIVE-MINUTE SPEAKING LIMIT*****

PLEASE COMPLETE THIS FORM FOR THE PUBLIC RECORD

HOUSE OF REPRESENTATIVES

Please PRINT Clearly

Committee on
Bill Number
Date Sept 5. 2023
☐ Support ☐ Oppose ☐ Neutral
Name Yvonne Cahill
Need to Speak? ☐ Yes ☐ No
Representing Voter in ARIZONA
Are you a registered lobbyist?
Complete Address 7987 E Via Linda Scottsdale
E-mail Address yvonne@azgopgirl.com
Phone Number 602-684-3788
Comments: Censorship of political speech –

*****FIVE-MINUTE SPEAKING LIMIT*****

PLEASE COMPLETE THIS FORM FOR THE PUBLIC RECORD

HOUSE OF REPRESENTATIVES

Please PRINT Clearly

Committee on Interim Comm on Big Tech Bill Number ______

Date 9/5/23 ☐ Support ☐ Oppose ☐ Neutral

Name Kari Lake Need to Speak? ☒ Yes ☐ No

Representing ______ Are you a registered lobbyist? NO

Complete Address 5225 N. 31st Pl.

E-mail Address Kari@karilake.com Phone Number ______

Comments: ______

FIVE-MINUTE SPEAKING LIMIT

PLEASE COMPLETE THIS FORM FOR THE PUBLIC RECORD

HOUSE OF REPRESENTATIVES

Please PRINT Clearly

Committee on Big Tech Ad Hoc Bill Number ______

Date 9/5/23 ☒ Support ☐ Oppose ☐ Neutral

Name Jeff Caldwell Need to Speak? ☒ Yes ☐ No

Representing Self Are you a registered lobbyist? No

Complete Address ______

E-mail Address Jeff4Liberty@me.com Phone Number 913-484-8404

Comments: ______

FIVE-MINUTE SPEAKING LIMIT

Information Registered on the Request to Speak System

House Ad Hoc Committee on Oversight, Accountability and Big Tech (9/5/2023)

6, Public Testimony

Support:

Merissa Hamilton, representing self; Jeff Caldwell, representing self; Mathia Hassett, representing self; John Hassett, representing self

Oppose:

Nikki Colletti, representing self; Andrea Young, representing self; RTS Warrior, representing self; Pam Throw, representing self

All Comments:

Merissa Hamilton, Self: Big Tech has interfered in our elections in Arizona. They've also assisted tyrants in AZ to use our government resources to do the same.; Jeff Caldwell, Self: ANB Systems has deployed eTRACK+ with NE utilities to allow data recorded by EVs is transferred when plugged at home to charge. 15 Min cities have "Traffic Filters." WBCSD is a CEO led organization of over 200 co's pushing for climate lockdowns.; Mathia Hassett, Self: Members of the Committee, We need to stop how Big Tech is controlling our searches, interfering in elections, and censoring social media platforms. This cannot be allowed to continue if we want to have free and fair elections in the United States.; John Hassett, Self: Big tech companies control enough of the information we seek to create a bias. They also have the ability to censor people and info that go against their narrative and have done so. Together, those two things threaten free and fair elections. Thanks!; Nikki Colletti, Self: This is election interference, whether specifically stated in the law or not. At the very least, what Hobbs did was suppression of speech and big tech allowed or encouraged it. There is no room for this in a constitutional republic!; RTS Warrior, Self: Katie Hobbs has NO business asking Twitter to censor tweets in regards to accusing Trump's base as Neo Nazi's. she should be censored-this is against freedom of speech and exceeded her authority to even request this action. THIS HAS GOT TO STOP; Pam Throw, Self: It is obvious that big tech has played a role in swaying public opinion in our elections & it must be stopped.

ARIZONA STATE LEGISLATURE

Fifty-sixth Legislature

COMMITTEE ATTENDANCE RECORD

HOUSE AD HOC
COMMITTEE ON **OVERSIGHT, ACCOUNTABILITY AND BIG TECH**

CHAIRMAN: Alexander Kolodin

	DATE	09/05/23	
	CONVENED	10:04 a m	
	RECESSED	12:12 pm	
	RECONVENED	1:10pm	
	ADJOURNED	3:01 pm	
MEMBERS			
Aguilar		✓	✓
Carter		✓	✓
Kolodin, Chairman		✓	✓

√ Present --- Absent exc Excused

Attachment 3

AIBRT American Institute for Behavioral Research and Technology

August 24, 2023

CONFIDENTIAL PROGRESS REPORT: ARIZONA SUPPLEMENT
ANALYSIS OF 2022 U.S. MIDTERM ELECTION DATA
CONTACT: DR. ROBERT EPSTEIN (re@aibrt.org)

Protecting Our Elections and Our Children from Manipulation by Big Tech

We're Building a Nationwide "Digital Shield" to Make Big Tech Companies Accountable to the Public

Summary

Midterms: Based on our analysis of more than 2.5 million "ephemeral experiences" we preserved in the days leading up to the November 8, 2022 elections through the computers of a politically-balanced group of 2,742 registered voters in 10 swing states – our "field agents" – we have so far documented four powerful forms of manipulation sufficient to have shifted upwards of 80 million votes (spread across hundreds of midterm elections): (1) substantial liberal bias in Google search results (but not Bing), (2) liberal bias in more than 75 percent of news-related videos suggested by YouTube (even though only 38 percent of news videos have such bias), (3) significantly more go-vote reminders sent to liberals than to conservatives on Google's home page on Election Day, and (4) significantly more election updates sent to liberals than to conservatives on Twitter. *Ephemeral experiences* are interactions we all have with temporary online content that is normally lost forever and that Google and other tech companies use to influence thinking and behavior. Over time, we will release more detailed data, along with state-by-state and race-by-race analyses.

Building a "Digital Shield": In the months following the 2022 election, we have continued to expand our nationwide panel of field agents and are on track to have a large-scale, permanent "digital shield" in place – more than 25,000 field agents – by the end of 2023. As of August 24, 2023, we are preserving and analyzing ephemeral content 24 hours a day through the computers of a politically-balanced group of 10,963 registered voters in all 50 states, and we have also enrolled 2,250 children and teens into our monitoring system. We have so far preserved more than **33 million ephemeral experiences** on multiple platforms. We also have evidence that our monitoring may be causing YouTube to back off on its manipulations. The more data we capture, the greater the pressure on tech companies to cease their manipulations. Without a large-scale digital shield in place, we will be surrendering our democracy and the impressionable minds of our children to the world's new Tech Lords.

Attachment 4

CONTENTS

1. Overview of the Project

Contact: Robert Epstein, Ph.D., Senior Research Psychologist, email: re@aibrt.org, mobile/text: 760-917-8152

Organization: AIBRT is a nonprofit, nonpartisan, 501(c)(3) public charity founded in 2012 to conduct behavioral research that advances the public good (see https://aibrt.org/index.php/about). EIN: 45-5623500

Background: Since 2013, AIBRT has been conducting and publishing rigorous controlled experiments that demonstrate the unprecedented power that Google and other tech companies have to shift opinions and votes on a massive scale worldwide. The experiments also show that these new forms of manipulation – among the most powerful types of influence that have ever been discovered in the behavioral sciences – can be employed without people's awareness, which makes them especially dangerous (https://SearchEngineManipulationEffect.com).

Since 2016, AIBRT has also been a leader in developing systems and software that allow us to monitor, preserve, and analyze ephemeral online content being sent to internet users by tech companies (https://TamingBigTech.com). In 2019, AIBRT researcher Dr. Robert Epstein testified before a Congressional committee (chaired by Sen. Ted Cruz) about our progress in this area (see 7-min. video at https://EpsteinTestimony.com).

In 2020, we preserved and analyzed more than 1.5 million online election-related ephemeral experiences controlled by Google and other tech companies – brief, influential experiences (such as showing people personalized search results or YouTube videos) that are normally lost forever and that can be used to manipulate people without their knowledge. We did so by recruiting 1,735 registered voters (our "field agents"), mainly in four swing states, and then, with their permission, installing custom "passive monitoring software" on their computers that allowed us to aggregate and analyze ephemeral content being sent to them (see 15-min. video at https://TheCaseForMonitoring.com).

Our 2020 findings echoed those from our 2016 and 2018 monitoring projects: Google was showing politically biased content to US voters. Among other things, its search results were highly biased in favor of one Presidential candidate, and it was sending more go-vote reminders (on its home page) to members of that candidate's party. The level of bias was sufficient in the 2020 Presidential election to have, over time, shifted six million votes among undecided voters toward the favored Presidential candidate.

On November 5, 2020, our findings prompted three U.S. Senators to send a letter to Google's CEO (https://LetterToGoogleCEO.com) expressing concern about Google's possible interference in the Presidential election. As a result, Google immediately *stopped* its online interference in Georgia's two runoff Senate races (https://TheCaseForMonitoring.com). Bias in search results disappeared, and no go-vote reminders were sent to Georgia voters.

The knowledge that their ephemeral content was being tracked and analyzed apparently caused Google's leaders to stop their manipulations. As Supreme Court Justice Louis Brandeis said a century ago, "Sunlight is the best disinfectant." To protect Americans from manipulation by emerging technologies, a permanent, nationwide monitoring system must be built.

Current project: In early 2022, we began building a larger, more comprehensive monitoring system as a step toward creating a permanent, nationwide "Digital Shield" (see https://AmericasDigitalShield.com).

By Election Day for the midterm elections – November 8, 2022 – we had recruited 2,742 field agents, again mainly in 10 swing states, and we preserved more than 2.5 million ephemeral experiences on the Google and Bing search engines, Google's home page, YouTube, Twitter, and Facebook.

For the first time – and with the permission and cooperation of their parents – we also began recruiting minors as field agents, and we began to develop custom passive monitoring software that will allow us to capture content that young people are viewing on their phones and other mobile devices.

A growing body of research suggests that a variety of internet content – content that is normally not monitored by parents, authorities, or researchers – is harming young people in the U.S. and elsewhere, especially young girls. A permanent monitoring system that continuously detects and exposes harmful content will force tech companies to modify or delete such content.

Please note: We do *not* violate the privacy of either adults or children with our monitoring systems. Unlike companies like Google and Facebook, we do *not* look at the data of individuals. When data are transmitted to us, they contain *no identifying information*, and we analyze only aggregated data.

And this time, we did not stop our recruiting efforts or shut down our monitoring system after the election. Instead, we have continued to expand our panel of field agents, the goal being to have a large, representative sample in all 50 U.S. states by the end of 2023.

This will be the world's first nationwide "Digital Shield" (see mockup on the following page). It will operate 24 hours a day, collecting and analyzing evidence of bias, indoctrination, censorship, and manipulation on multiple tech platforms, including the popular TikTok and Instagram platforms, as well as on personal assistants such as Alexa and Siri.

On an ongoing basis, we will report anomalies to journalists and public officials: members of Congress, the U.S. Attorney General, state attorneys general, members of the Federal Election Commission, parenting groups, and so on. Exposure will force the tech companies to stop their manipulations, and we will detect and report such changes as soon as they occur. A vast and ever-growing archive of data will also be stored, available for future analysis and use by regulators, legislators, advocacy groups, and the general public.

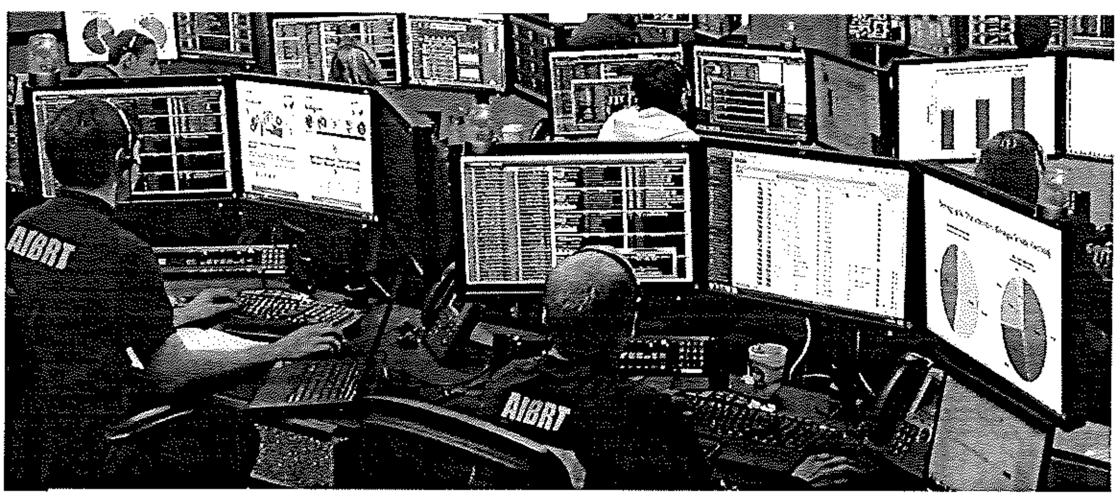

Real-time monitoring of Big Tech content 24 hours a day, with evidence of bias, censorship, manipulation, and indoctrination reported continuously to the public, journalists, and public officials.

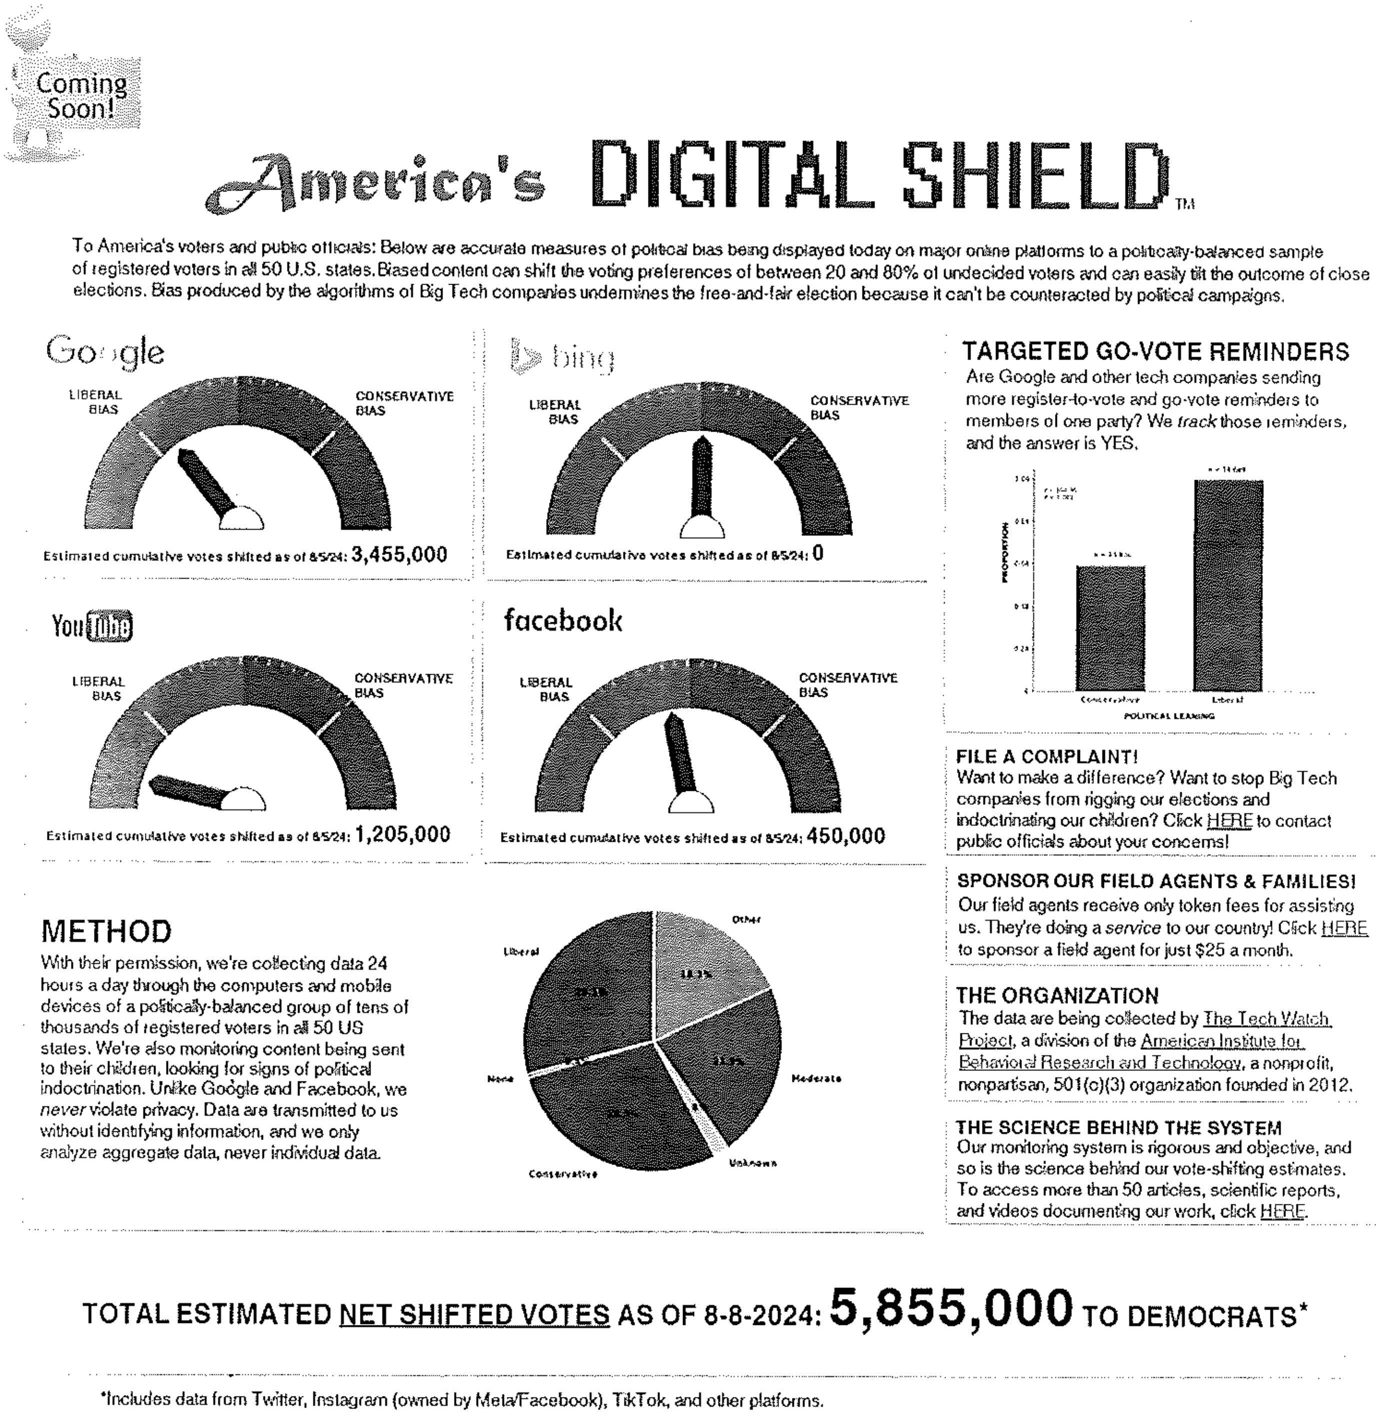

A mockup of a public dashboard that will show up-to-the-minute summaries of bias, censorship, indoctrination, and manipulation on Google and other Big Tech platforms 24 hours a day. The mockup can be viewed at https://AmericasDigitalShield.com. It will go live late in 2023.

Supplement: Arizona Midterms Data (Nov. 3-8, 2022)

a. Arizona Field Agents as of Nov. 8, 2022: 328

b. Statewide: Bias in Google Search Results

Bias in search results in Arizona followed the trend of bias in search results in other swing states: almost no bias on the Bing search engine and substantial liberal bias on the Google search engine (**Figure A_{AZ}** and **Figure B_{AZ}**). This bias was being shown to liberals, conservatives, and moderates (**Figure C_{AZ}**). The difference between the bias levels for Google and Bing was statistically significant ($P < 0.001$, $d = 0.38$). Below the statewide data, we show data for the Arizona senate race and governor's race.

Figure A_{AZ}

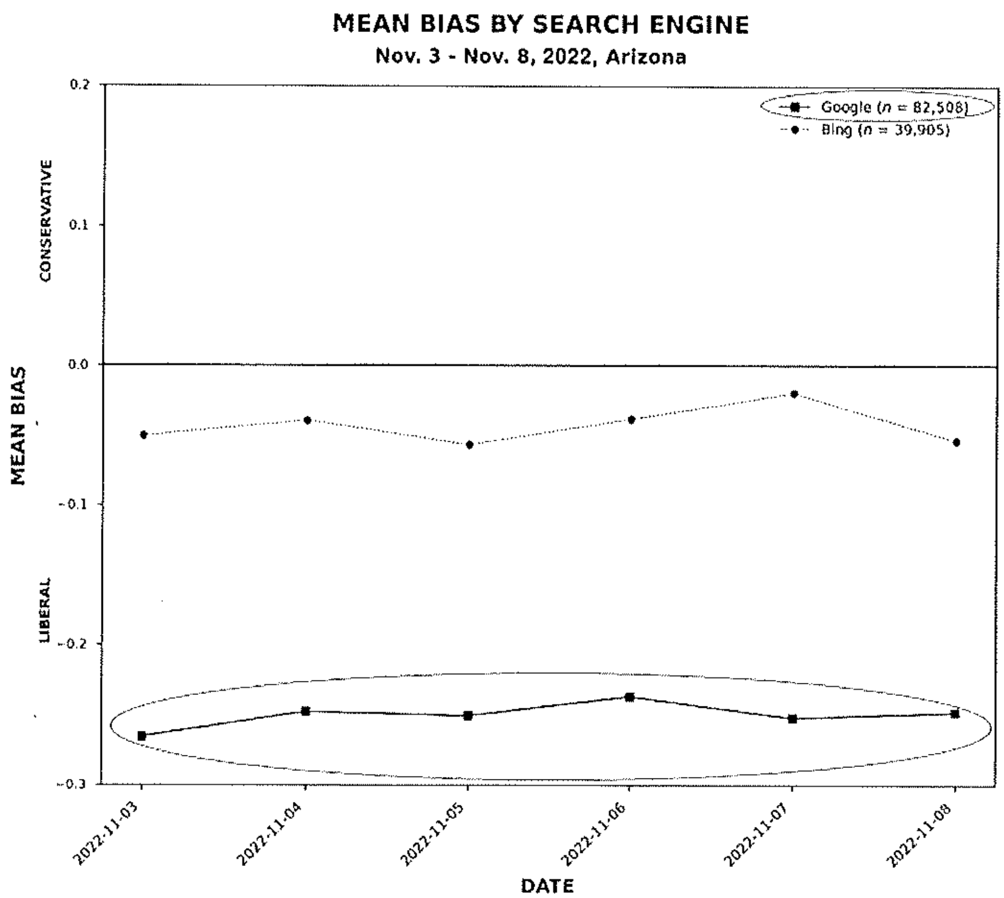

Figure B_{AZ}

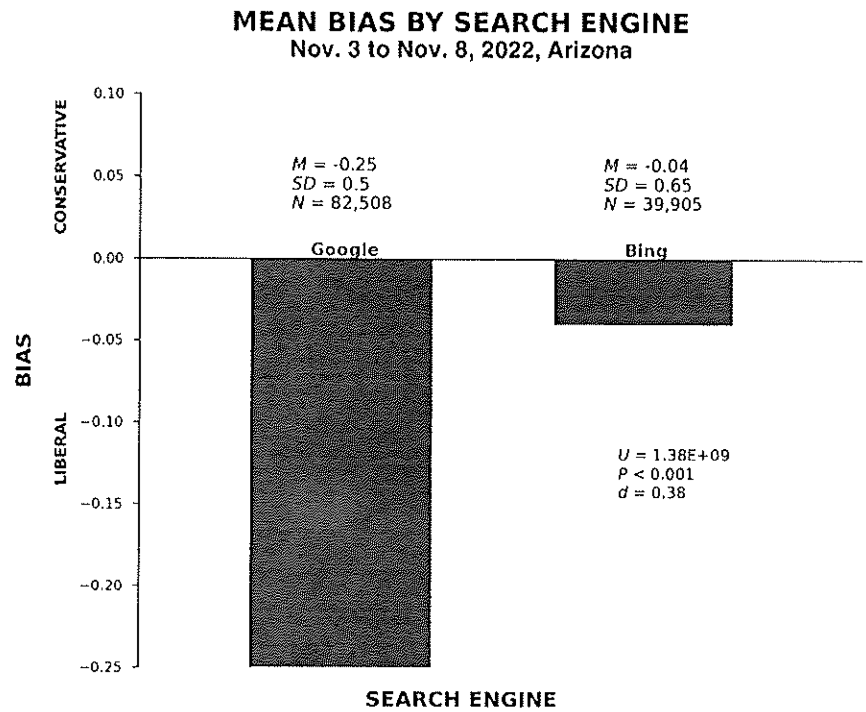

Figure C_{AZ}

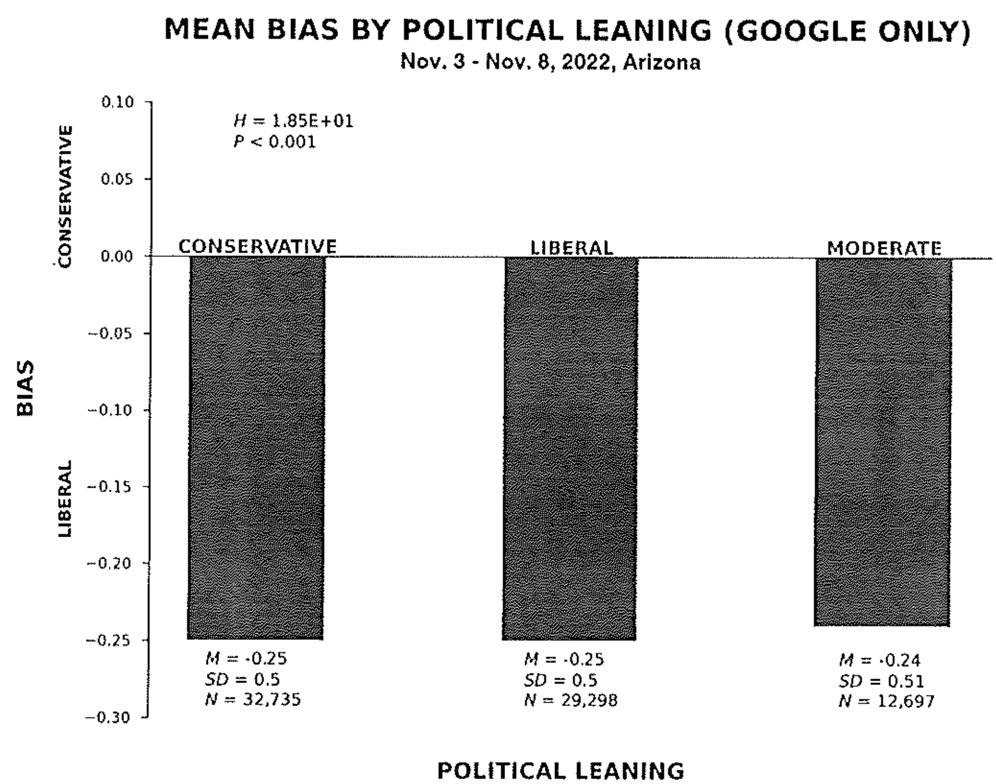

c. Statewide: Bias in YouTube Suggestions

Summary: Our preliminary analysis of the YouTube data we captured shows that 64.9% of the news-related videos YouTube (owned by Google) was suggesting to voters came from liberal news sources. Because bias in videos dramatically shifts the opinions and votes of undecided voters (see our new research at https://YouTubeManipulationEffect.com), this is an especially powerful and subtle manipulation. See data summary in the pie chart below (**Figure D$_{AZ}$**). Was YouTube simply displaying a representative sample of online news videos? Apparently not. See **Figure D'** (next page) to see the actual distribution of available online news videos categorized by political leaning. Our initial look at bias in the videos being recommend to minors shows even greater political bias than adults are seeing (on next page, see **Figure D''**).

Figure D$_{AZ}$

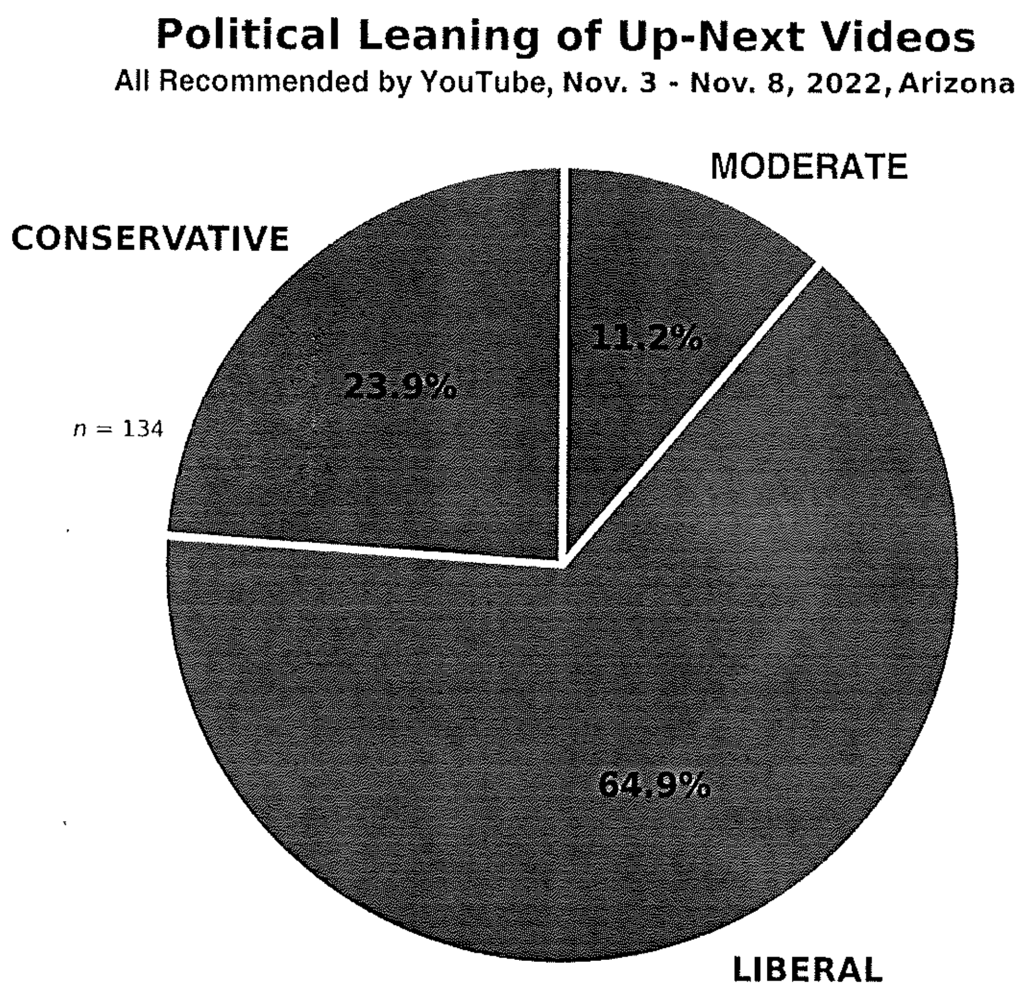

Figure D'

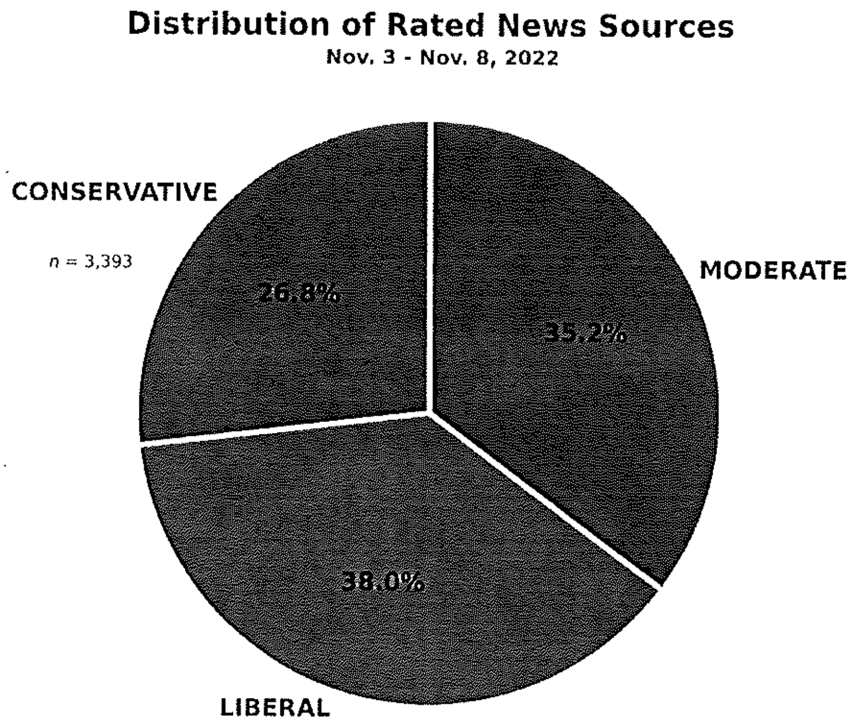

Figure D''

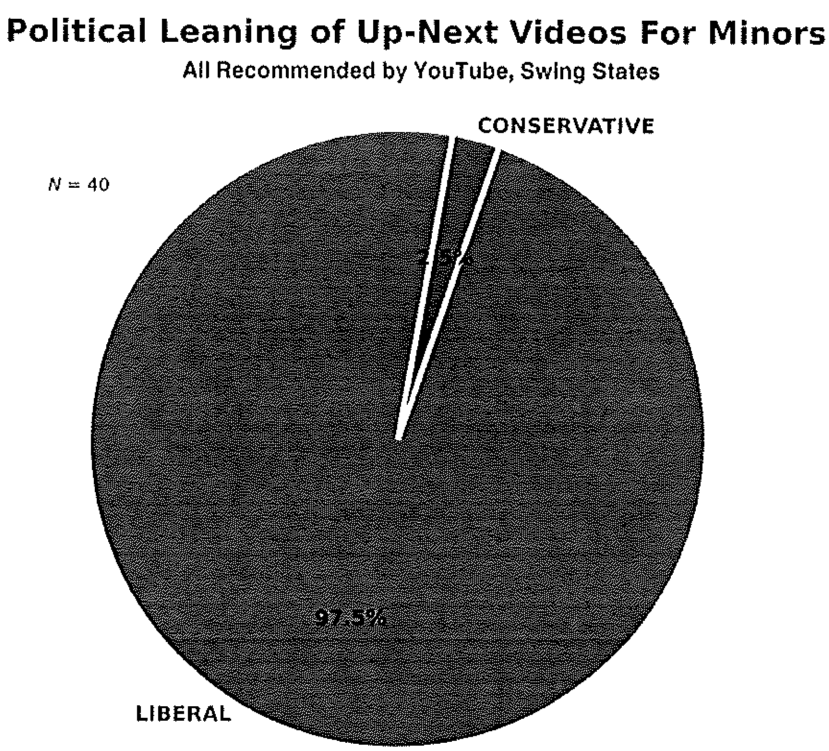

d. <u>Senate Race: Bias in Google Search Results</u>

Bias in search results in the Arizona senate race followed the trend of bias in search results statewide: almost no bias on the Bing search engine and substantial liberal bias on the Google search engine (**Figures** $\mathbf{E_{AZ}}$ and $\mathbf{F_{AZ}}$). This bias was being shown to liberals, conservatives, and moderates (**Figure** $\mathbf{G_{AZ}}$). The difference between the bias levels for Google and Bing was statistically significant ($P < 0.001$, $d = 0.44$).

Figure E_{AZ}

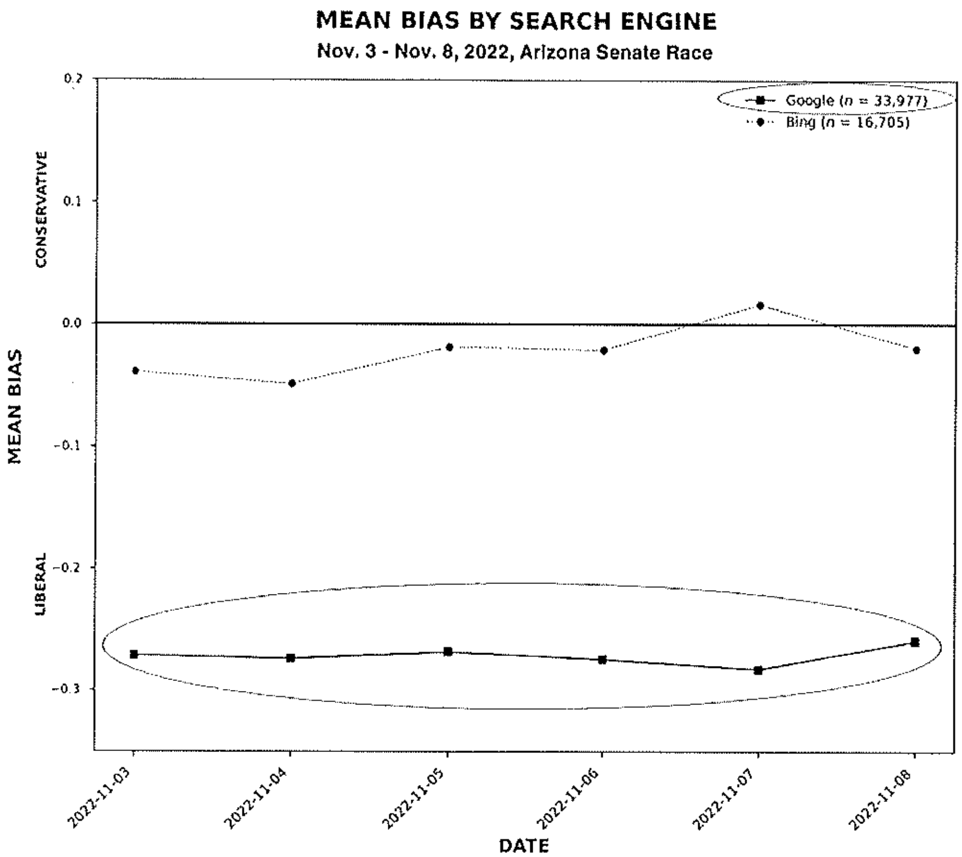

Figure F_{AZ}

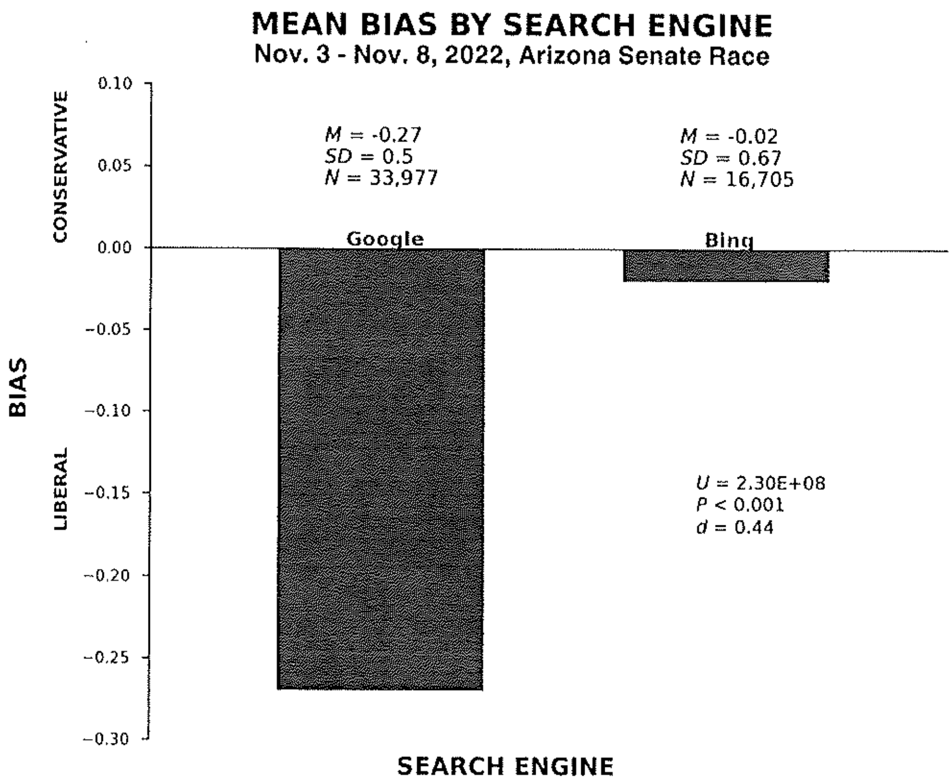

Figure G_{AZ}

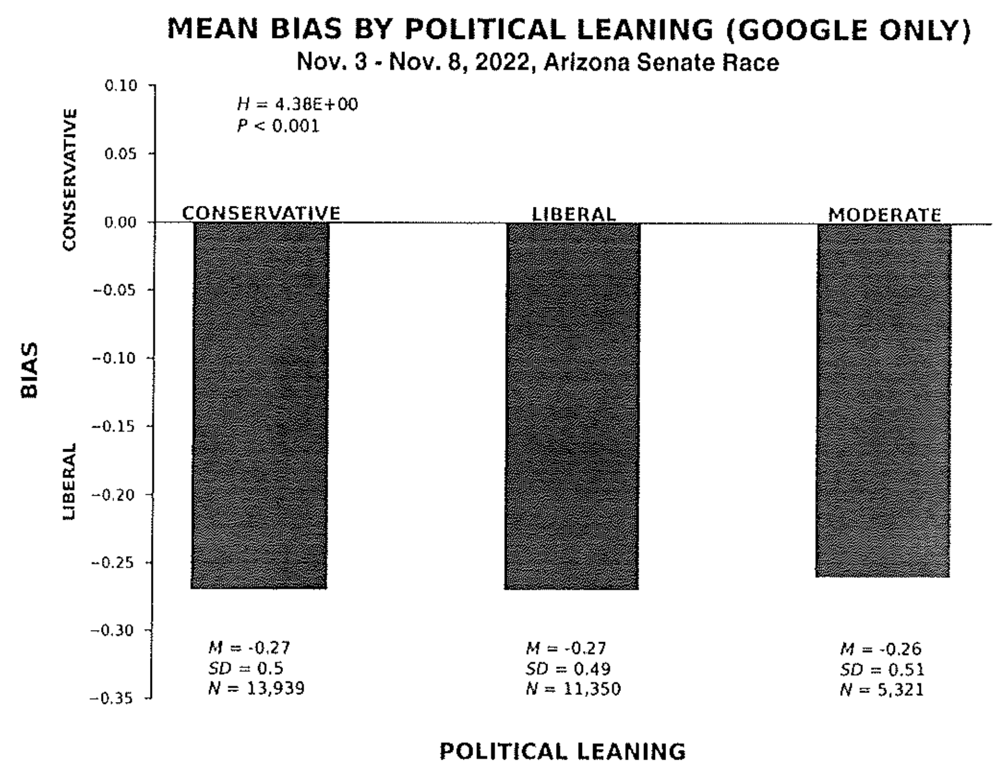

e. Governor's Race: Bias in Google Search Results

Bias in search results in the Arizona governor's race followed the trend of bias in search results statewide: almost no bias on the Bing search engine and substantial liberal bias on the Google search engine (**Figures H_{AZ}** and I_{AZ}). This bias was being shown to liberals, conservatives, and moderates (**Figure J_{AZ}**). The difference between the bias levels for Google and Bing was statistically significant ($P < 0.001$, $d = 0.34$).

Figure H_{AZ}

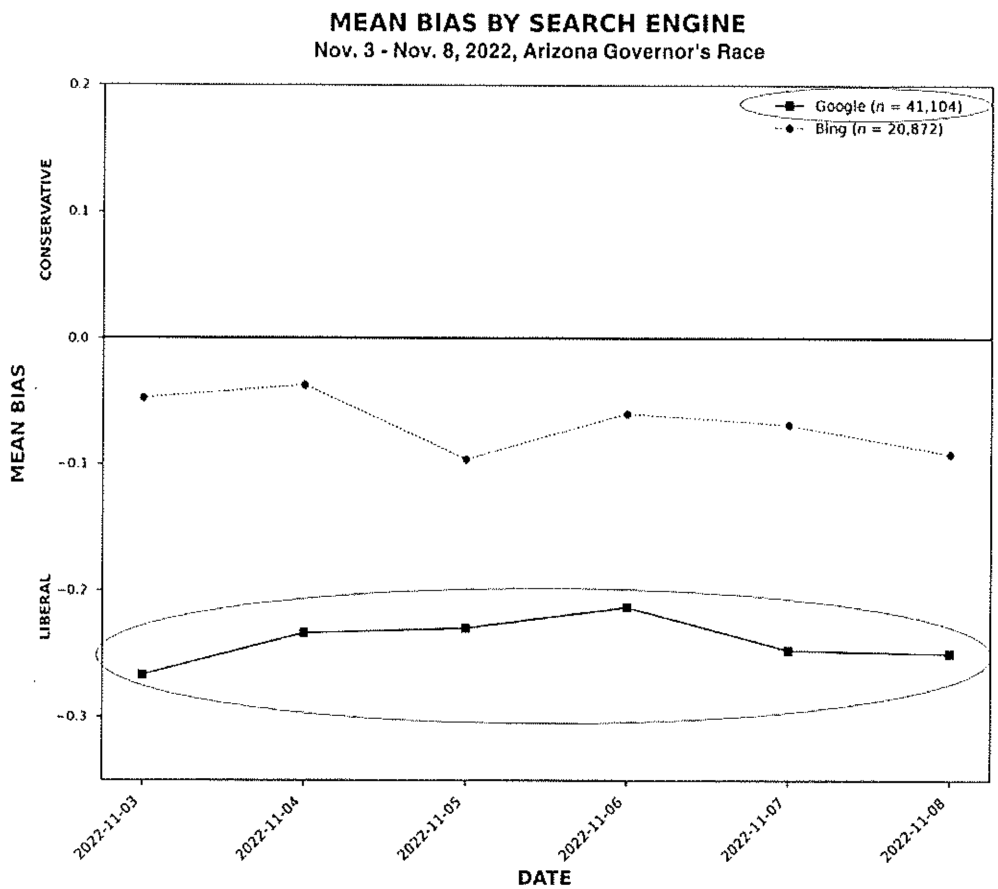

Figure I_{AZ}

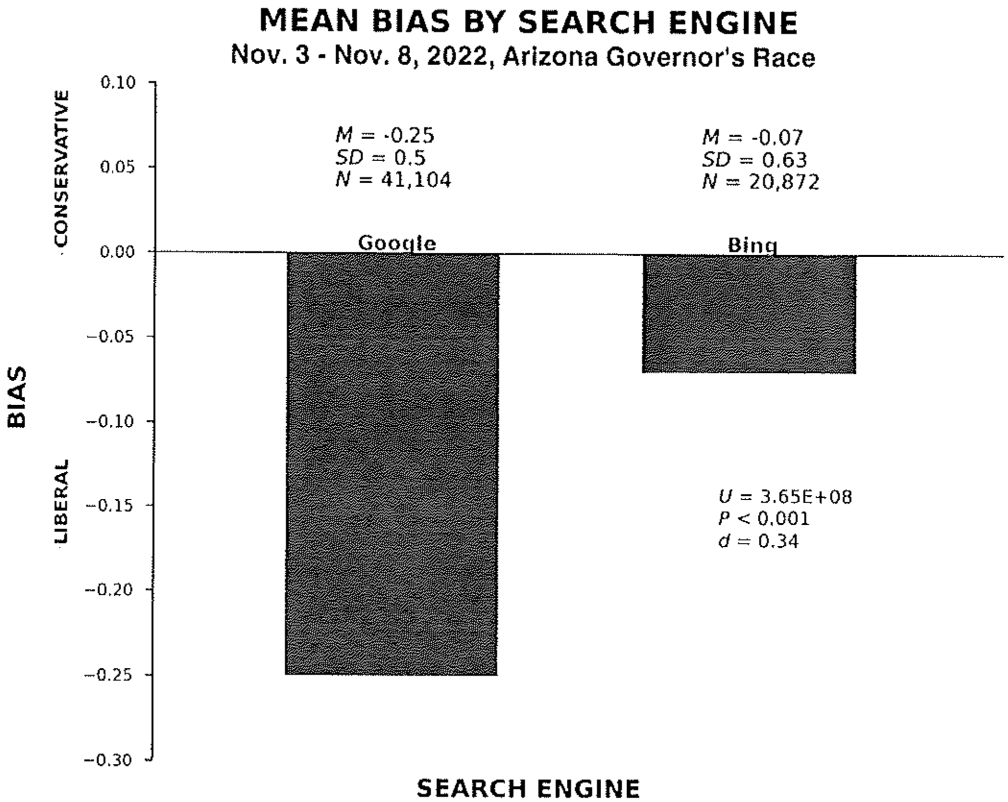

Figure J_{AZ}

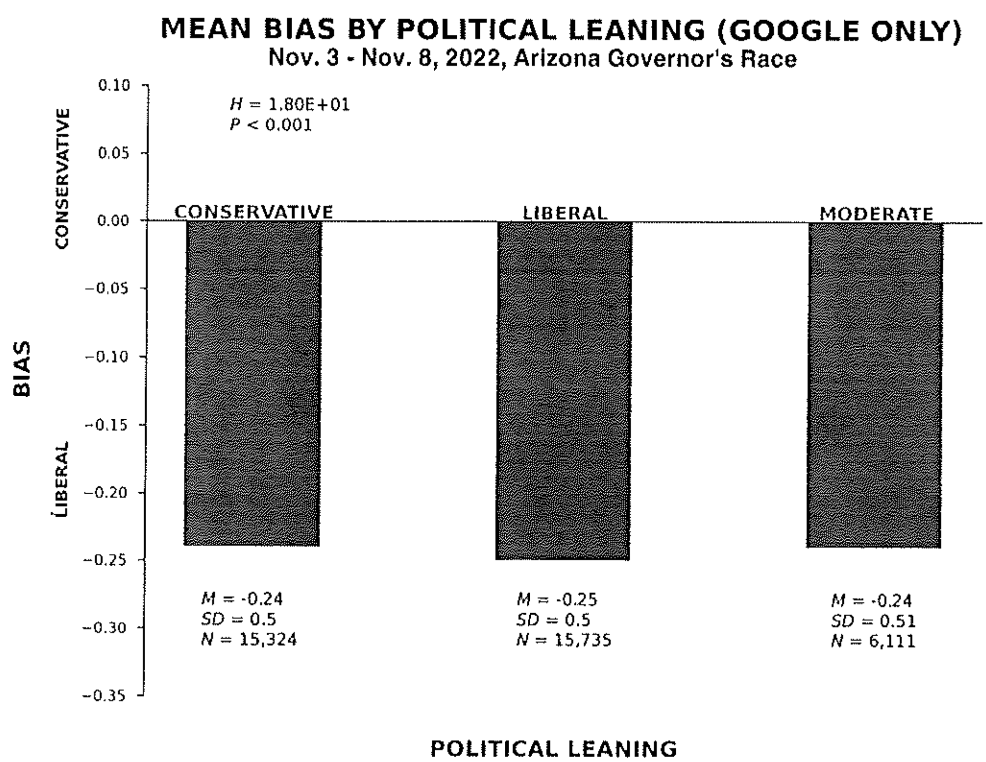

Additional Resources

- https://EpsteinTestimony.com (7-min. video, Epstein 2019 Congressional testimony)
- https://GooglesTripleThreat.com (Updated and expanded version of Epstein's 2019 Congressional testimony, with 50+ references to Epstein talks and publications on Big Tech)
- https://AmericasDigitalShield.com (mockup of real-time dashboard showing Big Tech manipulations, to be released in late 2023)
- https://TechWatchProject.org (new website focusing on the nationwide monitoring project)
- https://MyGoogleResearch.com (to support & learn more about AIBRT research on Big Tech)
- https://HowGoogleStoppedTheRedWave.com (Epstein in *The Epoch Times*)
- https://MyPrivacyTips.com (Epstein article on privacy, updated Feb. 2022)
- https://TheCaseForMonitoring.com (15-min. video about AIBRT's 2020 election findings)
- https://EpsteinOnTuckerCarlson.com (Epstein on Tucker Carlson, 56-min. video, password: "epstein" – private link, not for publication)
- https://EpsteinOnRogan.com (Epstein on Joe Rogan Experience, 160-min. video)
- https://EpsteinOnAmericanThoughtLeaders.com (90-min. video, interview on Big Tech)
- https://EpsteinOnSTEMTalks (90-min. biographical audio interview)
- https://CreepyLine.org (80-min. documentary about Big Tech, features Epstein research)
- https://TheNewCensorship.com (Epstein article on Google's blacklists, in *U.S. News & World Report*)
- https://TamingBigTech.com (article about AIBRT's 2016 monitoring project)
- https://TheAnswerBotEffect.com (new research on personal assistants)
- https://TargetedMessagingEffect.com (new research on Twitter)
- https://YouTubeManipulationEffect.com (new research on YouTube)
- https://SearchSuggestionEffect.com (new research on Google search suggestions)
- https://SearchEngineManipulationEffect.com (2015 SEME paper in *PNAS*)
- https://OpinionMatchingEffect.com (new research on biased online quizzes)
- https://LetterToGoogleCEO.com (Nov. 5, 2020 letter from U.S. Senators to Google CEO)

NEW YORK POST

MIRANDA DEVINE

Please see the last line on page 9: "Only Epstein is standing in the way."

OPINION

How Google manipulates search to favor liberals and tip elections

By Miranda Devine

May 24, 2023 7:48pm Updated

Dr. Robert Epstein is tracking Google's bias in its search results.
REUTERS

While the focus has been on Twitter and Facebook's censorship and liberal bias, the worst Big Tech culprit of all has been getting a free pass — and now it's coming for our children.

That's the warning from research psychologist Dr. Robert Epstein, a Californian Democrat with a Harvard Ph.D, who has spent the last decade monitoring Google's manipulation of newsfeeds, search results and YouTube suggestions.

He shared his latest research with The Post when he was in New York this week to raise donations for the next stage of his project.

Epstein's research shows that Google has the power to change minds and move elections to suit its liberal corporate worldview.

And despite regular protestations of innocence to Congress, the $1 trillion multinational tech giant is using its virtual monopoly as a search engine to elevate liberal views, stifle conservatives and manipulate the impressionable minds of our children.

Shifting 6 million votes

You thought Twitter and Facebook censoring The Post's Hunter Biden laptop stories was bad?

How about 6 million votes shifted to Joe Biden by Google in the 2020 election by manipulating what we read and see online?

That's the electoral impact Epstein, 69, claims Google secretly had in 2020, using biased algorithms which skewed search results towards positive links for Biden and negative links for Trump, as well as Get Out The Vote messaging on Google's home page targeted primarily at Democrat voters.

Preliminary results from Epstein's new project, monitoring how Google's massive psy-op is targeting children through YouTube and other products, show liberal bias is even more pervasive.

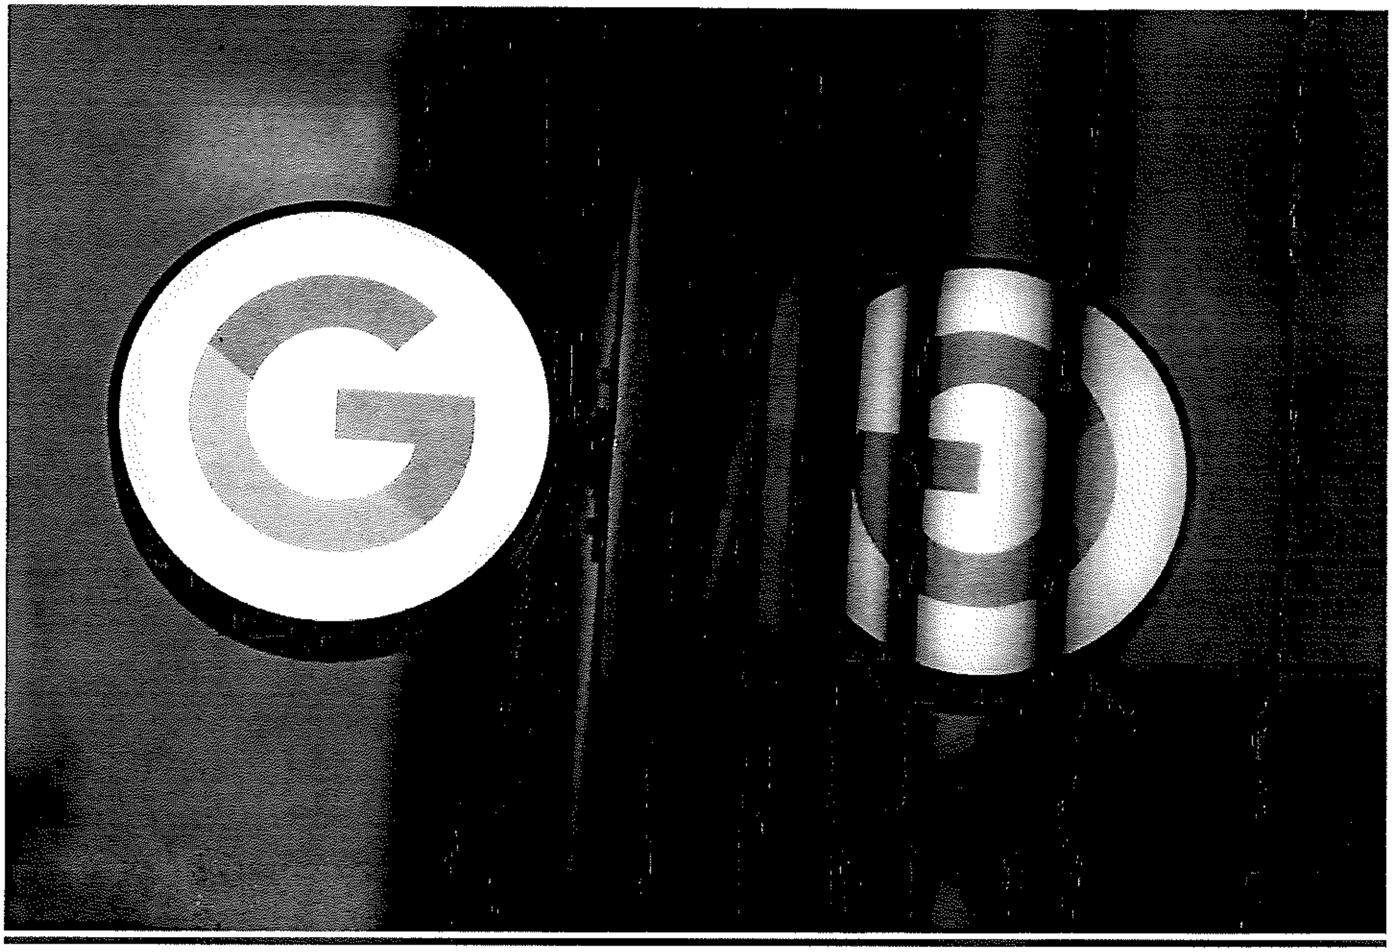

Epstein claims Google secretly had an electoral impact in 2020, using biased algorithms which skewed search results towards positive links for Biden and negative links for Trump.

Reuters

For instance, he found that YouTube's "Up Next" suggestions to adults for the next video to watch were biased towards liberal sources 76% of the time.

But for children and teens, initial data from the past three months shows the percentage of suggested videos on YouTube which come from liberal sources is 96%.

"That's how aggressive they are with our kids," he said this week. "Because they think they're gods. And no one has ever taken them to task, ever."

He did not evaluate the content of the suggested videos but used an average ranking from three nonpartisan organizations which measure media bias to assign liberal, conservative or centrist labels to the sources of the videos.

Out in the wilds of cyber space just 38% of video sources are liberal. The rest are conservative or centrist. But YouTube wildly skews its content in favor of the minority viewpoint.

While Epstein can't tell you what exactly is in the videos being promoted to children, it's not hard to guess the sorts of toxic woke ideas being pushed, thus normalizing the abnormal and rendering wholesome childhood fare the outlier.

The problem for the average person trying to guard against the manipulation is that you can't catch Google in the act because its search results or YouTube suggestions are "ephemeral," meaning they disappear once you click off onto one of the links provided, and can never be recovered.

It really is the perfect crime.

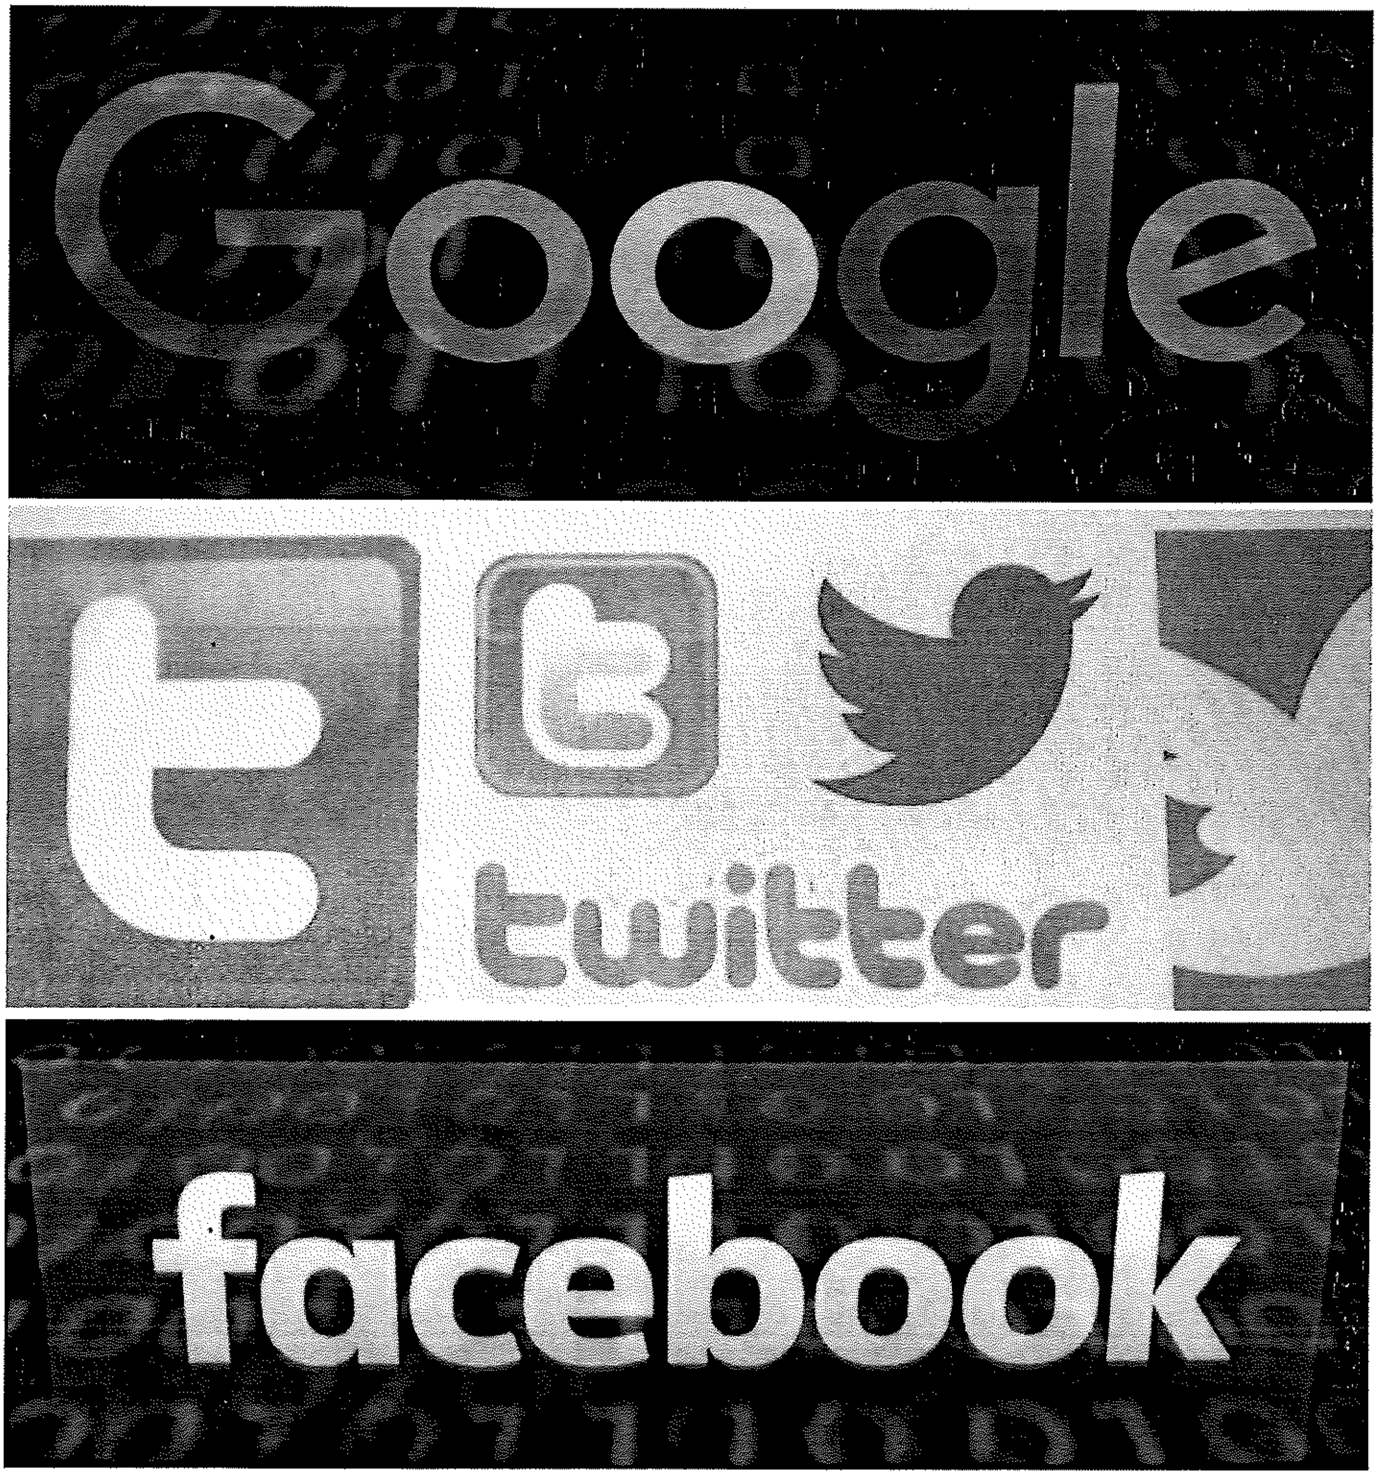

Epstein's data will track Twitter, Google, Bing and Facebook.
AFP via Getty Images

Capturing data

But Epstein has developed a way to capture that ephemeral data by effectively "looking over the shoulders" of real users, whom he calls field agents.

He now has 7,566 registered voters in all 50 states, who have given him permission to monitor and record their every Google interaction, in much the same way that Nielsen monitors TV ratings.

More recently he has added 1,600 children, aged 5 to 17, and has big plans to expand his panel to more than 25,000 field agents of all ages and political bents.

He says the best way to stop Google is to expose what they do, so he has built a public dashboard which he hopes will go live later this year at americasdigitalshield.org, in time for the 2024 election.

It will feature live tracking of bias on Google, YouTube, Facebook and Bing using real time data from his field agents across the country.

It also will show how many estimated votes would be shifted based on the level of bias.

Moderate Republicans are the most susceptible to having their minds changed, he says.

Before the 2020 election Google CEO Sundar Pichai promised Congress: "Google does not modify any products, including Search, to promote a particular political viewpoint . . . [We] will not do so for the upcoming 2020 presidential election."

But he's not fooling anyone. In terms of motive, Google is renowned for its liberal bias.

You can see that in political donations by employees of Google, YouTube, and other subsidiaries of parent company Alphabet which went 94 percent to Democrats in 2016.

Hillary Clinton's largest donor was Google/Alphabet.

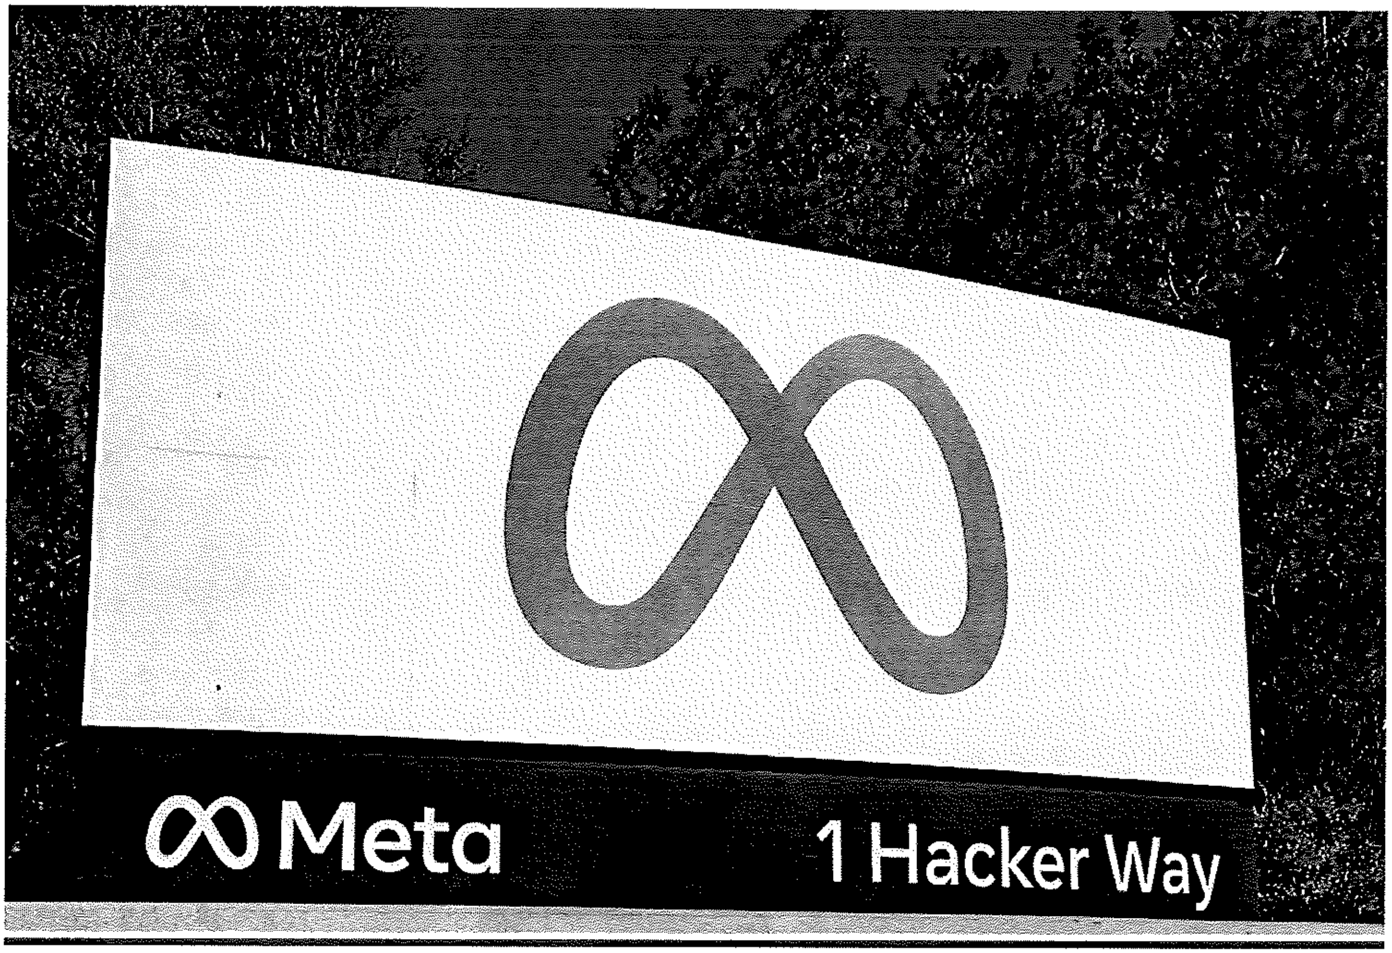

Meta censored The Post's Hunter Biden story.
AP

Google's boss of almost two decades, the $25 billion man, Eric Schmidt, played a crucial role in Hillary's campaign, as emails released by Wikileaks showed.

He also was Barack Obama's 2012 campaign tech adviser. Obama's analytics director, Elan Kriegel, told psmag.com that he credited that tech team for up to half of Obama's four point winning margin in 2012: nearly 2.5 million votes.

A leaked video of a Google executive meeting after the 2016 election gives an insight into the extreme partisan mindset of co-founders Larry Page and Sergey Brin, CEO Sundar Pichai, and other bosses lamenting Trump's win and vowing to do better next time.

Conflicting values

Brin brands Trump supporters "extremists" and says the election outcome "conflicts with many of [Google's] values."

In a portent of censorship to come, Pichai says "investments in machine learning and AI" are a "big opportunity" to address what another employee describes as "misinformation" shared by "low-information voters."

CFO Ruth Porat promises Google will "use the great strength and resources and reach we have to continue to advance really important values."

Leaked emails after the 2016 election also indicated Google had a get-out-the-vote effort with Hispanic voters in swing states in an effective in-kind contribution to the Hillary Clinton campaign.

More leaked emails, published by the Wall Street Journal, show Google employees in 2017 discussing ways to tweak search functions to change people's views about Trump's travel bans, including prioritizing pro-immigration organizations "to donate to" and "actively counter Islamophobic . . . results."

Google subsequently claimed it did not follow through with the plan.

But there is no doubt Google's culture is monomaniacally liberal and that Google executives have shown a willingness to use the power of their algorithms to intervene in elections to change voters' minds on a massive scale.

There is no reason to think it won't happen again in 2024. Only Epstein is standing in the way.

THE EPOCH TIMES

PREMIUM VIEWPOINTS

How Google Stopped the Red Wave

Google and other tech companies want you obsessing about conspiracy theories so you won't look at how they tampered with the 2022 midterm elections

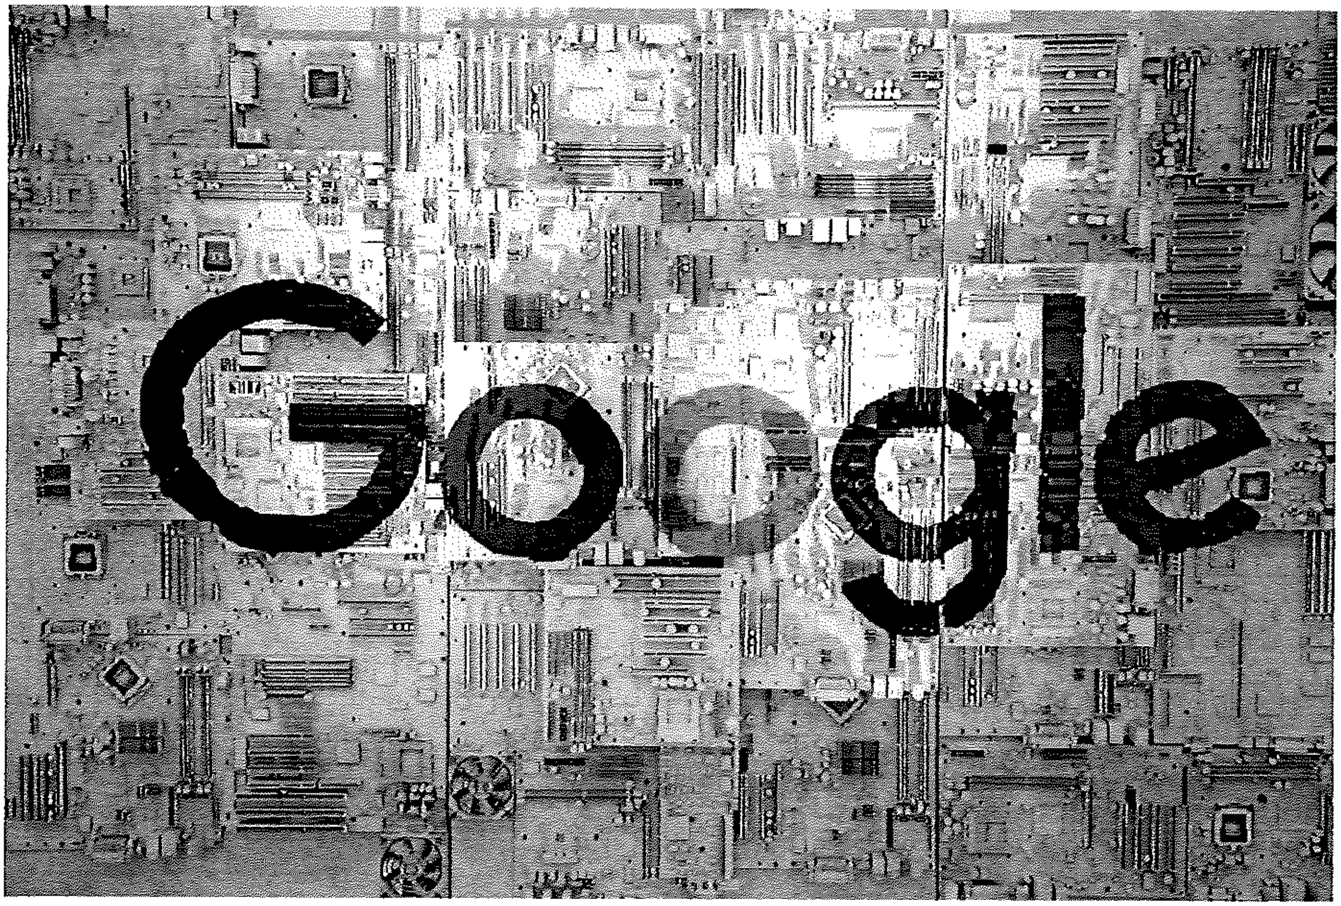

The Google logo at a data center in Belgium on Oct. 21, 2022. (NICOLAS MAETERLINCK/BELGA MAG/AFP via Getty Images)

By Robert Epstein
November 15, 2022 Updated: November 16, 2022

Commentary

What happened to the gigantic red wave that was supposed to crush the Democrats in the midterm elections? Every Republican in the country is blaming everyone else for this disaster, but almost no one is looking in the right place—and that's exactly how the Big Tech companies like it.

Based on my team's research, Google, and to a lesser extent, Facebook and other tech monopolies, not only took steps to shift millions of votes to Democrats in the midterms, but they are using their influence to spread rumors and conspiracy theories to make sure people look everywhere for explanations—except at them.

Two days before the 2022 midterm elections, I published an article explaining how Google and other tech companies were shifting millions of votes without people knowing, and I also explained how I knew, without doubt, that this was occurring.

Google isn't the only culprit, but since they're the biggest, most aggressive, and most arrogant culprit, I'll focus on them in this article.

Over a period of months, Google nudged undecided voters toward voting blue by showing people politically biased content in their search engine, suppressing content they didn't want people to see, recommending left-leaning videos on YouTube (pdf) (which Google owns), allegedly sending tens of millions of emails to people's spam boxes, and sending go-vote reminders on their home page mainly to liberal and moderate voters.

These manipulations (and others) don't affect voters with strong points of view, but they can have an enormous impact on voters who are undecided (pdf)—the people who decide the outcomes of close elections.

I know Google did these things (and more!) because, in 2022, my team and I were doing to them exactly what they do to us and our kids 24/7: We were monitoring the politically related content that Google and other tech companies were showing to actual voters—our politically diverse panel of 2,742 "field agents," who were located mainly in swing states.

In particular, we were tracking what Google employees call "ephemeral experiences"—content that appears briefly, affects people, and then disappears. In 2018, in emails that leaked from the company, Googlers were discussing how they might use ephemeral experiences to change people's views about Trump's travel ban. They know how powerful ephemeral experiences can be. That's one of the most closely held secrets of Google's management.

Ephemeral content is ideal for manipulation purposes. If you get a go-vote reminder on Google's home page (see the image below for an actual go-vote reminder sent to a liberal voter on Election Day), how would you know whether anyone else was getting it? You wouldn't, and if you didn't receive such a reminder, how would you know that anyone else had?

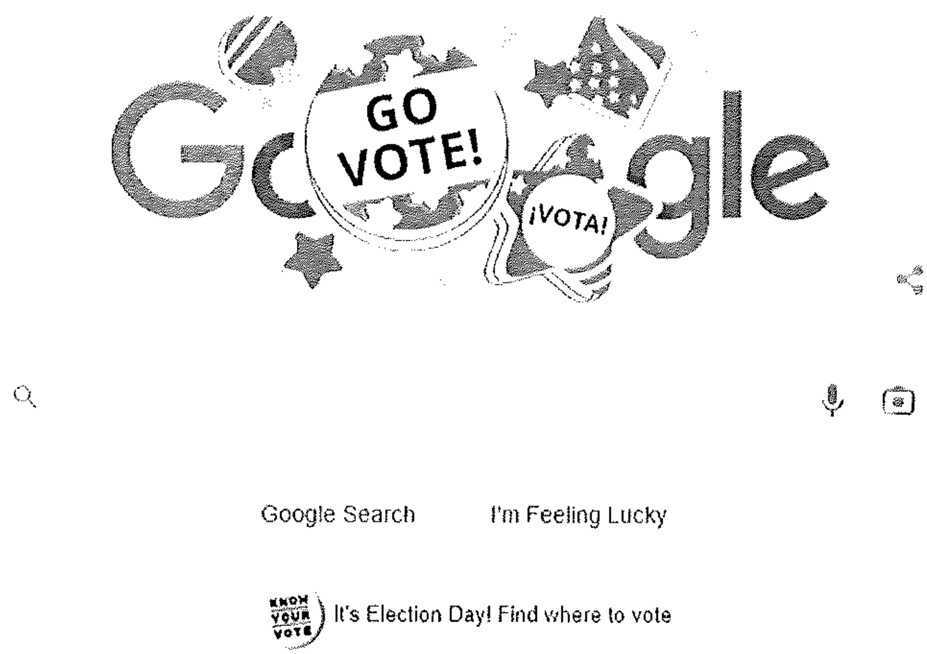

A go-vote reminder sent to a liberal voter at 11:25 a.m. on Nov. 8, 2022 (Screenshot via Google)

But we were capturing, aggregating, and analyzing the content that Google and other companies were sending to the computers of our field agents, so we could accurately estimate how many go-vote reminders Google was sending to liberals, moderates, and conservatives. In all, in the weeks leading up to the 2022 midterms, we preserved more than 2.5 million of those persuasive ephemeral experiences.

When we used similar methods to monitor content being sent by tech companies to voters before the 2020 presidential election, we found that Google was sending fewer go-vote reminders to conservatives than to moderates and liberals. Targeted messaging of this sort is a blatant manipulation that can, on Election Day in a national election in the United States, generate 450,000 extra votes for the favored candidate.

In 2020, we reported our findings to members of Congress, and on Nov. 5, 2020, three U.S. senators sent an intimidating letter (pdf) to the CEO of Google that summarized our data. As a result, Google turned off its manipulations. In the Georgia Senate runoffs that followed the presidential election, no one received a go-vote reminder from Google.

But we weren't so lucky this time around. The article I published just before the election had no effect on Google, and this year, we couldn't find a member of Congress to send a warning letter, although we came close.

As a result, Google search results remained politically biased on Election Day, and so did the up-next recommendations on YouTube. Google also sent out targeted go-vote reminders in most swing states.

If manipulations like these were being used nationwide in the months leading up to the midterm elections, Google alone might have shifted 80 million votes over time (with those votes scattered over hundreds of elections). We'll have a more precise estimate of the extent of vote shifting that occurred as we dig into our data in the coming weeks.

That's why the red wave fizzled—because Google had its digital thumb on the scale for months before the elections.

Look at history. Given inflation, the faltering economy, and President Joe Biden's low approval rating—not to mention the extensive vote redistricting that Republicans engineered in many states recently (also called gerrymandering)—the Republicans should easily have dominated the Senate races and picked up 60 or more seats in the House (as they did in the 2010 midterm elections when Barack Obama was in office). This time around, they'll be lucky to end up with a slim majority in the House and an even split in the Senate (which means that it remains under Democratic control).

Although we were unable to stop the manipulations in 2022, the good news is that we were able to preserve a treasure trove of incriminating evidence—those 2.5 million politically related ephemeral experiences. In 2023, this large dataset might be used by authorities to go after Big Tech. This will almost certainly occur if Republicans control the House.

And we're continuing to build a digital shield. By late 2023, we'll be monitoring the content that tech companies are sending to a representative sample of more than 20,000 voters and children in all 50 U.S. states 24 hours a day, and we'll report suspect content to authorities and journalists as we find it.

This digital shield—the first of its kind in the world—will protect our democracy and our children from potential manipulation by current and emerging technologies for many years to come.

Finally, a word of advice: In the coming weeks and months, you'll probably be bombarded with scary stories about how the midterm elections were tainted by rigged voting machines, fake ballots and other dirty tricks, just as you were after the 2020 presidential election. Please try your best to ignore those stories.

Dirty tricks like these are competitive; if one party can use them, so can the other. And even if some of the stories prove to be true (and most won't), dirty tricks of the sort people talk about online make little difference in election outcomes. Sometimes they shift only hundreds of votes; it's rare for them to shift thousands.

What's more, if these stories are spreading like wildfire on social media platforms, that's only because the tech companies want them to spread. Platforms such as Facebook and Instagram (both part of Meta), Twitter, and YouTube (owned by Google) have complete and absolute control over whether stories go viral.

Remember when Twitter and Facebook suppressed stories about Hunter Biden's laptop in 2020? Again, these companies can spread stories or suppress them as they please.

When you see a conspiracy theory spreading, you are often seeing an example of large-scale manipulation by misdirection. The tech companies allow such stories to spread—or even force them to spread—to turn your attention away from the companies themselves. If you think that there were fake ballots, you won't pay attention to the fact that tech companies might have shifted millions of votes in the midterms.

Sure, ballot stuffing sounds a lot more diabolical than "sending people targeted go-vote reminders," but don't let yourself be fooled. Ballot stuffing is a competitive activity that has little net effect. But targeted register-to-vote and go-vote reminders on Google's home page, which is viewed more than 500 million times a day in the United States, can shift votes by the million.

And that kind of manipulation can't be counteracted, because it's controlled exclusively by the platform. People can't even see that kind of manipulation, and—except for the monitoring my team is doing—it also leaves no paper trail for authorities to trace.

If you see a conspiracy theory blowing up on a tech platform such as Facebook or even on Fox News (which often amplifies scary stories that are spreading online), ask yourself this: Is this story real, or am I being manipulated yet again by the tech lords who have taken control of our democracy (pdf)?

The chances are good that you're being manipulated—yet again.

Views expressed in this article are the opinions of the author and do not necessarily reflect the views of The Epoch Times.

SULLIVAN & ASSOCIATES

ATTORNEYS-AT- LAW, PLLC

601 PENNSYLVANIA AVE. NW SUITE 900 WASHINGTON D.C. 20004

PHONE 703-309-7576

ATTORNEY CLIENT PRIVILEGED COMMUNICATION
ATTORNEY WORK PRODUCT

PAUL E. SULLIVAN, CA & DC

MEMORANDUM

TO:	Jason Gilbert
FR:	Paul Sullivan
DATE:	June 29, 2023
RE:	Nonpartisan Activities

Introduction: In accordance with your request and as a follow-up to our recent conversation, I have set out below a brief and generic summary of the compliance standards which are required under the Federal Election Campaign Act, as amended (FECA) and the Federal Election Commission (FEC) regulations, Title 11 (Regulations) for a corporation to undertake communications and activities which pertain to candidate and political party voter registration, get-out-the-vote (GOTV), endorsements, voter guides and voting records (jointly "Activities"). This memo is not intended to constitute a legal brief or opinion related to the issues discussed and should not be relied upon for taking any legal action of any type. Rather, the memo is structured as a guideline for a conversation; however, the points made are generally. supported by the Regulations.

Issue Presented: Your question, as I understand it, is whether entities may deliver a targeted nonpartisan[1] message pertaining to Activities. Targeting is generally understood to consist of delivering the communication to only what is perceived to be a favorable audience (e.g., specific Congressional precincts or demographic groups) or withholding the communication from specific Congressional precincts or demographic groups who are perceived to be unfavorable to a candidate or political party.

FECA Terms & Concepts: There are some basic FECA terms and concepts which are required to be understood as a foundation for this discussion. As a general rule, corporations, whether for profit or nonprofit, and labor organizations, are prohibited from utilizing their treasury funds to contribute directly or to provide in-kind contributions (making use of corporate goods or services) to influence a candidate's election or to benefit a political party.[2] There are a host of

[1] There are situations in which corporations may undertake "partisan" Activities, but though permitted, those communications would require disclosure to the FEC of the expenses associated with that partisan Activity (see fn. 2 infra).

[2] The Regulations permit corporations to use treasury funds to make independent expenditures, which could include each of the Activities discussed in this memo. That approach would require the filing of disclosure reports with the

exceptions to this general rule, but the focus of this memo will be the exceptions to that rule which enables corporations and labor organizations to use treasury funds to communicate to the general public which promotes or advocates Activities on a "nonpartisan" basis (see 11 CFR 114.4(c & d)).

Analysis: Each of the exemption Activities are separately addressed in the Regulations with distinctive compliance qualifications. However, for purposes of this generic memo, the focus must be on some common qualifications which apply to each category of the Activities. The overarching qualification is that these Activities cannot be undertaken in support or opposition, directly or indirectly, of a candidate or political party, to include not coordinating the Activities with a candidate or political party. This approach has been the basic standard since the FECA was first considered by Congress.[3]

With that general background, let me address the specific concerns you have raise. There are a couple of potential scenarios that should be addressed. Each one, however, goes back to the basic principle that utilizing a biased or prejudicial targeting of a demographic population or geographic area is prohibited. Such targeting is the anthesis of an "equitable and nonpartisan" communication (see fn. 3 infra.).

If a corporation communicates a nonpartisan voter registration message to a Congressional district or a state in the event of a Senate campaign, then the message cannot be targeted to only include geographic areas/precincts whose registered voters favor a particular political party or candidate. Correspondingly, the communication cannot be withheld from those areas whose registered voters do not favor the political party or candidate whom corporation favors. The Regulations related to this issue states:

> "The voter registration drive shall not be directed primarily to individuals previously registered with, or intending to register with, the political Party favored by the corporation or labor organization. The get-out-the-vote drive shall not be directed primarily to individuals currently registered with the political party favored by the corporation or labor organization." (11 CFR 114.4(d)(2)(ii)).

The Regulations go on to state that in regard to voter-registration or GOTV activities:

> "These services shall be made available without regard to the voter's political preference. Information and other assistance regarding registering

FEC, the details of which would be a rather extensive discussion. However, for purposes of this memo it is not needed at this time in order to evaluate the primary issue you have raised.

[3] During the May 3, 1976 floor debate of this §114.4 of the regulations, then Congressmen Hays and Wiggins agreed that corporations and labor organizations had a "free speech" right to encourage the general public to register to vote and to turn-out the vote. However, that recognition was qualified by the recognition that the activities must be carried out in an "equitable and nonpartisan basis". See January 12, 1977, Report by Federal Election Commission to Committee on House Administration, transmitting proposed regulations, pursuant to Federal Election Campaign Act, of 1971, as amended.

or voting, including transportation and other services offered, shall not be withheld or refused on the basis of support for or opposition to particular political party." (11 CFR §14.4 (d) (2) (iii)).[4]

The Regulations related to partisan communications[5] (advocating the election of a candidate or support of a political party) recognizes and reinforces these points. The partisan communication Regulations related to voter registration and GOTV communications states:

> "Disbursements for a voter registration or get-out-the-vote drive conducted under paragraph (c)(4)(i) of this section are not contributions or expenditures if the drive is nonpartisan (see 52 U.S.C. 30118(b)(2)(B)). A drive is nonpartisan if it is conducted so that information and other assistance regarding registering or voting, including transportation and other services offered, is not withheld or refused on the basis of support for or opposition to particular candidates or a particular political party." 11 CFR §114.3 (4)(ii).

Ramifications of Violations: As noted above, these various provisions related to the Activities are categorized as an exemption to the definition of a contribution and expenditure. But for these noted Activities exemptions, if the Activities were to be undertaken and paid for by by a corporation or labor organization, the value of those Activities would constitute a prohibited contribution/expenditure, an excessive contribution and a failure to file FEC disclosure reports if the entity attempted to claim they were independent expenditures (see fn. 3 infra.). Violations would be subject to civil fines negotiated with the FEC, enforced by Federal courts and in the most severe cases, referral to Department of Justice for criminal prosecution.

This memo provides a very brief summary of the applicable law addressing the issue you raise. In my opinion, I believe this type of targeting occurs on a routine basis. However, there are very few FEC enforcement cases pertaining to the issue. It is a very difficult issue to substantiate from a factual perspective and I think therein lies the problem from not only determining that a violation has occurred but even assembling the factual evidence that would support the allegations in a complaint filed with the FEC.

Moving Forward: As a last point, you have requested me to opine on several additional procedural items. This will be a heavily factual driven case. Therefore, getting the supporting factual information to support the legal argument is paramount. I am not in a position to make a determination of how long that process would take.

---

[4] A question may surface as to whether targeting demographic groups (racial, ethnic, financial) would be included in the scope of the prohibitions discussed herein. Though the Regulations do not specifically address that point, the bottom line is that those demographic groups, regardless of the demographic grouping, are being targeted because they are likely supporters of opponents or a candidate or political party and on that basis, that level of targeting would come within the scope of the Regulations general.

[5] 11 CFR §114.3.

A complaint would then be filed with the FEC, submitting the legal brief and factual information, the latter in the form of an affidavit from Dr. Epstein. The information in the complaint will not be made public until the complaint is resolved which will likely be a couple of years. After the case is closed all of the information in the complaint will be placed on the public record, so the parties need to be sensitive to that point as we assemble the factual information to be submitted with the complaint. Once the complaint is filed the Regulations prohibit the parties from publicly commenting on the case. A press conference is permitted to be held PRIOR to the filing of the complaint along with an explanation of the basis for the complaint.

In the event the FEC makes an adverse ruling, there is the opportunity to bring a case in federal district court making the same allegations that the law was violated. Litigation would likely take a couple of years since an appeal of any judgement would need to be anticipated given the parties likely involved.

You have requested a best- and worst-case scenario. Best, is that the FEC makes a finding that there was a violation, and a civil penalty is assessed and if sufficiently egregious, a referral to DOJ for a potential criminal investigation. Worst case is FEC finds no violation and a case brought in the district court is also dismissed or a ruling that there was no violation.

In response to your question of cost, at the initial complaint stage it will depend largely upon assembling the factual data to support the legal arguments. Not knowing how long it would take for me to work with Dr. Epstein on that issue is not a cost I could reasonably assess. The actual research and filing of the legal brief I would anticipate it to be in the $25-50,000 arena. To litigate this case in federal court against a major corporation, the cost would be substantial.

Once the complaint is filed there are generally no other cost that will be incurred. At that point it is up to the FEC to conduct their investigation based upon the documents that were submitted. However, the FEC may request depositions or further clarification of the information proffered in the complaint. Most of the time that process is handled through interrogatories or depositions. Legal services would be required for either of those situations.

As to a likely outcome at the FEC, I would not venture a guess at this time until I understand the factual record that Dr. Epstein will be able to submit. If, and as we proceed through the process, it can be re-evaluated as we work through the process.

There is a substantial amount of legal information and analysis that is raised with this issue, but my hope was to provide you with an initial foothold of the basic issue. Please reach out to me with any follow-up questions or clarification of the issues raised herein.

****

United States Senate

WASHINGTON, DC 20510

November 5, 2020

Sundar Pichai
Chief Executive Officer
Google LLC
1600 Amphitheater Parkway
Mountain View, CA 94043

Dear Mr. Pichai,

On August 6, 2020, Senators Ron Johnson and Mike Lee wrote you regarding your testimony to the House Committee on the Judiciary in which you stated, *"We won't do any work, you know, to politically tilt anything one way or the other. It's against our core values."*[1] In your response to the letter, you reiterated this assertion and further stated, *"Google does not modify any products, including Search, to promote a particular political viewpoint...[we] will not do so for the upcoming 2020 presidential election."*[2]

On November 4, 2020, Senator Johnson spoke with Dr. Robert Epstein to determine if his 2020 Election monitoring of Google's activities had uncovered anything that could be viewed as promoting a particular political viewpoint. Dr. Epstein provided Senator Johnson the following response.

> In our election monitoring project this year, we recruited a politically-diverse group of 733 field agents in Arizona, Florida, and North Carolina. Through their computers, we were able to preserve more than 400,000 ephemeral experiences that tech companies use to shift opinions and votes and that normally are lost forever.
>
> One of our most disturbing findings so far is that between Monday, October 26th (the day our system became fully operational) and Thursday, October 29th, only our liberal field agents received vote reminders on Google's home page. Conservatives did not receive even a single vote reminder. This kind of targeting, if present nationwide, could shift millions of votes, in part because Google's home page is seen 500 million times a day in the U.S.
>
> The good news is that it appears that we got Google to stop this manipulation four days before Election Day. On Thursday, October 29th, I sent materials about the monitoring project to Ebony Bowden, a reporter at the New York Post, who was

---

[1] Letter from U.S. Sen. Ron Johnson and U.S. Sen. Michael Lee, to Mr. Sundar Pichai, Chief Exec. Officer, Google LLC (Aug. 6, 2020).
[2] Letter from Mr. Sundar Pichai, Chief Exec. Officer, Google LLC, to U.S. Sen. Ron Johnson and U.S. Sen. Michael Lee (Aug. 26, 2020).

writing a story about the project. I did so knowing that all nypost.com emails are shared with algorithms and employees at Google.[3]

Late night on the 29th, two notable things happened: First, Ms. Bowden's article, which was about possible large-scale election rigging by Big Tech, was pulled by the Post. Second, Google's targeted messaging stopped completely. From midnight on the 29th to the end of Election Day, all of our field agents have received the vote reminder. Because of the demographics of the people who use Google, this is still a vote manipulation, but it is far more benign than the extreme targeting we detected last week.[4]

Based on Dr. Epstein's response, it would appear your assertion that *"We won't do any work, you know, to politically tilt anything one way or the other"* is not true.

We will be asking Dr. Epstein to provide Congress with the evidence of Google's politically biased activities his monitoring collected during this election cycle. We are writing to provide you another opportunity to conduct a thorough review with your management team to determine the veracity of your previous responses to congressional inquiry regarding this issue, and correct your answers if necessary.

Please provide your response as soon as possible, but no later than 5 p.m. on Thursday, November 12, 2020.

Sincerely,

Ron Johnson
United States Senator

Ted Cruz
United States Senator

Mike Lee
United States Senator

---

[3] *See* Robert Epstein, How Major News Organizations, Universities, and Businesses Surrender Their Privacy to Google, THE DAILY CALLER (Aug. 27, 2018), *available at* https://dailycaller.com/2018/08/27/surrender-privacy-google.

[4] *See* Robert Epstein, How Google Shifts Votes: A 'Go Vote' Reminder Is Not Always What You Think, THE EPOCH TIMES (Jan. 2, 2019), *available at* https://www.theepochtimes.com/another-way-google-manipulates-votes-without-us-knowing-a-go-vote-reminder-is-not-what-you-think-it-is_2754073.html.

PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES
https://SearchEngineManipulationEffect.com

The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections

Robert Epstein[1] and Ronald E. Robertson

American Institute for Behavioral Research and Technology, Vista, CA 92084

Edited by Jacob N. Shapiro, Princeton University, Princeton, NJ, and accepted by the Editorial Board July 8, 2015 (received for review October 16, 2014)

Internet search rankings have a significant impact on consumer choices, mainly because users trust and choose higher-ranked results more than lower-ranked results. Given the apparent power of search rankings, we asked whether they could be manipulated to alter the preferences of undecided voters in democratic elections. Here we report the results of five relevant double-blind, randomized controlled experiments, using a total of 4,556 undecided voters representing diverse demographic characteristics of the voting populations of the United States and India. The fifth experiment is especially notable in that it was conducted with eligible voters throughout India in the midst of India's 2014 Lok Sabha elections just before the final votes were cast. The results of these experiments demonstrate that (*i*) biased search rankings can shift the voting preferences of undecided voters by 20% or more, (*ii*) the shift can be much higher in some demographic groups, and (*iii*) search ranking bias can be masked so that people show no awareness of the manipulation. We call this type of influence, which might be applicable to a variety of attitudes and beliefs, the search engine manipulation effect. Given that many elections are won by small margins, our results suggest that a search engine company has the power to influence the results of a substantial number of elections with impunity. The impact of such manipulations would be especially large in countries dominated by a single search engine company.

search engine manipulation effect | search rankings | Internet influence | voter manipulation | digital bandwagon effect

Recent research has demonstrated that the rankings of search results provided by search engine companies have a dramatic impact on consumer attitudes, preferences, and behavior (1–12); this is presumably why North American companies now spend more than 20 billion US dollars annually on efforts to place results at the top of rankings (13, 14). Studies using eye-tracking technology have shown that people generally scan search engine results in the order in which the results appear and then fixate on the results that rank highest, even when lower-ranked results are more relevant to their search (1–5). Higher-ranked links also draw more clicks, and consequently people spend more time on Web pages associated with higher-ranked search results (1–9). A recent analysis of ~300 million clicks on one search engine found that 91.5% of those clicks were on the first page of search results, with 32.5% on the first result and 17.6% on the second (7). The study also reported that the bottom item on the first page of results drew 140% more clicks than the first item on the second page (7). These phenomena occur apparently because people trust search engine companies to assign higher ranks to the results best suited to their needs (1–4, 11), even though users generally have no idea how results get ranked (15).

Why do search rankings elicit such consistent browsing behavior? Part of the answer lies in the basic design of a search engine results page: the list. For more than a century, research has shown that an item's position on a list has a powerful and persuasive impact on subjects' recollection and evaluation of that item (16–18). Specific order effects, such as primacy and recency, show that the first and last items presented on a list, respectively, are more likely to be recalled than items in the middle (16, 17).

Primacy effects in particular have been shown to have a favorable influence on the formation of attitudes and beliefs (18–20), enhance perceptions of corporate performance (21), improve ratings of items on a survey (22–24), and increase purchasing behavior (25). More troubling, however, is the finding that primacy effects have a significant impact on voting behavior, resulting in more votes for the candidate whose name is listed first on a ballot (26–32). In one recent experimental study, primacy accounted for a 15% gain in votes for the candidate listed first (30). Although primacy effects have been shown to extend to hyperlink clicking behavior in online environments (33–35), no study that we are aware of has yet examined whether the deliberate manipulation of search engine rankings can be leveraged as a form of persuasive technology in elections. Given the power of order effects and the impact that search rankings have on consumer attitudes and behavior, we asked whether the deliberate manipulation of search rankings pertinent to candidates in political elections could alter the attitudes, beliefs, and behavior of undecided voters.

It is already well established that biased media sources such as newspapers (36–38), political polls (39), and television (40) sway voters (41, 42). A 2007 study by DellaVigna and Kaplan found, for example, that whenever the conservative-leaning Fox television network moved into a new market in the United States, conservative votes increased, a phenomenon they labeled the Fox News Effect (40). These researchers estimated that biased coverage by Fox News was sufficient to shift 10,757 votes in Florida during the 2000 US Presidential election: more than enough to flip the deciding state in the election, which was carried by the Republican presidential candidate by only 537 votes. The Fox News Effect was also found to be smaller in television markets that were more competitive.

We believe, however, that the impact of biased search rankings on voter preferences is potentially much greater than the influence of traditional media sources (43), where parties compete in

Significance

We present evidence from five experiments in two countries suggesting the power and robustness of the search engine manipulation effect (SEME). Specifically, we show that (*i*) biased search rankings can shift the voting preferences of undecided voters by 20% or more, (*ii*) the shift can be much higher in some demographic groups, and (*iii*) such rankings can be masked so that people show no awareness of the manipulation. Knowing the proportion of undecided voters in a population who have Internet access, along with the proportion of those voters who can be influenced using SEME, allows one to calculate the win margin below which SEME might be able to determine an election outcome.

Author contributions: R.E. and R.E.R. designed research, performed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. J.N.S. is a guest editor invited by the Editorial Board.

Freely available online through the PNAS open access option.

[1]To whom correspondence should be addressed. Email: re@aibrt.org.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1419828112/-/DCSupplemental.

an open marketplace for voter allegiance. Search rankings are controlled in most countries today by a single company. If, with or without intervention by company employees, the algorithm that ranked election-related information favored one candidate over another, competing candidates would have no way of compensating for the bias. It would be as if Fox News were the only television network in the country. Biased search rankings would, in effect, be an entirely new type of social influence, and it would be occurring on an unprecedented scale. Massive experiments conducted recently by social media giant Facebook have already introduced other unprecedented types of influence made possible by the Internet. Notably, an experiment reported recently suggested that flashing "VOTE" ads to 61 million Facebook users caused more than 340,000 people to vote that day who otherwise would not have done so (44). Zittrain has pointed out that if Facebook executives chose to prompt only those people who favored a particular candidate or party, they could easily flip an election in favor of that candidate, performing a kind of "digital gerrymandering" (45).

We evaluated the potential impact of biased search rankings on voter preferences in a series of experiments with the same general design. Subjects were asked for their opinions and voting preferences both before and after they were allowed to conduct research on candidates using a mock search engine we had created for this purpose. Subjects were randomly assigned to groups in which the search results they were shown were biased in favor of one candidate or another, or, in a control condition, in favor of neither candidate. Would biased search results change the opinions and voting preferences of undecided voters, and, if so, by how much? Would some demographic groups be more vulnerable to such a manipulation? Would people be aware that they were viewing biased rankings? Finally, what impact would familiarity with the candidates have on the manipulation?

Study 1: Three Experiments in San Diego, CA

To determine the potential for voter manipulation using biased search rankings, we initially conducted three laboratory-based experiments in the United States, each using a double-blind control group design with random assignment. For each of the experiments, we recruited 102 eligible voters through newspaper and online advertisements, as well through notices in senior recreation centers, in the San Diego, CA, area.* The advertisements offered USD$25 for each subject's participation, and subjects were prescreened in an attempt to match diverse demographic characteristics of the US voting population (46).

Each of the three experiments used 30 actual search results and corresponding Web pages relating to the 2010 election to determine the prime minister of Australia. The candidates were Tony Abbott and Julia Gillard, and the order in which their names were presented was counterbalanced in all conditions. This election was used to minimize possible preexisting biases by US study participants and thus to try to guarantee that our subjects would be truly "undecided." In each experiment, subjects were randomly assigned to one of three groups: (*i*) rankings favoring Gillard (which means that higher-ranked search results linked to Web pages that portrayed Gillard as the better candidate), (*ii*) rankings favoring Abbott, or (*iii*) rankings favoring neither (Fig. 1 *A*–*C*). The order of these rankings was determined based on ratings of Web pages provided by three independent observers. Neither the subjects nor the research assistants who supervised them knew either the hypothesis of the experiment or the groups to which subjects were assigned.

Initially, subjects read brief biographies of the candidates and rated them on 10-point Likert scales with respect to their overall impression of each candidate, how much they trusted each candidate, and how much they liked each candidate. They were also asked how likely they would be to vote for one candidate or the other on an 11-point scale ranging from −5 to +5, as well as to indicate which of the two candidates they would vote for if the election were held that day.

The subjects then spent up to 15 min gathering more information about the candidates using a mock search engine we had created (called Kadoodle), which gave subjects access to five pages of search results with six results per page. As is usual with search engines, subjects could click on any search result to view the corresponding Web page, or they could click on numbers at the bottom of each results page to view other results pages. The same search results and Web pages were used for all subjects in each experiment; only the order of the search results was varied (Fig. 1). Subjects had the option to end the search whenever they felt they had acquired sufficient information to make a sound decision. At the conclusion of the search, subjects rated the candidates again. When their ratings were complete, subjects were asked (on their computer screens) whether anything about the search rankings they had viewed "bothered" them; they were then given an opportunity to write at length about what, if anything, had bothered them. We did not ask specifically whether the search rankings appeared to be "biased" to avoid false positives typically generated by leading or suggestive questions (47).

Regarding the ethics of our study, our manipulation could have no impact on a past election, and we were also not concerned that it could affect the outcome of future elections, because the number of subjects we recruited was small and, to our knowledge, included no Australian voters. Moreover, our study was designed so that it did not favor any one candidate, so there was no overall bias. The study presented no more than minimal risk to subjects and was approved by the Institutional Review Board (IRB) of the American Institute for Behavioral Research and Technology (AIBRT). Informed consent was obtained from all subjects.

In aggregate for the first three experiments in San Diego, CA, the demographic characteristics of our subjects (mean age, 42.5 y; SD = 18.1 y; range, 18–95 y) did not differ from characteristics of the US voting population by more than the following

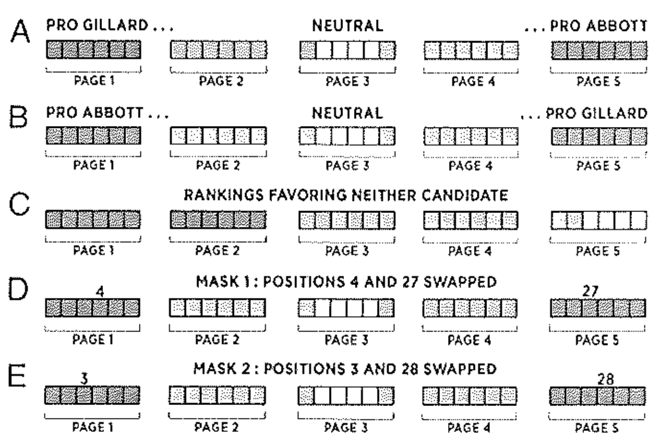

Fig. 1. Search rankings for the three experiments in study 1. (*A*) For subjects in group 1 of experiment 1, 30 search results that linked to 30 corresponding Web pages were ranked in a fixed order that favored candidate Julia Gillard, as follows: those favoring Gillard (from highest to lowest rated pages), then those favoring neither candidate, then those favoring Abbott (from lowest to highest rated pages). (*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 fourth position—favored one candidate. (*E*) For subjects in groups 1 and 2 of experiment 3, a more aggressive mask was used by swapping results that had originally appeared in positions 3 and 28.

*Although all participants claimed to be eligible voters in the prescreening, we later discovered that 6.9% of subjects marked "I don't know" and 5.2% of subjects marked "No" in response to a question asking "If you are not currently registered, are you eligible to register for elections?"

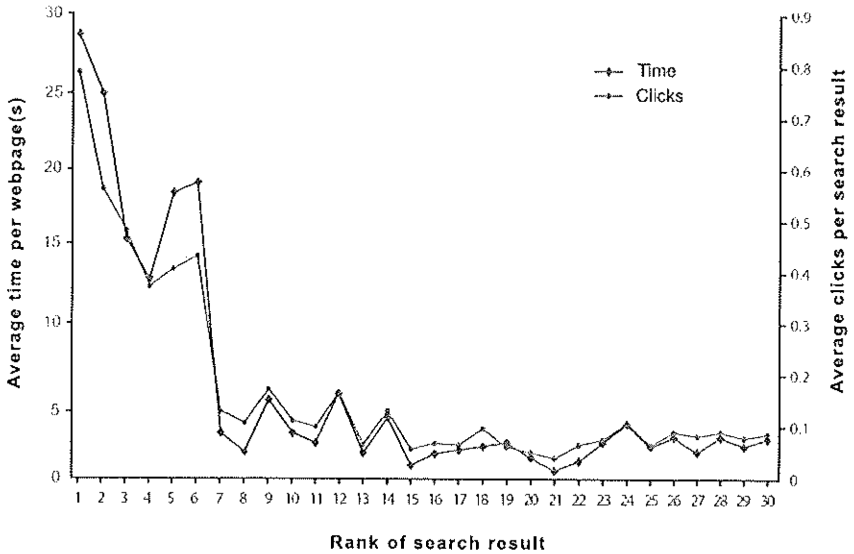

Fig. 2. Clicks on search results and time allocated to Web pages as a function of search result rank, aggregated across the three experiments in study 1. Subjects spent less time on Web pages corresponding to lower-ranked search results (blue curve) and were less likely to click on lower-ranked results (red curve). This pattern is found routinely in studies of Internet search engine use (1–12).

margins: 6.4% within any category of the age or sex measures; 14.1% within any category of the race measure; 18.7% within any category of the income or education measures; and 21.1% within any category of the employment status measure (Table S1). Subjects' political inclinations were fairly balanced, with 20.3% identifying themselves as conservative, 28.8% as moderate, 22.5% as liberal, and 28.4% as indifferent. Political party affiliation, however, was less balanced, with 21.6% identifying as Republican, 19.6% as Independent, 44.8% as Democrat, 6.2% as Libertarian, and 7.8% as other. In aggregate, subjects reported conducting an average of 7.9 searches (SD = 17.5) per day using search engines, and 52.3% reported having conducted searches to learn about political candidates. They also reported having little or no familiarity with the candidates (mean familiarity on a scale of 1–10, 1.4; SD = 0.99). On average, subjects in the first three experiments spent 635.9 s (SD = 307.0) using our mock search engine.

As expected, higher search rankings drew more clicks, and the pattern of clicks for the first three experiments correlated strongly with the pattern found in a recent analysis of ~300 million clicks [$r(13) = 0.90$, $P < 0.001$; Kolmogorov–Smirnov test of differences in distributions: $D = 0.033$, $P = 0.31$; Fig. 2] (7). In addition, subjects spent more time on Web pages associated with higher-ranked results (Fig. 2), as well as substantially more time on earlier search pages (Fig. 3).

In experiment 1, we found no significant differences among the three groups with respect to subjects' ratings of the candidates before Web research (Table S2). Following the Web research, all candidate ratings in the bias groups shifted in the predicted directions compared with candidate ratings in the control group (Table 1).

Before Web research, we found no significant differences among the three groups with respect to the proportions of people who said that they would vote for one candidate or the other if the election were held today (Table 2). Following Web research, significant differences emerged among the three groups for this measure (Table 2), and the number of subjects who said they would vote for the favored candidate in the two bias groups combined increased by 48.4% (95% CI, 30.8–66.0%; McNemar's test, $P < 0.01$).

We define the latter percentage as vote manipulation power (VMP). Thus, before the Web search, if a total of x subjects in the bias groups said they would vote for the target candidate, and if, following the Web search, a total of x' subjects in the bias groups said they would vote for the target candidate, VMP = $(x' - x)/x$. The VMP is, we believe, the key measure that an administrator would want to know if he or she were trying to manipulate an election using SEME.

Using a more sensitive measure than forced binary choice, we also asked subjects to estimate the likelihood, on an 11-point scale from −5 to +5, that they would vote for one candidate or the other if the election were held today. Before Web research, we found no significant differences among the three groups with respect to the likelihood of voting for one candidate or the other [Kruskal–Wallis (K–W) test: $\chi^2(2) = 1.384$, $P = 0.501$]. Following Web research, the likelihood of voting for either candidate in the bias groups diverged from their initial scale values by 3.71 points in the predicted directions [Mann–Whitney (M–W) test: $u = 300.5$, $P < 0.01$]. Notably, 75% of subjects in the bias groups showed no awareness of the manipulation. We counted subjects as showing awareness of the manipulation if (*i*) they had clicked on the box indicating that something bothered them about the rankings and (*ii*) we found specific terms or phrases in their open-ended comments suggesting that they were aware of bias in the rankings (*SI Text*).

In experiment 2, we sought to determine whether the proportion of subjects who were unaware of the manipulation could be increased with voter preferences still shifting in the predicted directions. We accomplished this by masking our manipulation to some extent. Specifically, the search result that had appeared in the fourth position on the first page of the search results favoring Abbott in experiment 1 was swapped with the corresponding search result favoring Gillard (Fig. 1*D*). Before Web research, we found no significant differences among the three

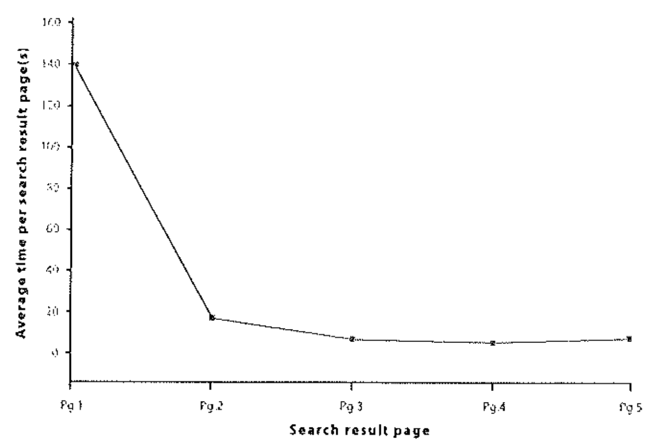

Fig. 3. Amount of time, aggregated across the three experiments in study 1, that subjects spent on each of the five search pages. Subjects spent most of their time on the first search page, a common finding in Internet search engine research (1–12).

Table 1. Postsearch shifts in voting preferences for study 1

			Mean deviation from control (SE)			
Experiment	Candidate	Rating	Gillard bias	u	Abbott bias	u
1	Gillard	Impression	1.44 (0.56)*	761.0	−1.52 (0.56)**	380.5
		Trust	1.26 (0.53)**	779.0	−1.85 (0.48)**	330.5
		Like	0.26 (0.54)	615.5	−1.73 (0.65)**	387.0
	Abbott	Impression	−2.29 (0.73)**	373.0	1.11 (0.72)**	766.5
		Trust	−2.02 (0.63)**	384.0	0.67 (0.76)	679.0
		Like	−1.55 (0.71)	460.5	1.17 (0.64)*	733.0
2	Gillard	Impression	0.97 (0.65)	704.0	−2.38 (0.79)***	325.0
		Trust	0.94 (0.72)	691.5	−2.17 (0.74)**	332.5
		Like	0.55 (0.76)	639.5	−1.82 (0.66)**	378.0
	Abbott	Impression	−1.44 (0.81)*	395.5	1.17 (0.75)*	742.0
		Trust	−0.79 (0.81)	453.5	1.85 (0.72)**	774.5
		Like	−1.44 (0.70)*	429.0	0.64 (0.71)	690.0
3	Gillard	Impression	1.44 (0.73)*	717.5	−0.55 (0.69)	507.5
		Trust	0.47 (0.70)	620.0	−0.23 (0.56)	466.5
		Like	0.44 (0.65)	623.5	−0.41 (0.70)	528.5
	Abbott	Impression	−0.32 (0.70)	534.0	1.26 (0.60)*	750.5
		Trust	−0.73 (0.65)	498.5	1.50 (0.58)**	795.0
		Like	−0.50 (0.61)	496.0	0.88 (0.62)	681.5

*$P < 0.05$, **$P < 0.01$, and ***$P < 0.001$: Mann–Whitney u tests were conducted between the control group and each of the bias groups.

groups with respect to subjects' ratings of the candidates (Table S2). Following the Web research, all candidate ratings in the bias groups shifted in the predicted directions compared with candidate ratings in the control group (Table 1).

Before Web research, we found no significant differences among the three groups with respect to voting proportions (Table 2). Following Web research, significant differences emerged among the three groups for this measure (Table 2), and the VMP was 63.3% (95% CI, 46.1–80.6%; McNemar's test, $P < 0.001$).

For the more sensitive measure (the 11-point scale), we found no significant differences among the three groups with respect to the likelihood of voting for one candidate or the other before Web research [K-W test: $\chi^2(2) = 0.888$, $P = 0.642$]. Following Web research, the likelihood of voting for either candidate in the bias groups diverged from their initial scale values by 4.44 points in the predicted directions (M-W test: $u = 237.5$, $P < 0.001$). In addition, the proportion of people who showed no awareness of the manipulation increased from 75% in experiment 1 to 85% in experiment 2, although the difference between these percentages was not significant ($\chi^2 = 2.264$, $P = 0.07$).

In experiment 3, we sought to further increase the proportion of subjects who were unaware of the manipulation by using a more aggressive mask. Specifically, the search result that had appeared in the third position on the first page of the search results favoring Abbott in experiment 1 was swapped with the corresponding search result favoring Gillard (Fig. 1*E*). This mask is a more aggressive one because higher ranked results are viewed more and taken more seriously by people conducting searches (1–12).

Before Web research, we found no significant differences among the three groups with respect to subjects' ratings of candidates (Table S2). Following the Web research, all candidate ratings in the bias groups shifted in the predicted directions compared with candidate ratings in the control group (Table 1).

Before Web research, we found no significant differences among the three groups with respect to voting proportions (Table 2). Following Web research, significant differences did not emerge among

Table 2. Comparison of voting proportions before and after Web research by group for studies 1 and 2

			Simulated vote before Web research			Simulated vote after Web research			
Study	Experiment	Group	Gillard	Abbott	χ^2	Gillard	Abbott	χ^2	VMP
1	1	1	8	26	5.409	22	12	8.870*	48.4%**
		2	11	23		10	24		
		3	17	17		14	20		
	2	1	16	18	2.197	27	7	14.274***	63.3%***
		2	20	14		12	22		
		3	14	20		22	12		
	3	1	17	17	2.199	22	12	3.845	36.7%*
		2	21	13		15	19		
		3	15	19		15	19		
2	4	1	317	383	1.047	489	211	196.280***	37.1%***
		2	316	384		228	472		
		3	333	367		377	323		

McNemar's test was conducted to assess VMP significance. VMP, percent increase in subjects in the bias groups combined who said that they would vote for the favored candidate.

*$P < 0.05$; **$P < 0.01$; and ***$P < 0.001$: Pearson χ^2 tests were conducted among all three groups.

the three groups for this measure (Table 2); the VMP, however, was 36.7% (95% CI, 19.4–53.9%; McNemar's test, $P < 0.05$).

For the more sensitive measure (the 11-point scale), we found no significant differences among the three groups with respect to the likelihood of voting for one candidate or the other before Web research [K-W test: $\chi^2(2) = 0.624$, $P = 0.732$]. Following Web research, the likelihood of voting for either candidate in the bias groups diverged from their initial scale values by 2.62 points in the predicted directions (M-W test: $u = 297.0$, $P < 0.001$). Notably, in experiment 3, no subjects showed awareness of the rankings bias, and the difference between the proportions of subjects who appeared to be unaware of the manipulations in experiments 1 and 3 was significant ($\chi^2 = 19.429$, $P < 0.001$).

Although the findings from these first three experiments were robust, the use of small samples from one US city limited their generalizability and might even have exaggerated the effect size (48).

Study 2: Large-Scale National Online Replication of Experiment 3

To better assess the generalizability of SEME to the US population at large, we used a diverse national sample of 2,100 individuals[†] from all 50 US states (Table S1), recruited using Amazon's Mechanical Turk (mturk.com), an online subject pool that is now commonly used by behavioral researchers (49, 50). Subjects (mean age, 33.9 y; SD = 11.9 y; range, 18–81 y) were exposed to the same aggressive masking procedure we used in experiment 3 (Fig. 1*E*). Each subject was paid USD$1 for his or her participation.

Regarding ethical concerns, as in study 1, our manipulation could have no impact on a past election, and we were not concerned that it could affect the outcome of future elections. Moreover, our study was designed so that it did not favor any one candidate, so there was no overall bias. The study presented no more than minimal risk to subjects and was approved by AIBRT's IRB. Informed consent was obtained from all subjects.

Subjects' political inclinations were less balanced than those in study 1, with 19.5% of subjects identifying themselves as conservative, 24.2% as moderate, 50.2% as liberal, and 6.3% as indifferent; 16.1% of subjects identified themselves as Republican, 29.9% as Independent, 43.2% as Democrat, 8.0% as Libertarian, and 2.9% as other. Subjects reported having little or no familiarity with the candidates (mean, 1.9; SD = 1.7). As one might expect in a study using only Internet-based subjects, self-reported search engine use was higher in study 2 than in study 1 [mean searches per day, 15.3; SD = 26.3; $t(529.5)^{\ddagger} = 6.9$, $P < 0.001$], and more subjects reported having previously used a search engine to learn about political candidates (86.0%, $\chi^2 = 204.1$, $P < 0.001$). Subjects in study 2 also spent less time using our mock search engine [mean total time, 309.2 s; SD = 278.7; $t(381.9)^{\ddagger} = -17.6$, $P < 0.001$], but patterns of search result clicks and time spent on Web pages were similar to those we found in study 1 [clicks: $r(28) = 0.98$, $P < 0.001$; Web page time: $r(28) = 0.98$, $P < 0.001$] and to those routinely found in other studies (1–12).

Before Web research, we found no significant differences among the three groups with respect to subjects' ratings of the candidates (Table S3). Following the Web research, all candidate ratings in the bias groups shifted in the predicted directions compared with candidate ratings in the control group (Table 3).

Before Web research, we found no significant differences among the three groups with respect to voting proportions (Table 2). Following Web research, significant differences emerged among the three groups for this measure (Table 2), and the VMP was 37.1% (95% CI, 33.5–40.7%; McNemar's test, $P < 0.001$). Using post-stratification and weights obtained from the 2010 US Census (46) and a 2011 study from Gallup (51), which were scaled to size for age, sex, race, and education, the VMP was 36.7% (95% CI, 33.2–40.3%; McNemar's test, $P < 0.001$). When weighted using the same demographics via classical regression poststratification (52) (Table S4), the VMP was 33.5% (95% CI, 30.1–37.0%, McNemar's test, $P < 0.001$).

For the more sensitive measure (the 11-point scale), we found no significant differences among the three groups with respect to the likelihood of voting for one candidate or the other before Web research [K-W test: $\chi^2(2) = 2.790$, $P = 0.248$]. Following Web research, the likelihood of voting for either candidate in the bias groups diverged from their initial scale values by 3.03 points in the predicted directions (M-W test: $u = 1.29 \times 10^5$, $P < 0.001$). As one might expect of a more Internet-fluent sample, the proportion of subjects showing no awareness of the manipulation dropped to 91.4%.

The number of subjects in study 1 was too small to look at demographic differences. In study 2, we found substantial differences in how vulnerable different demographic groups were to SEME. Consistent with previous findings on the moderators of order effects (30–32), for example, we found that subjects reporting a low familiarity with the candidates (familiarity less than 5 on a scale from 1 to 10) were more vulnerable to SEME (VMP = 38.7%; 95% CI, 34.9–42.4%; McNemar's test, $P < 0.001$) than were subjects who reported high familiarity with the candidates (VMP = 19.3%; 95% CI, 9.1–29.5%; McNemar's test, $P < 0.05$), and this difference was significant ($\chi^2 = 8.417$, $P < 0.01$).

We found substantial differences in vulnerability to SEME among a number of different demographic groups (*SI Text*). Although the groups we examined were overlapping and somewhat arbitrary, if one were manipulating an election, information about such differences would have enormous practical value. For example, we found that self-labeled Republicans were more vulnerable to SEME (VMP = 54.4%; 95% CI, 45.2–63.5%; McNemar's test, $P < 0.001$) than were self-labeled Democrats (VMP = 37.7%; 95% CI, 32.3–43.1%; McNemar's test, $P < 0.001$) and that self-labeled divorcees were more vulnerable (VMP = 46.7%; 95% CI, 32.1–61.2%; McNemar's test, $P < 0.001$) than were self-labeled married subjects (VMP = 32.4%; 95% CI, 26.8–38.1%; McNemar's test, $P < 0.001$). Among the most vulnerable groups we identified were Moderate Republicans (VMP = 80.0%; 95% CI, 62.5–97.5%; McNemar's test, $P < 0.001$), whereas among the least vulnerable groups were people who reported a household income of $40,000 to $49,999 (VMP = 22.5%; 95% CI, 13.8–31.1%; McNemar's test, $P < 0.001$).

Notably, awareness of the manipulation not only did not nullify the effect, it seemed to enhance it, perhaps because people trust search order so much that awareness of the bias serves to confirm the superiority of the favored candidate. The VMP for people who showed no awareness of the biased search rankings ($n = 1,280$) was 36.3% (95% CI, 32.6–40.1%; McNemar's test, $P < 0.001$), whereas the VMP for people who showed awareness of the bias ($n = 120$) was 45.0% (95% CI, 32.4–57.6%; McNemar's test, $P < 0.001$).

Having now replicated the effect with a large and diverse sample of US subjects, we were concerned about the weaknesses associated with testing subjects on a somewhat abstract election (the election in Australia) that had taken place years before and in which subjects were unfamiliar with the candidates. In real elections, people are familiar with the candidates and are influenced, sometimes on a daily basis, by aggressive campaigning. Presumably, either of these two factors—familiarity and outside influence—could potentially minimize or negate the influence of biased search rankings on voter preferences. We therefore asked if SEME could be replicated with a large and diverse sample of real voters in the midst of a real election campaign.

Study 3: SEME Evaluated During the 2014 Lok Sabha Elections in India

In our fifth experiment, we sought to manipulate the voting preferences of undecided eligible voters in India during the 2014 national Lok Sabha elections there. This election was the largest democratic election in history, with more than 800 million eligible voters and more than 430 million votes ultimately cast. We accomplished this by randomly assigning undecided English-speaking

[†]As in study 1, although all participants claimed to be eligible voters in the prescreening, we later discovered that 4.7% of subjects marked "I don't know" and 2.6% of subjects marked "No" in response to a question asking "If you are not currently registered, are you eligible to register for elections?"

[‡]Degrees of freedom adjusted for significant inequality of variances (Welch's t test).

Table 3. Postsearch shifts in voting preferences for study 2

Candidate	Rating	Mean deviation from control (SE)			
		Gillard bias	u	Abbott bias	u
Gillard	Impression	0.65 (0.10)***	288,299.5	−1.25 (0.12)***	168,203.5
	Trust	0.61 (0.10)***	283,491.0	−1.21 (0.11)***	167,658.5
	Like	0.50 (0.10)***	279,967.0	−1.25 (0.11)***	166,544.0
Abbott	Impression	−0.96 (0.13)***	189,290.5	1.35 (0.12)***	326,067.0
	Trust	−1.09 (0.14)***	183,993.0	1.31 (0.12)***	318,740.5
	Like	−0.85 (0.13)***	195,088.5	0.94 (0.11)***	302,318.0

***$P < 0.001$: Mann–Whitney u tests were conducted between the control group and each of the bias groups.

voters throughout India who had not yet voted (recruited through print advertisements, online advertisements, and online subject pools) to one of three groups in which search rankings favored either Rahul Gandhi, Arvind Kejriwal, or Narendra Modi, the three major candidates in the election.[§]

Subjects were incentivized to participate in the study either with payments between USD$1 and USD$4 or with the promise that a donation of approximately USD$1.50 would be made to a prominent Indian charity that provides free lunches for Indian children. (At the close of the study, a donation of USD$1,457 was made to the Akshaya Patra Foundation.)

Regarding ethical concerns, because we recruited only a small number of subjects relative to the size of the Indian voting population, we were not concerned that our manipulation could affect the election's outcome. Moreover, our study was designed so that it did not favor any one candidate, so there was no overall bias. The study presented no more than minimal risk to subjects and was approved by AIBRT's IRB. Informed consent was obtained from all subjects.

The subjects ($n = 2,150$) were demographically diverse (Table S5), residing in 27 of 35 Indian states and union territories, and political leanings varied as follows: 13.3% identified themselves as politically right (conservative), 43.8% as center (moderate), 26.0% as left (liberal), and 16.9% as indifferent. In contrast to studies 1 and 2, subjects reported high familiarity with the political candidates (mean familiarity Gandhi, 7.9; SD = 2.5; mean familiarity Kejriwal, 7.7; SD = 2.5; mean familiarity Modi, 8.5; SD = 2.1). The full dataset for all five experiments is accessible at Dataset S1.

Subjects reported more frequent search engine use compared with subjects in studies 1 or 2 (mean searches per day, 15.7; SD = 30.1), and 71.7% of subjects reported that they had previously used a search engine to learn about political candidates. Subjects also spent less time using our mock search engine (mean total time, 277.4 s; SD = 368.3) than did subjects in studies 1 or 2. The patterns of search result clicks and time spent on Web pages in our mock search engine was similar to the patterns we found in study 1 [clicks, $r(28) = 0.96$; $P < 0.001$; Web page time, $r(28) = 0.91$; $P < 0.001$] and study 2 [clicks, $r(28) = 0.96$; $P < 0.001$; Web page time, $r(28) = 0.92$; $P < 0.001$].

Before Web research, we found one significant difference among the three groups for a rating pertaining to Kejriwal, but none for Gandhi or Modi (Table S6). Following the Web research, most of the subjects' ratings of the candidates shifted in the predicted directions (Table 4).

Before Web research, we found no significant differences among the three groups with respect to voting proportions (Table 5). Following Web research, significant differences emerged among the three groups for this measure (Table 5), and the VMP was 10.6% (95% CI, 8.3–12.8%; McNemar's test, $P < 0.001$). Using poststratification and weights obtained from the 2011 India Census data on literate Indians (53)—scaled to size for age, sex, and location (grouped into state or union territory)—the VMP was 9.4% (95% CI, 8.2–10.6%; McNemar's test, $P < 0.001$). When weighted using the same demographics via classical regression poststratification (Table S7), the VMP was 9.5% (95% CI, 8.3–10.7%; McNemar's test, $P < 0.001$).

To obtain a more sensitive measure of voting preference in study 3, we asked subjects to estimate the likelihood, on three separate 11-point scales from −5 to +5, that they would vote for each of the candidates if the election were held today. Before Web research, we found no significant differences among the three groups with respect to the likelihood of voting for any of the candidates (Table S6). Following Web research, significant differences emerged among the three groups with respect to the likelihood of voting for Rahul Gandhi and Arvind Kejriwal but not Narendra Modi (Table S6), and all likelihoods shifted in the predicted directions (Table 4). The proportion of subjects showing no awareness of the manipulation in experiment 5 was 99.5%.

In study 3, as in study 2, we found substantial differences in how vulnerable different demographic groups were to SEME (*SI Text*). Consistent with the findings of study 2 and previous findings on the moderators of order effects (30–32), for example, we found that subjects reporting a low familiarity with the candidates (familiarity less than 5 on a scale from 1 to 10) were more vulnerable to SEME (VMP = 13.7%; 95% CI, 4.3–23.2%; McNemar's test, $P = 0.17$) than were subjects who reported high familiarity with the candidates (VMP = 10.3%; 95% CI, 8.0–12.6%; McNemar's test, $P < 0.001$), although this difference was not significant ($\chi^2 = 0.575$, $P = 0.45$).

As in study 2, although the demographic groups we examined were overlapping and somewhat arbitrary, if one was manipulating an election, information about such differences would have enormous practical value. For example, we found that subjects between ages 18 and 24 were less vulnerable to SEME (VMP = 8.9%; 95% CI, 5.0–12.8%; McNemar's test, $P < 0.05$) than were subjects between ages 45 and 64 (VMP = 18.9%; 95% CI, 6.3–31.5%; McNemar's test, $P = 0.10$) and that self-labeled Christians were more vulnerable (VMP = 30.7%; 95% CI, 20.2–41.1%; McNemar's test, $P < 0.001$) than self-labeled Hindus (VMP = 8.7%; 95% CI, 6.3–11.1%; McNemar's test, $P < 0.001$). Among the most vulnerable groups we identified were unemployed males from Kerala (VMP = 72.7%; 95% CI, 46.4–99.0%; McNemar's test, $P < 0.05$), whereas among the least vulnerable groups were female conservatives (VMP = −11.8%; 95% CI, −29.0%–5.5%; McNemar's test, $P = 0.62$).

A negative VMP might suggest oppositional attitudes or an underdog effect for that group (54). No negative VMPs were found in the demographic groups examined in study 2, but it is understandable that they would be found in an election in which people are highly familiar with the candidates (study 3). As a practical matter, where a search engine company has the ability to send people customized rankings and where biased search rankings are likely to produce an oppositional response with certain voters, such rankings would probably not be sent to them. Eliminating the 2.6% of our sample ($n = 56$) with oppositional responses, the overall VMP in this experiment increases from 10.6% to 19.8% (95% CI, 16.8–22.8%; $n = 2,094$; McNemar's test: $P < 0.001$).

As we found in study 2, awareness of the manipulation appeared to enhance the effect rather than nullify it. The VMP for people

[§]English is one of India's two official languages, the other being Hindi.

Table 4. Postsearch shifts in voting preferences for study 3

Candidate	Rating	χ^2	Mean (SE) Gandhi bias	Kejriwal bias	Modi bias
Gandhi	Impression	3.61	−0.16 (0.06)	−0.21 (0.06)	−0.30 (0.06)
	Trust	21.19***	0.14 (0.06)	−0.04 (0.07)	−0.20 (0.06)
	Like	12.99**	−0.09 (0.07)	−0.17 (0.06)	−0.34 (0.06)
	Voting likelihood	10.79**	0.16 (0.07)	−0.04 (0.07)	−0.18 (0.07)
Kejriwal	Impression	17.75***	−0.30 (0.06)	−0.11 (0.06)	−0.39 (0.05)
	Trust	26.69***	−0.17 (0.07)	0.15 (0.06)	−0.16 (0.06)
	Like	24.74***	−0.31 (0.06)	0.05 (0.06)	−0.23 (0.06)
	Voting likelihood	13.22**	−0.03 (0.06)	0.17 (0.07)	−0.12 (0.06)
Modi	Impression	24.98***	−0.22 (0.06)	−0.21 (0.06)	0.12 (0.05)
	Trust	18.78***	−0.04 (0.06)	−0.10 (0.06)	0.23 (0.06)
	Like	16.89***	−0.16 (0.05)	−0.09 (0.06)	0.19 (0.06)
	Voting likelihood	31.07***	−0.07 (0.07)	−0.10 (0.06)	0.33 (0.06)

**$P < 0.01$ and ***$P < 0.001$: for each rating, a Kruskal–Wallis χ^2 test was used to assess significance of group differences.

who showed no awareness of the biased search rankings (n = 2,140) was 10.5% (95% CI, 8.3–12.7%; McNemar's test, $P < 0.001$), whereas the VMP for people who showed awareness of the bias (n = 10) was 33.3%.

The rankings and Web pages we used in study 3 were selected by the investigators based on our limited understanding of Indian politics and perspectives. To optimize the rankings, midway through the election process we hired a native consultant who was familiar with the issues and perspectives pertinent to undecided voters in the 2014 Lok Sabha Election. Based on the recommendations of the consultant, we made slight changes to our rankings on 30 April, 2014. In the preoptimized rankings group (n = 1,259), the VMP was 9.5% (95% CI, 6.8–12.2%; McNemar's test, $P < 0.001$); in the postoptimized rankings group (n = 891), the VMP increased to 12.3% (95% CI, 8.5–16.1%; McNemar's test, $P < 0.001$). Eliminating the 3.1% of the subjects in the postoptimization sample with oppositional responses (n = 28), the VMP increased to 24.5% (95% CI, 19.3–29.8%; n = 863).

Discussion

Elections are often won by small vote margins. Fifty percent of US presidential elections were won by vote margins under 7.6%, and 25% of US senatorial elections in 2012 were won by vote margins under 6.0% (55, 56). In close elections, undecided voters can make all of the difference, which is why enormous resources are often focused on those voters in the days before the election (57, 58). Because search rankings biased toward one candidate can apparently sway the voting preferences of undecided voters without their awareness and, at least under some circumstances, without any possible competition from opposing candidates, SEME appears to be an especially powerful tool for manipulating elections. The Australian election used in studies 1 and 2 was won by a margin of only 0.24% and perhaps could easily have been turned by such a manipulation. The Fox News Effect, which is small compared with SEME, is believed to have shifted between 0.4% and 0.7% of votes to conservative candidates: more than enough, according to the researchers, to have had a "decisive" effect on a number of close elections in 2000 (40).

Political scientists have identified two of the most common methods political candidates use to try to win elections. The core voter model describes a strategy in which resources are devoted to mobilizing supporters to vote (59). As noted earlier, Zittrain recently pointed out that a company such as Facebook could mobilize core voters to vote on election day by sending "get-out-and-vote" messages en masse to supporters of only one candidate. Such a manipulation could be used undetectably to flip an election in what might be considered a sort of digital gerrymandering (44, 45). In contrast, the swing voter model describes a strategy in which candidates target their resources toward persuasion—attempting to change the voting preferences of undecided voters (60). SEME is an ideal method for influencing such voters.

Although relatively few voters have actively sought political information about candidates in the past (61), the ease of obtaining information over the Internet appears to be changing that: 73% of online adults used the Internet for campaign-related purposes during the 2010 US midterm elections (61), and 55% of all registered voters went online to watch videos related to the 2012 US election campaign (62). Moreover, 84% of registered voters in the United States were Internet users in 2012 (62). In our nationwide study in the United States (study 2), 86.0% of our subjects reported having used search engines to get information about candidates. Meanwhile, the number of people worldwide with Internet access is increasing rapidly, predicted to increase to nearly 4 billion by 2018 (63). By 2018, Internet access in India is expected to rise from the 213 million users who had access in 2013 to 526 million (63). Worldwide, it is reasonable to conjecture that both proportions will increase substantially in future years; that is, more people will have Internet access, and more people will obtain information about candidates from the Internet. In the context of the experiments we have presented, this suggests that whatever the effect sizes we have observed now, they will likely be larger in the future.

Table 5. Comparison of voting proportions before and after Web research for study 3

Group	Simulated vote before Web research: Gandhi	Kejriwal	Modi	χ^2	Simulated vote after Web research: Gandhi	Kejriwal	Modi	χ^2	VMP
1	115	164	430	3.070	144	152	413	16.935**	10.6%***
2	112	183	393		113	199	376		
3	127	196	430		117	174	462		

McNemar's test was conducted to assess VMP significance. VMP, percent increase in subjects in the bias groups combined who said that they would vote for the favored candidate.
**$P < 0.01$; and ***$P < 0.001$: Pearson χ^2 tests were conducted among all three groups.

The power of SEME to affect elections in a two-person race can be roughly estimated by making a small number of fairly conservative assumptions. Where i is the proportion of voters with Internet access, u is the proportion of those voters who are undecided, and VMP, as noted above, is the proportion of those undecided voters who can be swayed by SEME, W—the maximum win margin controllable by SEME—can be estimated by the following formula: $W = i*u*VMP$.

In a three-person race, W will vary between 75% and 100% of its value in a two-person race, depending on how the votes are distributed between the two losing candidates. (Derivations of formulas in the two-candidate and three-candidate cases are available in *SI Text*.) In both cases, the size of the population is irrelevant.

Knowing the values for i and u for a given election, along with the projected win margin, the minimum VMP needed to put one candidate ahead can be calculated (Table S8). In theory, continuous online polling would allow search rankings to be optimized continuously to increase the value of VMP until, in some instances, it could conceivably guarantee an election's outcome, much as "conversion" and "click-through" rates are now optimized continuously in Internet marketing (64).

For example, if (*i*) 80% of eligible voters had Internet access, (*ii*) 10% of those individuals were undecided at some point, and (*iii*) SEME could be used to increase the number of people in the undecided group who were inclined to vote for the target candidate by 25%, that would be enough to control the outcome of an election in which the expected win margin was as high as 2%. If SEME were applied strategically and repeatedly over a period of weeks or months to increase the VMP, and if, in some locales and situations, i and u were larger than in the example given, the controllable win margin would be larger. That possibility notwithstanding, because nearly 25% of national elections worldwide are typically won by margins under 3%,[¶] SEME could conceivably impact a substantial number of elections today even with fairly low values of i, u, and VMP.

Given our procedures, however, we cannot rule out the possibility that SEME produces only a transient effect, which would limit its value in election manipulation. Laboratory manipulations of preferences and attitudes often impact subjects for only a short time, sometimes just hours (65). That said, if search rankings were being manipulated with the intent of altering the outcome of a real election, people would presumably be exposed to biased rankings repeatedly over a period of weeks or months. We produced substantial changes not only in voting preferences but in multiple ratings of attitudes toward candidates given just one exposure to search rankings linking to Web pages favoring one candidate, with average search times in the 277- to 635-s range. Given hundreds or thousands of exposures of this sort, we speculate not only that the resulting attitudes and preferences would be stable, but that they would become stronger over time, much as brand preferences become stronger when advertisements are presented repeatedly (66).

Our results also suggest that it is a relatively simple matter to mask the bias in search rankings so that it is undetectable to virtually every user. In experiment 3, using only a simple mask, none of our subjects appeared to be aware that they were seeing biased rankings, and in our India study, only 0.5% of our subjects appeared to notice the bias. When people are subjected to forms of influence they can identify—in campaigns, that means speeches, billboards, television commercials, and so on—they can defend themselves fairly easily if they have opposing views. Invisible sources of influence can be harder to defend against (67–69), and for people who are impressionable, invisible sources of influence not only persuade, they also leave people feeling that they made up their own minds—that no external force was applied (70, 71). Influence is sometimes undetectable because key stimuli act subliminally (72–74), but search results and Web pages are easy to perceive; it is the pattern of rankings that people cannot see. This invisibility makes SEME especially dangerous as a means of control, not just of voting behavior but perhaps of a wide variety of attitudes, beliefs, and behavior. Ironically, and consistent with the findings of other researchers, we found that even those subjects who showed awareness of the biased rankings were still impacted by them in the predicted directions (75).

One weakness in our studies was the manner in which we chose to determine whether subjects were aware of bias in the search rankings. As noted, to not generate false-positive responses, we avoided asking leading questions that referred specifically to bias; rather, we asked a rather vague question about whether anything had bothered subjects about the search rankings, and we then gave subjects an opportunity to type out the details of their concerns. In so doing, we probably underestimated the number of detections (47), and this is a matter that should be studied further. That said, because people who showed awareness of the bias were still vulnerable to our manipulation, people who use SEME to manipulate real elections might not be concerned about detection, except, perhaps, by regulators.

Could regulators in fact detect SEME? Theoretically, by rating pages and monitoring search rankings on an ongoing basis, search ranking bias related to elections might be possible to identify and track; as a practical matter, however, we believe that biased rankings would be impossible or nearly impossible for regulators to detect. The results of studies 2 and 3 suggest that vulnerability to SEME can vary dramatically from one demographic group to another. It follows that if one were using biased search rankings to manipulate a real election, one would focus on the most vulnerable demographic groups. Indeed, if one had access to detailed online profiles of millions of individuals, which search engine companies do (76–78), one would presumably be able to identify those voters who appeared to be undecided and impressionable and focus one's efforts on those individuals only—a strategy that has long been standard in political campaigns (79–84) and continues to remain important today (85). With search engine companies becoming increasingly adept at sending users customized search rankings (76–78, 86–88), it seems likely that only customized rankings would be used to influence elections, thus making it difficult or impossible for regulators to detect a manipulation. Rankings that appear to be unbiased on the regulators' screens might be highly biased on the screens of select individuals.

Even if a statistical analysis did show that rankings consistently favored one candidate over another, those rankings could always be attributed to algorithm-guided dynamics driven by market forces—so-called "organic" forces (89)—rather than by deliberate manipulation by search engine company employees. This possibility suggests yet another potential danger of SEME. What if election-related search results are indeed being left to the vagaries of market forces? Do such forces end up pushing some candidates to the top of search rankings? If so, it seems likely that those high rankings are cultivating additional supporters for those candidates in a kind of digital bandwagon effect. In other words, for several years now and with greater impact each year (as more people get election-related information through the Internet), SEME has perhaps already been affecting the outcomes of close elections. To put this another way, without human intention or direction, algorithms have perhaps been having a say in selecting our leaders.

Because search rankings are based, at least in part, on the popularity of Web sites (90), it is likely that voter preferences impact those rankings to some extent. Given our findings that search rankings can in turn affect voter preferences, these phenomena might interact synergistically, causing a substantial increase in support for one candidate at some point even when the effects of the individual phenomena are small.[||]

Our studies produced a wide range of VMPs. In a real election, what proportion of undecided voters could actually be

[¶]Some of the data applied in the analysis in this publication are based on material from the "European Election Database." The data are collected from original sources and prepared and made available by the Norwegian Social Science Data Services (NSD). NSD is not responsible for the analyses/interpretation of the data presented here.

[||]A mathematical model we developed—highly conjectural, we admit, and at this point unverifiable—shows the possible dynamics of such synergy (Fig. S1).

shifted using SEME? Our first two studies, which relied on a campaign and candidates that were unfamiliar to our subjects, produced overall VMPs in the range 36.7–63.3%, with demographic shifts occurring with VMPs as high as 80.0%. Our third study, with real voters in the midst of a real election, produced, overall, a lower VMP: just 10.6%, with optimizing our rankings raising the VMP to 12.3% and with the elimination of a small number of oppositional subjects raising the VMP to 24.5%, which is the value we would presumably have found if our search rankings had been optimized from the start and if we had advance knowledge about oppositional groups. In the third study, VMPs in some demographic groups were as high as 72.7%. If a search engine company optimized rankings continuously and sent customized rankings only to vulnerable undecided voters, there is no telling how high the VMP could be pushed, but it would almost certainly be higher than our modest efforts could achieve. Our investigation suggests that with optimized, targeted rankings, a VMP of at least 20% should be relatively easy to achieve in real elections. Even if only 60% of a population had Internet access and only 10% of voters were undecided, that would still allow control of elections with win margins up to 1.2%—five times greater than the win margin in the 2010 race between Gillard and Abbott in Australia.

Conclusions

Given that search engine companies are currently unregulated, our results could be viewed as a cause for concern, suggesting that such companies could affect—and perhaps are already affecting—the outcomes of close elections worldwide. Restricting search ranking manipulations to voters who have been identified as undecided while also donating money to favored candidates would be an especially subtle, effective, and efficient way of wielding influence.

Although voters are subjected to a wide variety of influences during political campaigns, we believe that the manipulation of search rankings might exert a disproportionately large influence over voters for four reasons:

First, as we noted, the process by which search rankings affect voter preferences might interact synergistically with the process by which voter preferences affect search rankings, thus creating a sort of digital bandwagon effect that magnifies the potential impact of even minor search ranking manipulations.

Second, campaign influence is usually explicit, but search ranking manipulations are not. Such manipulations are difficult to detect, and most people are relatively powerless when trying to resist sources of influence they cannot see (66–68). Of greater concern in the present context, when people are unaware they are being manipulated, they tend to believe they have adopted their new thinking voluntarily (69, 70).

Third, candidates normally have equal access to voters, but this need not be the case with search engine manipulations. Because the majority of people in most democracies use a search engine provided by just one company, if that company chose to manipulate rankings to favor particular candidates or parties, opponents would have no way to counteract those manipulations. Perhaps worse still, if that company left election-related search rankings to market forces, the search algorithm itself might determine the outcomes of many close elections.

Finally, with the attention of voters shifting rapidly toward the Internet and away from traditional sources of information (12, 61, 62), the potential impact of search engine rankings on voter preferences will inevitably grow over time, as will the influence of people who have the power to control such rankings.

We conjecture, therefore, that unregulated election-related search rankings could pose a significant threat to the democratic system of government.

Materials and Methods

We used 102 subjects in each of experiments 1–3 to give us an equal number of subjects in all three groups and both counterbalancing conditions of the experiments.

Nonparametric statistical tests such as the Mann–Whitney u and the Kruskal–Wallis H are used throughout the present report because Likert scale scores, which were used in each of the studies, are ordinal.

In study 3, the procedure was identical to that of studies 1 and 2; only the Web pages and search results were different: that is, Web pages and search results were pertinent to the three leading candidates in the 2014 Lok Sabha general elections. The questions we asked subjects were also adjusted for a three-person race.

ACKNOWLEDGMENTS. We thank J. Arnett, E. Clemons, E. Fantino, S. Glenn, M. Hovell, E. Key, E. Loftus, C. McKenzie, B. Meredith, N. Metaxas, D. Moriarty, D. Peel, M. Runco, S. Stolarz-Fantino, and J. Wixted for comments; K. Robertson for image editing; V. Sharan for advice on optimizing search rankings in the India study; K. Duncan and F. Tran for assistance with data analysis; J. Hagan for technical assistance; and S. Palacios and K. Huynh for assistance in conducting the experiments in study 1. This work was supported by the American Institute for Behavioral Research and Technology, a nonpartisan, nonprofit organization.

1. Joachims T, et al. (2007) Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *Association for Computing Machinery Transactions on Information Systems* 25(2):7.
2. Pan B, et al. (2007) In Google we trust: User's decisions on rank, position, and relevance. *J Comput Mediat Commun* 12(3):801–823.
3. Guan Z, Cutrell E (2007) An eye tracking study of the effect of target rank on web search. *Proceedings of the Special Interest Group for Computer-Human Interaction Conference on Human Factors in Computing Systems* (ACM, New York), pp 417–420.
4. Lorigo L, et al. (2008) Eye tracking and online search: Lessons learned and challenges ahead. *J Am Soc Inf Sci Technol* 59(7):1041–1052.
5. Granka LA, Joachim T, Gay G (2004) Eye-tracking analysis of user behavior in www search. *Proceedings of the 27th Annual International Association for Computing Machinery Special Interest Group on Information Retrieval Conference on Research and Development in Information Retrieval* (ACM, New York, NY), pp 478–479.
6. Agichtein E, Brill E, Dumais S, Ragno R (2006) Learning user interaction models for predicting web search result preferences. *Proceedings of the 29th Annual International Association for Computing Machinery Special Interest Group on Information Retrieval Conference on Research and Development in Information Retrieval* (ACM, New York, NY), pp 3–10.
7. Chitika (2013) The value of Google result positioning. Available at perma.cc/7AGC-HTDH. Accessed June 30, 2015.
8. Optify (2011) The changing face of SERPS: Organic CTR. Available at perma.cc/KZ5X-78TL. Accessed June 30, 2015.
9. Jansen BJ, Spink A, Saracevic T (2000) Real life, real users, and real needs: A study and analysis of user queries on the web. *Inf Process Manage* 36(2):207–227.
10. Spink A, Wolfram D, Jansen BJ, Saracevic T (2001) Searching the web: The public and their queries. *J Am Soc Inf Sci Technol* 53(3):226–234.
11. Silverstein C, Marais H, Henzinger M, Moricz A (1999) Analysis of a very large web search engine query log. *Association for Computing Machinery Special Interest Group on Information Retrieval Forum* 33(1):6–12.
12. Purcell K, Brenner J, Rainie L (2012) Pew Research Center's Internet & American Life project: Search engine use 2012. Available at perma.cc/NF6C-JCPW. Accessed June 30, 2015.
13. Learmonth M (2010) What big brands are spending on Google. Available at perma.cc/5L3B-5PTX. Accessed June 30, 2015.
14. Econsultancy (2012) SEMPO state of search marketing report 2012. Available at perma.cc/A3KH-6QVB. Accessed June 30, 2015.
15. Gerhart S (2004) Do Web search engines suppress controversy? *First Monday* 9(1):1111.
16. Ebbinghaus H (1913) *Memory: A Contribution to Experimental Psychology (No. 3)* (Teachers College, Columbia Univ, New York).
17. Murdock BB (1962) The serial position effect of free recall. *J Exp Psychol* 64(5):482–488.
18. Asch SE (1946) Forming impressions of personality. *J Abnorm Psychol* 41:258–290.
19. Lund FH (1925) The psychology of belief. *J Abnorm Soc Psych* 20(1):63–81.
20. Hovland CI (1957) *The Order of Presentation in Persuasion* (Yale Univ Press, New Haven, CT).
21. Baird JE, Zelin RC (2000) The effects of information ordering on investor perceptions: An experiment utilizing presidents' letters. *J Financial Strategic Decisions* 13(3):71–80.
22. Tourangeau R, Couper MP, Conrad FG (2013) "Up means good" The effect of screen position on evaluative ratings in web surveys. *Public Opin Q* 77(1, Suppl 1):69–88.
23. Krosnick JA, Alwin DF (1987) An evaluation of a cognitive theory of response order effects in survey measurement. *Pub Op Quarterly* 51(2):201–219.
24. Stern MJ, Dillman DA, Smyth JD (2007) Visual design, order effects, and respondent characteristics in a self-administered survey. *Surv Res Methods* 1(3):121–138.
25. Miller JE (1980) *Menu Pricing and Strategy* (CBI Publishing, Boston).
26. Ho DE, Imai K (2008) Estimating causal effects of ballot order from a randomized natural experiment the California alphabet lottery, 1978–2002. *Pub Op Quarterly* 72(2):216–240.
27. Chen E, Simonovits G, Krosnick JA, Pasek J (2014) The impact of candidate name order on election outcomes in North Dakota. *Elect Stud* 35:115–122.

28. Krosnick JA, Miller JM, Tichy MP (2004) An unrecognized need for ballot reform. *Rethinking the Vote*, eds Crigler A, Just M, McCaffery E (Oxford Univ Press, New York), pp 51–74.
29. Koppell JG, Steen JA (2004) The effects of ballot position on election outcomes. *J Polit* 66(1):267–281.
30. Kim N, Krosnick JA, Casasanto D (2014) Moderators of candidate name-order effects in elections: An experiment. *Polit Psychol*, 10.1111/pops.12178.
31. Miller JM, Krosnick JA (1998) The impact of candidate name order on election outcomes. *Pub Op Quarterly* 62(3):291–330.
32. Pasek J, et al. (2014) Prevalence and moderators of the candidate name-order effect evidence from statewide general elections in California. *Pub Op Quarterly* 78(2):416–439.
33. Murphy J, Hofacker C, Mizerski R (2006) Primacy and recency effects on clicking behavior. *J Comput Mediat Commun* 11(2):522–535.
34. Ansari A, Mela CF (2003) E-customization. *J Mark Res* 40(2):131–145.
35. Drèze X, Zufryden F (2004) Measurement of online visibility and its impact on Internet traffic. *J Interact Market* 18:20–37.
36. Chiang CF, Knight BG (2011) Media bias and influence: Evidence from newspaper endorsements. *Rev Econ Stud* 78(3):795–820.
37. Druckman JN, Parkin M (2005) The impact of media bias: How editorial slant affects voters. *J Polit* 67(4):1030–1049.
38. Gerber A, Karlan DS, Bergan D (2009) Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions. *Am Econ J Appl Econ* 1(2):35–52.
39. Morwitz VG, Pluzinski C (1996) Do polls reflect opinions or do opinions reflect polls? The impact of political polling on voters' expectation, preferences, and behavior. *J Consum Res* 23(1):53–67.
40. DellaVigna S, Kaplan E (2007) The Fox News effect: Media bias and voting. *Q J Econ* 122(3):1187–1234.
41. Lakoff G (1987) *Women, Fire and Dangerous Things* (Univ of Chicago Press, Chicago).
42. DeMarzo PM, Vayanos D, Zwiebel J (2003) Persuasion bias, social influence, and unidimensional opinions. *Q J Econ* 118(3):909–968.
43. Fogg BJ (2002) *Persuasive Technology: Using Computers to Change What We Think and Do* (Morgan Kaufmann, San Francisco).
44. Bond RM, et al. (2012) A 61-million-person experiment in social influence and political mobilization. *Nature* 489(7415):295–298.
45. Zittrain J (2014) Engineering an election. Digital gerrymandering poses a threat to democracy. *Harvard Law Review Forum* 127(8):335–341.
46. US Census Bureau (2011) Voting and registration in the election of November 2010. Available at perma.cc/DU9N-C6XT. Accessed June 30, 2015.
47. Loftus EF (1996) *Eyewitness Testimony* (Harvard Univ Press, Boston).
48. Kühberger A, Fritz A, Scherndl T (2014) Publication bias in psychology: A diagnosis based on the correlation between effect size and sample size. *PLoS One* 9(9):e105825.
49. Paolacci G, Chandler J (2014) Inside the Turk: Understanding mechanical Turk as a participant pool. *Curr Dir Psychol Sci* 23(3):184–188.
50. Berinsky AJ, Huber GA, Lenz GS (2012) Evaluating online labor markets for experimental research: Amazon.com's mechanical turk. *Polit Anal* 20(3):351–368.
51. Gates GJ (2011) How many people are lesbian, gay, bisexual and transgender? Available at perma.cc/N7P2-TXHV. Accessed June 30, 2015.
52. Gelman A, Hill J (2006) *Data Analysis Using Regression and Multilevel/Hierarchical Models* (Cambridge Univ Press, Boston).
53. Census of India (2011) Population enumeration data. Available at perma.cc/4JKF-UWQZ. Accessed June 30, 2015.
54. Simon HA (1954) Bandwagon and underdog effects and the possibility of election predictions. *Pub Op Quarterly* 18(3):245–253.
55. Leip D (2012) Dave Leip's atlas of U.S. presidential elections. Available at perma.cc/G9YF-6CPE. Accessed June 30, 2015.
56. Rogers D, Cage F (2012) US house and senate elections 2012. Available at perma.cc/5QSK-FNCL. Accessed June 30, 2015.
57. Cabletelevision Advertising Bureau (2012) CAB undecided voters study 2012. Available at perma.cc/T82K-YC38. Accessed June 30, 2015.
58. Rutenberg J (2012) Obama campaign endgame: Grunt work and cold math. Available at perma.cc/X74D-YNM5. Accessed June 30, 2015.
59. Cox GW, McCubbins MD (1986) Electoral politics as a redistributive game. *J Polit* 48(2):370–389.
60. Lindbeck A, Weibull JW (1987) Balanced-budget redistribution as the outcome of political competition. *Public Choice* 52(3):273–297.
61. Smith A (2011) Pew Research Center's Internet & American Life Project: The Internet and campaign 2010. Available at perma.cc/3U9R-EUQQ. Accessed June 30, 2015.
62. Smith A, Duggan M (2012) Pew Research Center's Internet & American life project: Online political videos and campaign 2012. Available at perma.cc/5UX2-LELJ. Accessed June 30, 2015.
63. Cisco (2014) VNI Global IP Traffic Forecast, 2013–2018. Available at perma.cc/AE9R-J5HL. Accessed June 30, 2015.
64. Ash T, Ginty M, Page R (2012) *Landing Page Optimization: The Definitive Guide to Testing and Tuning for Conversions* (John Wiley & Sons, Indianapolis).
65. Kraus SJ (1995) Attitudes and the prediction of behavior: A meta-analysis of the empirical literature. *Pers Soc Psychol Bull* 21(1):58–75.
66. Zajonc RB (2001) Mere exposure: A gateway to the subliminal. *Curr Dir Psychol Sci* 10(6):224–228.
67. Berger J, Fitzsimons G (2008) Dogs on the street, pumas on your feet: How cues in the environment influence product evaluation and choice. *J Mark Res* 45(1):1–14.
68. Brasel SA, Gips J (2011) Red bull "gives you wings" for better or worse: A double-edged impact of brand exposure on consumer performance. *J Consum Psychol* 21(1):57–64.
69. Fransen ML, Fennis BM, Pruyn ATH, Das E (2008) Rest in peace? Brand-induced mortality salience and consumer behavior. *J Bus Res* 64(10):1053–1061.
70. Bargh JA, Gollwitzer PM, Lee-Chai A, Barndollar K, Trötschel R (2001) The automated will: Nonconscious activation and pursuit of behavioral goals. *J Pers Soc Psychol* 81(6):1014–1027.
71. Pronin E, Kugler MB (2007) Valuing thoughts, ignoring behavior: The introspection illusion as a source of the bias blind spot. *J Exp Soc Psychol* 43(4):565–578.
72. Winkielman P, Berridge KC, Wilbarger JL (2005) Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Pers Soc Psychol Bull* 31(1):121–135.
73. Karremans JC, Stroebe W, Claus J (2006) Beyond Vicary's fantasies: The impact of subliminal priming and brand choice. *J Exp Soc Psychol* 42(6):792–798.
74. Légal JB, Chappé J, Coiffard V, Villard-Forest A (2012) Don't you know that you want to trust me? Subliminal goal priming and persuasion. *J Exp Soc Psychol* 48(1):358–360.
75. Mills J, Aronson E (1965) Opinion change as a function of the communicator's attractiveness and desire to influence. *J Pers Soc Psychol* 1(2):173–177.
76. Bi B, Shokouhi M, Kosinski M, Graepel T (2013) Inferring the demographics of search users: Social data meets search queries. *Proceedings of the 22nd International Conference on World Wide Web* (International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland), pp 131–140.
77. Yan J, et al. (2009, April). How much can behavioral targeting help online advertising? *Proceedings of the 18th International Conference on World Wide Web* (ACM, New York), pp 261–270.
78. Weber I, Castillo C (2010) The demographics of web search. *Proceedings of the 33rd Annual International Association for Computing Machinery Special Interest Group on Information Retrieval Conference on Research and Development in Information Retrieval* (ACM, New York), pp 523–530.
79. Mayer WG (2008) *The Swing Voter in American Politics* (Brookings Institution Press, Washington, DC).
80. Cox GW (2010) Swing voters, core voters, and distributive politics. *Political Representation*, eds Shapiro I, Stokes SC, Wood EJ, Kirshner AS (Cambridge Univ Press, Cambridge, MA), pp 342–357.
81. Bickers KN, Stein RM (1996) The electoral dynamics of the federal pork barrel. *Am J Pol Sci* 40(4):1300–1325.
82. Dahlberg M, Johansson E (2002) On the vote-purchasing behavior of incumbent governments. *Am Polit Sci Rev* 96(1):27–40.
83. Stein RM, Bickers KN (1994) Congressional Elections and the Pork Barrel. *J Polit* 56(2):377–399.
84. Stokes SC (2005) Perverse accountability: A formal model of machine politics with evidence from Argentina. *Am Polit Sci Rev* 99(3):315–325.
85. Mayer WG (2012) The disappearing–but still important–swing voter. *Forum* 10(3):1520.
86. Agichtein E, Brill E, Dumais S (2006) Improving web search ranking by incorporating user behavior information. *Proceedings of the 29th Annual International Association for Computing Machinery Special Interest Group on Information Retrieval Conference on Research and Development in Information Retrieval* (ACM, New York), pp 19–26.
87. Ahmed A, Aly M, Das A, Smola AJ, Anastasakos T (2012) Web-scale multi-task feature selection for behavioral targeting. *Proceedings of the 21st Association for Computing Machinery International Conference on Information and Knowledge Management* (ACM, New York), pp 1737–1741.
88. Chen Y, Pavlov D, Canny JF (2009) Large-scale behavioral targeting. *Proceedings of the 15th Association for Computing Machinery Special Interest Group on Knowledge Discovery in Data International Conference on Knowledge Discovery and Data Mining* (ACM, New York), pp 209–218.
89. Höchstötter N, Lewandowski D (2009) What users see—Structures in search engine results pages. *Inf Sci* 179(12):1796–1812.
90. Page L, Brin S, Motwani R, Winograd T (1999) *The PageRank Citation Ranking: Bringing Order to the Web. Technical Report* (Stanford InfoLab, Palo Alto, CA).

Supporting Information

Epstein and Robertson 10.1073/pnas.1419828112

SI Text

Demographic Differences in VMP. In study 2, we found substantial differences in how vulnerable different demographic groups were to SEME. Although the groups we examined are somewhat arbitrary, overlapping, and by no means definitive, they do establish a range of vulnerability to SEME. Ten groups ($n \geq 50$) that appeared to be highly vulnerable in study 2, as indicated by their VMP scores, were, in order from highest to lowest, as follows:

i) Moderate Republicans (80.0%; 95% CI, 62.5–97.5%)
ii) People from North Carolina (66.7%; 95% CI, 42.8–90.5%)
iii) Moderate Libertarians (73.3%; 95% CI, 51–95.7%)
iv) Male Republicans (66.1%; 95% CI, 54–78.2%)
v) Female conservatives age 30 and over (67.7%; 95% CI, 52.5–82.7%)
vi) People from Virginia (60.0%; 95% CI, 38.5–81.5%)
vii) People earning between \$15,000 and \$19,999 (60.0%; 95% CI, 42.5–77.5%)
viii) Hispanics (59.4%; 95% CI, 42.4–76.4%)
ix) Independents with no political leaning (58.3%; 95% CI, 38.6–78.1%)
x) Female conservatives (54.7%; 95% CI, 41.3–68.1%)

Ten groups that appeared to show little vulnerability to SEME, as indicated by their VMP scores, were, in order from highest to lowest, as follows:

i) People from California (24.1%; 95% CI, 15.1–33.1%)
ii) Moderate independents (24.0%; 95% CI, 15.4–32.5%)
iii) Liberal independents (23.4%; 95% CI, 13.1–33.8%)
iv) People from Texas (22.9%; 95% CI, 11–34.8%)
v) Liberal Libertarians (22.7%; 95% CI, 5.2–40.2%)
vi) People earning between \$40,000 and \$49,999 (22.5%; 95% CI, 13.8–31.1%)
vii) Female independents (22.0%; 95% CI, 13.5–30.5%)
viii) Male moderates age 30 and over (19.3%; 95% CI, 9.1–29.5%)
ix) Female independent moderates (17.9%; 95% CI, 13.5–30.5%)
x) People with an uncommon political party (15.0%; 95% CI, –0.6% to 30.6%)

In study 3, as in study 2, we found substantial differences in how vulnerable different demographic groups were to SEME. Although the groups we examined are somewhat arbitrary, overlapping, and by no means definitive, they do establish a range of vulnerability to SEME. Ten groups ($n \geq 50$) that appeared to be highly vulnerable in study 3, as indicated by their VMP scores, were, in order from highest to lowest, as follows:

i) Unemployed males from Kerala (72.7%; 95% CI, 46.4–99.1%)
ii) Unemployed Christians (68.8%; 95% CI, 46.0–91.5%)
iii) Unemployed moderate males (50.0%; 95% CI, 33.2–66.8%)
iv) Moderate Christian males (47.6%; 95% CI, 26.3–69.0%)
v) Christian moderates (42.9%; 95% CI, 26.5–59.3%)
vi) Males from Kerala (40.4%; 95% CI, 26.4–54.5%)
vii) Unemployed moderates (33.3%; 95% CI, 22.0–44.7%)
viii) Male Christians (32.7%; 95% CI, 19.9–45.4%)
ix) People from Kerala (32.4%; 95% CI, 21.8–43.1%)
x) Unemployed females with no political ideology (31.6%; 95% CI, 10.7–52.5%)

Ten groups that appeared to show little vulnerability to SEME, as indicated by their VMP scores, were, in order from highest to lowest, as follows:

i) People from Tamil Nadu with no political ideology (0.0%; 95% CI, –0.01%–0.04%)
ii) Employed females with no political ideology (0.0%; 95% CI, –0.01%–0.06%)
iii) People earning between Rs 10,000 and Rs 29,999 (–3.2%; 95% CI, –7.6%–1.3%)
iv) Married people who are separated (–3.3%; 95% CI, –10.0%–3.3%)
v) People with a pre-university education (–4.3%; 95% CI, –10.5%–1.81%)
vi) Unemployed liberals (–4.3%; 95% CI, –10.5%–1.81%)
vii) Unemployed conservatives (–5.0%; 95% CI, –15.0%–5.0%)
viii) People from Gujarat (–5.9%; 95% CI, –17.8%–6.0%)
ix) Unemployed male liberals (–8.0%; 95% CI, –19.5%–3.5%)
x) Female conservatives (–11.8%; 95% CI, –29.0%–5.5%)

Bias Awareness. Subjects were counted as showing awareness of the manipulation if (*i*) they had clicked on a box indicating that something "bothered" them about the rankings and (*ii*) we found specific terms or phrases in their open-ended comments suggesting that they were aware of bias in the rankings, such as "biased," "bias," "leaning towards," "leaning toward," "leaning against," "slanted," "skewed," "favorable towards," "favorable toward," "favorable for," "favorable against," "favorable results," "favored towards," "favored toward," "favored for," "favored against," "favored results," "favor toward," "results favor," "favor Modi," "favor Kejriwal," "favor Gandhi," "negative toward," "negative for," "negative against," "all negative," "all positive," "mainly negative," "mainly positive," "nothing positive," "nothing negative," "more results for," "less results for," "most of the articles were negative," "most of the articles were positive," "pro Modi," "pro Kejriwal," "pro Gandhi," "Modi leaning," "Kejriwal leaning," "Gandhi leaning," "pro Gillard," "pro Abbott," "favor Gillard," "favor Abbott," "Gillard leaning," and "Abbott leaning."

Derivation of the Formulas for Computing *W*, the Maximum Win Margin Controllable Through SEME, in Two- and Three- Person Races.

Two-person race. Where T = total number of eligible voters in a population, i = proportion of T who are internet users, u = proportion of i who are undecided, p = proportion of u who are prone to vote for the target candidate, and VMP = proportion of p who can be shifted by SEME.

The number of votes that can be shifted by SEME is given by

$$n = T * i * u * p * \mathrm{VMP}.$$

In a two-person race, the number of votes for the candidate favored by SEME when the vote is initially evenly split is

$$\frac{T}{2} + n,$$

and the number of votes for the losing candidate is

$$\frac{T}{2} - n.$$

The vote margin in favor of the winning candidate is therefore the larger vote minus the smaller vote, or, simply: $2n$.

Therefore, the margin of voters, expressed as a proportion, that can be shifted by SEME is

$$\frac{2n}{T}=\frac{2*T*i*u*p*\mathrm{VMP}}{T}=2*i*u*p*\mathrm{VMP}.$$

Because the undecided voters in a two-person race have only two voting options, the value of p before outside influence is exercised can reasonably be assumed to be 0.5.

Therefore, W can be calculated as follows:

$$W=2*i*u*0.5*\mathrm{VMP},$$

and the calculation can be simplified as follows:

$$W=i*u*\mathrm{VMP}.$$

In other words, the maximum win margin controllable by SEME in a two-person race is equal to the proportion of people who can be influenced by SEME (the VMP) times the proportion of undecided Internet voters in the population. $(i*u)$.

Three-person race. Where $T=$ total number of voters in a population, $i=$ proportion of T who are internet users, $u=$ proportion of i who are undecided, $p=$ proportion of u who are prone to vote for the target candidate, and VMP $=$ proportion of p who can be shifted by SEME.

The number of votes that can be shifted by SEME is given by

$$n=T*i*u*p*\mathrm{VMP}.$$

In a three-person race, because the winning candidate can draw votes from either of the two losing candidates, W can vary between two extremes:

i) At one extreme, one of the two losing candidates draws zero votes, in which case the formula for the two-person case (above) is applicable.

ii) At the other extreme, voting preferences are initially split three ways evenly, and the winning candidate draws votes equally from the other two. This distribution will give us the lowest possible value of W in the three-person race, as follows.

The number of votes for the candidate favored by SEME will still be

$$\frac{T}{2}+n.$$

However, because of the split, the number of votes for each of the losing candidates will now be

$$\frac{T}{2}-\frac{n}{2}.$$

The vote margin in favor of the winning candidate will therefore be the larger vote minus either of the smaller votes or, simply, $1.5n$.

Therefore, the margin of voters, expressed as a proportion, that can be shifted by SEME is

$$\frac{2n}{T}=\frac{1.5*T*i*u*p*\mathrm{VMP}}{T}=1.5*i*u*p*\mathrm{VMP}.$$

Therefore, W can be calculated as follows:

$$W=1.5*i*u*0.5*\mathrm{VMP},$$

and the calculation can be simplified as follows:

$$W=0.75*i*u*\mathrm{VMP}.$$

Therefore, in a three-person race, W will vary between 75% and 100% of the W found in the two-person case, depending on how votes are distributed between the two losing candidates; the more even the split, the smaller the controllable win margin.

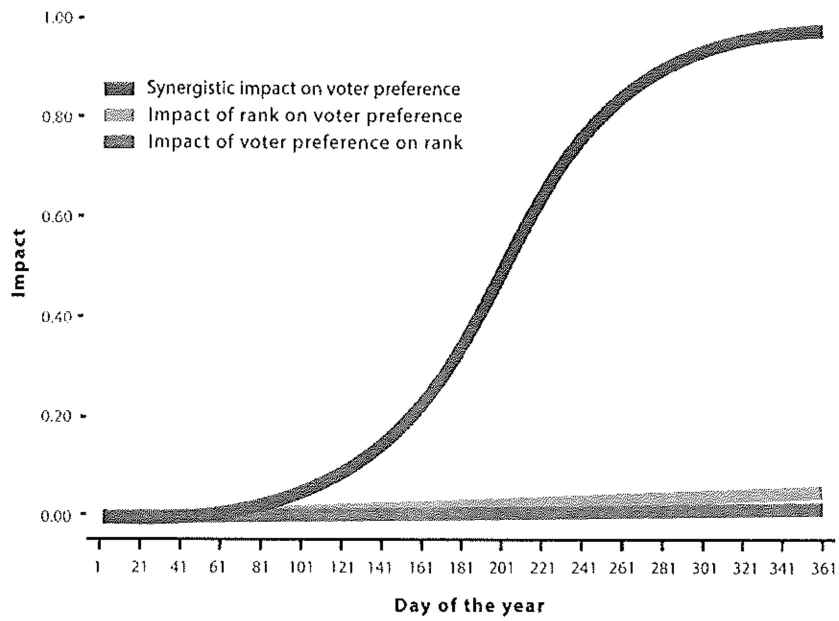

Fig. S1. A possible synergistic relationship between the impact that search rankings have on voter preferences and the impact that voter preferences have on search rankings. The lower curves (red and green) show slow increases that might occur if each of the processes acted alone over the course of a year (365 iterations of the model). The upper curve (blue) shows the result of a possible synergy between these two processes using the same parameters that generated the two lower curves. The curves are generated by an iterative model using equations of the general form $V_{n+1} = V_n + r[R_n \times (1 - V_n)] + r[O_n \times (1 - V_n)]$, where V is voter preference for one candidate, R is the impact of voter preferences on search rankings, O is the impact (randomized with each iteration) of other influences on voter preferences, and r is a rate-of-change factor. Because a change in voter preference alters the proportion of votes available, its value in the model cannot exceed 1.0.

Table S1. Demographics for studies 1 and 2

Category	Value	Census 2010[†]		Study 1		Census and study 1	Study 2	
		n	%	*n*	%	*Z*	*n*	%
Age	18–24	26,718	12.7%	51	16.7%	2.097*	446	21.2%
	25–44	70,472	33.4%	122	39.9%	2.385*	1,274	60.7%
	45–64	75,865	36.0%	95	31.0%	1.800	342	16.3%
	65–74	20,605	9.8%	20	6.5%	1.906	33	1.6%
	75+	17,140	8.1%	18	5.9%	1.438	5	0.2%
Race	White	152,929	72.5%	179	58.5%	5.502***	1,645	78.3%
	Black	25,632	11.8%	38	12.4%	0.349	126	6.0%
	Hispanic	21,285	9.8%	52	17.0%	4.169***	121	5.8%
	Asian	7,638	3.9%	7	2.3%	1.528	123	5.9%
	Other	3,316	2.0%	30	9.8%	10.977***	85	4.0%
Sex	Male	101,279	48.0%	162	52.9%	1.715	1,148	54.7%
	Female	109,521	52.0%	144	47.1%	1.715	947	45.1%
	Other	n/a	n/a	0	0.0%	n/a	5	0.2%
Education	Less than ninth grade	6,655	3.2%	2	0.7%	2.504*	0	0.0%
	Ninth to 12th grade	15,931	7.6%	45	14.7%	4.724***	22	1.0%
	High school graduate	65,951	31.3%	68	22.2%	3.417***	231	11.0%
	Some college or associate degree	62,655	29.7%	145	47.4%	6.753***	820	39.0%
	Bachelors	39,272	18.6%	30	9.8%	3.963***	752	35.8%
	Advanced	20,336	9.6%	16	5.2%	2.616**	275	13.1%
Used[‡]	Yes	126,477	60.0%	119	38.9%	7.531***	1,509	71.9%
	No	84,323	40.0%	187	61.1%	7.531***	591	28.1%
Income	Under $10,000	5,496	3.6%	67	21.9%	20.009***	137	6.5%
	$10,000 to $14,999	5,069	3.3%	33	10.8%	8.538***	131	6.2%
	$15,000 to $19,999	4,549	2.9%	28	9.2%	7.446***	124	5.9%
	$20,000 to $29,999	12,632	8.2%	45	14.7%	4.800***	282	13.4%
	$30,000 to $39,999	13,182	8.5%	34	11.1%	1.857	288	13.7%
	$40,000 to $49,999	10,807	7.0%	17	5.6%	1.143	239	11.4%
	$50,000 to $74,999	25,516	16.5%	30	9.8%	3.602***	405	19.3%
	$75,000 to $99,999	17,597	11.4%	11	3.6%	4.932***	235	11.2%
	$100,000 to $149,999	16,586	10.7%	5	1.6%	5.916***	148	7.0%
	$150,000 and over	12,102	7.8%	0	0.0%	5.893***	46	2.2%
	Prefer not to say	30,875	20.0%	36	11.8%	4.069***	65	3.1%
Marital status	Married	113,421	53.8%	48	15.7%	13.364***	751	35.8%
	Widowed	13,612	6.5%	27	8.8%	1.682	15	0.7%
	Divorced	23,035	10.9%	68	22.2%	6.324***	141	6.7%
	Separated	4,528	2.1%	15	4.9%	3.317***	33	1.6%
	Never married	56,203	26.7%	148	48.4%	8.576***	1,160	55.2%

$^{*}P < 0.05$; $^{**}P < 0.01$; and $^{***}P < 0.001$.

[†]Census numbers are in hundred thousands.

[‡]For census data, "No" includes "unemployed" and "not in labor force."

Table S2. **Voting preferences by group for study 1**

Experiment	Voting preferences	Mean (SE)			Kruskal–Wallis (χ^2)	Mann–Whitney *u*
		Group 1 (Gillard bias)	Group 2 (Abbott bias)	Group 3 (control)		
1	PreImpressionAbbott	8.09 (0.34)	7.74 (0.40)	7.41 (0.26)	3.979	525.0
	PreImpressionGillard	7.06 (0.42)	7.47 (0.35)	6.88 (0.32)	1.395	529.5
	PreTrustAbbott	7.82 (0.31)	7.85 (0.39)	7.35 (0.28)	3.275	538.5
	PreTrustGillard	6.38 (0.40)	7.56 (0.30)	6.88 (0.32)	5.213	407.0
	PreLikeAbbott	6.06 (0.52)	5.68 (0.47)	5.79 (0.38)	0.296	538.5
	PreLikeGillard	5.29 (0.48)	5.76 (0.41)	5.29 (0.37)	1.335	500.0
	PostImpressionAbbott	4.24 (0.49)	7.29 (0.51)	5.85 (0.38)	19.029***	252.0***
	PostImpressionGillard	7.26 (0.45)	4.71 (0.47)	5.65 (0.46)	14.667**	286.0**
	PostTrustAbbott	4.59 (0.43)	7.32 (0.51)	6.15 (0.38)	18.385***	260.5***
	PostTrustGillard	6.91 (0.42)	4.97 (0.43)	6.15 (0.40)	10.809**	326.5**
	PostLikeAbbott	3.88 (0.43)	6.24 (0.58)	5.18 (0.42)	11.026**	341.5**
	PostLikeGillard	5.68 (0.49)	4.15 (0.45)	5.41 (0.42)	5.836	403.0*
2	PreImpressionAbbott	6.76 (0.43)	7.50 (0.34)	6.76 (0.44)	1.761	477.0
	PreImpressionGillard	6.50 (0.36)	7.29 (0.43)	6.12 (0.45)	4.369	449.5
	PreTrustAbbott	6.41 (0.44)	7.12 (0.30)	7.32 (0.44)	2.700	499.0
	PreTrustGillard	6.56 (0.41)	7.32 (0.36)	6.35 (0.43)	3.094	465.0
	PreLikeAbbott	5.56 (0.46)	5.65 (0.43)	5.76 (0.49)	0.170	575.0
	PreLikeGillard	5.79 (0.44)	5.79 (0.48)	5.47 (0.45)	0.306	568.0
	PostImpressionAbbott	3.79 (0.41)	7.15 (0.49)	5.24 (0.48)	20.878***	226.5***
	PostImpressionGillard	7.35 (0.39)	4.79 (0.47)	6.00 (0.38)	15.270***	279.5***
	PostTrustAbbott	3.82 (0.40)	7.18 (0.47)	5.53 (0.51)	21.917***	207.5***
	PostTrustGillard	7.32 (0.41)	4.97 (0.46)	6.18 (0.36)	13.410**	302.0**
	PostLikeAbbott	3.91 (0.42)	6.09 (0.53)	5.56 (0.48)	9.822**	353.0**
	PostLikeGillard	6.68 (0.45)	4.29 (0.48)	5.79 (0.40)	12.905**	311.5**
3	PreImpressionAbbott	7.24 (0.39)	7.18 (0.39)	7.88 (0.27)	1.346	568.5
	PreImpressionGillard	6.12 (0.43)	7.09 (0.39)	7.26 (0.34)	4.134	452.0
	PreTrustAbbott	7.18 (0.35)	6.41 (0.41)	7.53 (0.32)	3.837	478.0
	PreTrustGillard	6.65 (0.38)	6.68 (0.40)	6.97 (0.33)	0.259	568.5
	PreLikeAbbott	6.59 (0.42)	5.94 (0.39)	6.59 (0.43)	2.301	491.0
	PreLikeGillard	5.85 (0.46)	5.85 (0.43)	6.26 (0.41)	1.065	576.5
	PostImpressionAbbott	5.29 (0.48)	6.82 (0.41)	6.26 (0.48)	5.512	384.0*
	PostImpressionGillard	6.50 (0.45)	5.47 (0.43)	6.21 (0.48)	3.027	445.5
	PostTrustAbbott	5.38 (0.49)	6.85 (0.45)	6.47 (0.47)	5.091	399.0*
	PostTrustGillard	6.44 (0.45)	5.76 (0.47)	6.29 (0.44)	1.365	493.0
	PostLikeAbbott	5.29 (0.48)	6.03 (0.48)	5.79 (0.53)	1.129	487.0
	PostLikeGillard	6.12 (0.47)	5.26 (0.54)	6.09 (0.51)	1.475	491.5

*$P < 0.05$; **$P < 0.01$; and ***$P < 0.001$: Kruskal–Wallis tests were conducted between all three groups, and Mann–Whitney *u* tests were conducted between groups 1 and 2. Preferences were measured for each candidate separately on 10-point Likert scales.

Table S3. **Voting preferences by group for study 2**

Voting preferences	Mean (SE)			Kruskal–Wallis (χ^2)	Mann–Whitney *u*
	Group 1 (Gillard bias)	Group 2 (Abbott bias)	Group 3 (control)		
PreImpressionAbbott	7.40 (0.07)	7.36 (0.08)	7.37 (0.07)	0.458	241,861.5
PreImpressionGillard	7.13 (0.07)	7.12 (0.08)	7.13 (0.07)	0.081	243,115.0
PreTrustAbbott	7.26 (0.07)	7.22 (0.08)	7.18 (0.07)	0.954	241,924.5
PreTrustGillard	6.95 (0.07)	6.89 (0.08)	6.92 (0.07)	0.222	241,779.0
PreLikeAbbott	6.42 (0.08)	6.39 (0.08)	6.23 (0.08)	2.987	243,677.5
PreLikeGillard	6.24 (0.08)	6.30 (0.08)	6.11 (0.08)	3.178	239,556.0
PostImpressionAbbott	4.61 (0.09)	6.88 (0.09)	5.53 (0.09)	289.065***	120,660.0***
PostImpressionGillard	6.87 (0.08)	4.95 (0.09)	6.21 (0.09)	237.034***	133,106.5***
PostTrustAbbott	4.56(0.10)	6.94 (0.09)	5.57 (0.10)	281.560***	121,786.5***
PostTrustGillard	6.84 (0.09)	4.95 (0.09)	6.19 (0.09)	221.709***	136,689.0***
PostLikeAbbott	4.55 (0.09)	6.31 (0.09)	5.21 (0.09)	177.225***	146,957.0***
PostLikeGillard	6.34(0.09)	4.64 (0.09)	5.71 (0.09)	176.066***	147,372.5***

***$P < 0.001$: Kruskal–Wallis tests were conducted between all three groups, and Mann–Whitney *u* tests were conducted between groups 1 and 2. Preferences were measured for each candidate separately on 10-point Likert scales.

Table S4. Treatment effect estimates for study 2 voting preferences

Predictor variable	Presearch vote		Postsearch vote	
	Coefficient	SE	Coefficient	SE
Intercept	−0.073	0.540	0.062	0.543
Sex				
Female	0	Referent	0	Referent
Male	0.039	0.110	−0.135	0.119
Other	−0.430	0.922	−0.568	0.924
Race/ethnicity				
White	0	Referent	0	Referent
Black	0.115	0.224	0.090	0.245
Hispanic	−0.435	0.235	−0.280	0.237
Asian	0.366	0.238	0.668	0.291*
Other	0.133	0.274	−0.072	0.291
Age group				
18–24	0	Referent	0	Referent
25–44	−0.024	0.144	−0.083	0.157
45–64	0.241	0.184	0.029	0.200
65+	0.258	0.411	0.685	0.519
Education level				
Less than ninth grade	0	Referent	0	Referent
Ninth to 12th grade	0.024	0.548	0.732	0.550
High school graduate	0.074	0.528	0.927	0.528
Bachelors	0.094	0.529	0.842	0.530
Advanced	−0.050	0.543	0.549	0.544

The presearch and postsearch columns report the estimate and variance for both treatment groups using classical regression poststratification. Data for sex, race/ethnicity, age group, and education level came from the 2010 US Census. Data on the number of people who identify their sex as "other" came from a 2011 Gallup study.
$^*P < 0.05$.

Table S5. **Demographics for study 3**

Category	Value	Study 3		Indian Census 2011 (literates)	
		n	%	*n*	%
Age	18–24	602	28.0%	160,241,457	21.0%
	25–44	1410	65.6%	347,587,712	45.6%
	45–64	124	5.8%	188,197,343	24.7%
	65+	14	0.7%	66,185,333	8.7%
Religion	Buddhism	14	0.7%	—	—
	Christianity	262	12.2%	—	—
	Hinduism	1512	70.3%	—	—
	Islam	314	14.6%	—	—
	Jainism	21	1.0%	—	—
	Other	15	0.7%	—	—
	Sikhism	12	0.6%	—	—
Sex	Male	1518	70.6%	388,428,872	51.0%
	Female	632	29.4%	373,782,973	49.0%
Education	None	0	0.0%	—	—
	Primary school	4	0.2%	—	—
	Higher secondary	71	3.3%	—	—
	Pre-university	136	6.3%	—	—
	Bachelors	1225	57.0%	—	—
	Masters	699	32.5%	—	—
	Doctorate	15	0.7%	—	—
Used	Yes	1635	76.0%	—	—
	No	515	24.0%	—	—
Income	Under Rs 10,000	121	5.6%	—	—
	Rs 10,000 to Rs 29,999	206	9.6%	—	—
	Rs 30,000 to Rs 49,999	131	6.1%	—	—
	Rs 50,000 to Rs 69,999	106	4.9%	—	—
	Rs 70,000 to Rs 89,999	146	6.8%	—	—
	Rs 90,000 to Rs 109,999	181	8.4%	—	—
	Rs 110,000 to Rs 129,999	172	8.0%	—	—
	Rs 130,000 to Rs 149,999	132	6.1%	—	—
	Rs 150,000 to Rs 169,999	124	5.8%	—	—
	Rs 170,000 to Rs 189,999	118	5.5%	—	—
	Rs 190,000 and over	486	22.6%	—	—
	I prefer not to say	227	10.6%	—	—
Marital status	Married	1,144	53.2%	—	—
	Widowed	5	0.2%	—	—
	Divorced	4	0.2%	—	—
	Separated	78	3.6%	—	—
	Never married	919	42.7%	—	—
Location	State	1,144	53.2%	749,758,470	98.4%
	Union Territory	5	0.2%	12,453,375	1.6%

Table S6. Voting Preferences by Group for Study 3

	Mean (SE)			
Voting preferences	Group 1 (Gandhi bias)	Group 2 (Kejriwal bias)	Group 3 (Modi bias)	Kruskal–Wallis (χ^2)
PreImpressionGandhi	5.94 (0.10)	5.73 (0.10)	5.65 (0.10)	4.782
PreImpressionKejriwal	6.80 (0.09)	7.07 (0.09)	7.09 (0.08)	6.230*
PreImpressionModi	7.49 (0.10)	7.46 (0.10)	7.48 (0.09)	0.188
PreLikableGandhi	5.71 (0.10)	5.64 (0.10)	5.61 (0.10)	0.722
PreLikableKejriwal	6.68 (0.09)	6.78 (0.09)	6.87 (0.09)	2.030
PreLikableModi	7.40 (0.10)	7.29 (0.10)	7.29 (0.10)	1.483
PreTrustGandhi	5.57 (0.11)	5.52 (0.11)	5.42 (0.10)	0.955
PreTrustKejriwal	6.54 (0.10)	6.74 (0.10)	6.85 (0.09)	4.546
PreTrustModi	7.22 (0.11)	7.31 (0.11)	7.27 (0.10)	0.159
PreLikelyToVoteGandhi	0.10 (0.12)	0.08 (0.12)	0.08 (0.12)	1.587
PreLikelyToVoteKejriwal	1.19 (0.11)	1.38 (0.11)	1.55 (0.10)	5.178
PreLikelyToVoteModi	2.15 (0.12)	2.12 (0.12)	2.06 (0.12)	0.202
PostImpressionGandhi	5.78 (0.10)	5.52 (0.10)	5.35 (0.10)	9.552**
PostImpressionKejriwal	6.50 (0.09)	6.96 (0.09)	6.70 (0.08)	14.288**
PostImpressionModi	7.27 (0.10)	7.26 (0.10)	7.60 (0.09)	7.860*
PostLikableGandhi	5.62 (0.10)	5.46 (0.10)	5.26 (0.10)	6.322*
PostLikableKejriwal	6.37 (0.09)	6.84 (0.09)	6.64 (0.08)	13.456**
PostLikableModi	7.24 (0.11)	7.20 (0.11)	7.47 (0.10)	3.874
PostTrustGandhi	5.71 (0.11)	5.48 (0.10)	5.22 (0.10)	11.386*
PostTrustKejriwal	6.38 (0.10)	6.89 (0.10)	6.68 (0.08)	15.840***
PostTrustModi	7.18 (0.11)	7.20 (0.11)	7.49 (0.10)	4.758

*$P < 0.05$; **$P < 0.01$; and ***$P < 0.001$: Kruskal–Wallis tests were conducted between all three groups. Preferences were measured for each candidate separately on 10-point Likert scales.

Table S7. Treatment effect estimates for study 3 voting preferences

	Presearch vote		Postsearch vote	
Predictor variable	Coefficient	SE	Coefficient	SE
Intercept	−0.716	0.090***	−0.552	0.088***
Sex				
Male	0	Referent	0	Referent
Female	0.168	0.100	0.030	0.099
Age group, y				
18–24	0	Referent	0	Referent
25–44	0.031	0.103	0.067	0.101
45–64	−0.222	0.217	−0.057	0.208
65+	−0.213	0.598	−0.366	0.598
Location				
State	0	Referent	0	Referent
Union Territory	−0.401	0.294	−0.321	0.279

The presearch and postsearch columns report the estimate and variance for both of the treatment groups using classical regression poststratification. Data for sex, age group, and location came from the 2011 India Census.

***$P < 0.001$.

Table S8. Minimum VMP levels needed to impact two-person races with various projected win margins and proportions of undecided Internet voters

Proportion of undecided Internet voters in the population ($i*u$)	Projected win margin									
	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	0.10
0.01	1.000	—	—	—	—	—	—	—	—	—
0.02	0.500	1.000	—	—	—	—	—	—	—	—
0.03	0.333	0.667	1.000	—	—	—	—	—	—	—
0.04	0.250	0.500	0.750	1.000	—	—	—	—	—	—
0.05	0.200	0.400	0.600	0.800	1.000	—	—	—	—	—
0.06	0.167	0.333	0.500	0.667	0.833	1.000	—	—	—	—
0.07	0.143	0.286	0.429	0.571	0.714	0.857	1.000	—	—	—
0.08	0.125	0.250	0.375	0.500	0.625	0.750	0.875	1.000	—	—
0.09	0.111	0.222	0.333	0.444	0.556	0.667	0.778	0.889	1.000	—
0.10	0.100	0.200	0.300	0.400	0.500	0.600	0.700	0.800	0.900	1.000
0.11	0.091	0.182	0.273	0.364	0.455	0.545	0.636	0.727	0.818	0.909
0.12	0.083	0.167	0.250	0.333	0.417	0.500	0.583	0.667	0.750	0.833
0.13	0.077	0.154	0.231	0.308	0.385	0.462	0.538	0.615	0.692	0.769
0.14	0.071	0.143	0.214	0.286	0.357	0.429	0.500	0.571	0.643	0.714
0.15	0.067	0.133	0.200	0.267	0.333	0.400	0.467	0.533	0.600	0.667
0.16	0.063	0.125	0.188	0.250	0.313	0.375	0.438	0.500	0.563	0.625
0.17	0.059	0.118	0.176	0.235	0.294	0.353	0.412	0.471	0.529	0.588
0.18	0.056	0.111	0.167	0.222	0.278	0.333	0.389	0.444	0.500	0.556
0.19	0.053	0.105	0.158	0.211	0.263	0.316	0.368	0.421	0.474	0.526
0.20	0.050	0.100	0.150	0.200	0.250	0.300	0.350	0.400	0.450	0.500

Other Supporting Information Files

Dataset S1 (XLS)

PLOS ONE

RESEARCH ARTICLE

The Answer Bot Effect (ABE): A powerful new form of influence made possible by intelligent personal assistants and search engines

Robert Epstein*, Vivian Lee, Roger Mohr, Vanessa R. Zankich

American Institute for Behavioral Research and Technology, Vista, California, United States of America

* re@aibrt.org

Check for updates

Citation: Epstein R, Lee V, Mohr R, Zankich VR (2022) The Answer Bot Effect (ABE): A powerful new form of influence made possible by intelligent personal assistants and search engines. PLoS ONE 17(6): e0268081. https://doi.org/10.1371/journal.pone.0268081

Editor: Lalit Chandra Saikia, National Institute of Technology Silchar, India, INDIA

Received: December 20, 2021

Accepted: April 21, 2022

Published: June 1, 2022

Data Availability Statement: An anonymized version of the data has been posted at https://doi.org/10.5281/zenodo.6537353. Data can also be requested from info@aibrt.org. The data have been anonymized to comply with requirements of the sponsoring institution's Institutional Review Board (IRB).

Funding: The author(s) received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

Abstract

We introduce and quantify a relatively new form of influence: the Answer Bot Effect (ABE). In a 2015 report in PNAS, researchers demonstrated the power that biased search results have to shift opinions and voting preferences without people's knowledge–by up to 80% in some demographic groups. They labeled this phenomenon the Search Engine Manipulation Effect (SEME), speculating that its power derives from the high level of trust people have in algorithmically-generated content. We now describe three experiments with a total of 1,736 US participants conducted to determine to what extent giving users "the answer"–either via an answer box at the top of a page of search results or via a vocal reply to a question posed to an intelligent personal assistant (IPA)–might also impact opinions and votes. Participants were first given basic information about two candidates running for prime minister of Australia (this, in order to assure that participants were "undecided"), then asked questions about their voting preferences, then given answers to questions they posed about the candidates–either with answer boxes or with vocal answers on an Alexa simulator–and then asked again about their voting preferences. The experiments were controlled, randomized, double-blind, and counterbalanced. Experiments 1 and 2 demonstrated that answer boxes can shift voting preferences by as much as 38.6% and that the appearance of an answer box can reduce search times and clicks on search results. Experiment 3 demonstrated that even a single question-and-answer interaction on an IPA can shift voting preferences by more than 40%. Multiple questions posed to an IPA leading to answers that all have the same bias can shift voting preferences by more than 65%. Simple masking procedures still produced large opinion shifts while reducing awareness of bias to close to zero. ABE poses a serious threat to both democracy and human autonomy because (a) it produces large shifts in opinions and voting preferences with little or no user awareness, (b) it is an ephemeral form of influence that leaves no paper trail, and (c) worldwide, it is controlled almost exclusively by just four American tech companies. ABE will become a greater threat as people increasingly rely on IPAs for answers.

1. Introduction

1.1 Search results

Multiple studies conducted in recent years have demonstrated the power that search engines have to alter thinking and behavior by showing people biased search results [1–8, cf. 9–14], and research has also shown that these shifts can be produced without people's awareness [2]. Bias in search results is difficult to see, and the few people who can spot it tend to shift their views even farther in the direction of the bias than people who cannot detect the bias [2, 15].

Search engines also influence people because of the trust people have in computer-generated output. Most people have no idea how search engines work [16–18] or, for that matter, how computers or algorithms work [19], and are oblivious to the various roles that humans play in generating computer output. Humans build the algorithms that computers use, for example, and those algorithms often produce biased content because of either the intentional or unconscious bias of the programmers [20–24]. Humans also modify existing programs–sometimes quite frequently. Recent reports suggest that Google's ubiquitous search algorithm is manually adjusted more than 3,000 times a year, and those adjustments change both the content and the ordering of search results [25, 26]. Employees also deliberately add or delete content from blacklists and whitelists, which again has the effect of suppressing or boosting content [27–29]. People try to resist manipulation when they can see the human hand–authors' names on news articles, guests on television and radio shows, videos on YouTube, and so on–but they think less critically when presented with algorithmic output, which they mistakenly believe to be inherently objective [30–34, cf. 35].

The human hand behind Big Tech companies is also invisible to users in another way. People are often oblivious to the many methods these companies are employing to collect personal data about them–the equivalent of more than three million pages of information about the average person who has been using the internet since its early days [36, cf. 37]. Monetizing that personal information is the bread and butter of Big Tech, which relies on the "surveillance business model" for nearly all its income [38–40]. Algorithms that match up users and vendors now direct the flow of hundreds of billions of dollars in purchases each year, but personal information can be used in other ways as well. As any con artist can tell you, the more you know about someone, the easier it is to manipulate him or her. Big Tech companies have accumulated massive databases about billions of people worldwide, and they are increasingly showing people personalized output that is optimized to draw clicks or impact a wide variety of thinking and behavior [15, 41–46, cf. 47, 48].

1.2 Search suggestions

Search results aren't the only tools a search engine can wield to control people. Recent research shows that search suggestions–the short lists of words and phrases users are shown as they type characters into the search bar–can also shift thinking and behavior [15, 49, cf. 50–57]. Because negative (or "low-valence") words draw far more attention and clicks than neutral or positive words [58, 59], one of the simplest ways to shift opinions to favor one candidate or cause is to suppress negative search terms for that candidate or cause. Google might have done so to support Hillary Clinton's candidacy in the 2016 Presidential election [49, 60, 61, cf. 62].

1.3 Answer boxes

In 2014, Google began displaying boxes above their search results which contain a single answer to a person's query, often accompanied by a link people can click to get more information [63]. Can these answers, now called "featured snippets" or "answer boxes," also impact thinking and behavior? This is an important question not only because bias in a featured snippet might enhance the impact of biased search results and biased search suggestions, but also because an answer box could be considered a simple variant of a wide range of new content sources. Intelligent personal assistants (IPAs) such as Amazon's Alexa, Apple's Siri, Microsoft's Cortana, and the Google Assistant (on Android devices and the Google Home device), all provide just one answer in response to a query. We are, in effect, moving away from search engines–platforms that provide thousands of possible answers in response to a query–toward the type of device we have seen portrayed in science fiction movies and television shows. On the original "Star Trek" episodes, when Captain Kirk wanted information, he didn't consult a search engine; he simply said things like, "Computer, who's the best looking captain in Star Fleet?" Why would one want a list of thousands of web pages when the computer can give you a simple answer?

Over time, Google–emulated to some extent by other, less popular search engines–has introduced several types of answer boxes, among them: a rich answer box (a type of featured snippet that includes additional information such as a graph, table, image, or interactive tool), a news stories box, a knowledge box (often information from Wikipedia displayed in the upper-right-hand corner of the search results page), a box suggesting related searches, and so on [64, 65]. Our focus, however, is on what Google calls the "featured snippet," a relatively small box that is unlabeled and contains a simple answer to a user's query [66]. On June 23, 2015, when people typed the query, "Who will be the next president?," into the Google search bar, a featured snippet appeared reading, in part, "Hillary Clinton is the next President of the United States. . .. 10 Reasons Why Hillary Clinton Will Be the Next President" [67]. On October 22, 2017, when one of the authors of this paper typed "google play vs spotify" into the Google search bar, an answer box appeared immediately below the search bar reading, in part, "Google Play Music is my top pick after months of research and testing. . .. Google Play Music is better than Spotify–Business Insider" (S1 Fig). A link was included in the box to the relevant *Business Insider* article.

1.4 Answer bots and intelligent personal assistants

1.4.1 An inevitable trend. For simplicity's sake, we will refer to all electronic devices that provide simple answers to queries posed by humans as "answer bots" and define the Answer Bot Effect (ABE) as the extent to which answers provided by answer bots can alter people's opinions and behaviors. It is important to measure this effect, we believe, because of what appears to be an inevitable trend: Worldwide, people are relying less and less on search results for their answers–just as, in the early 2000s, people began to rely less and less on books for their answers–and are simply accepting the answers they see in answer boxes or hear on their IPAs. Before answer boxes were introduced, people who used search engines had no choice but to click on search results and examine web pages to get their answers. As of 2016, approximately 43.9% of searches on mobile and desktop devices ended without a click; as of 2020, that percentage increased to 64.8% [68, 69; cf. 70]. Again, why click on a search result when the answer is right in front of you?

The shift toward answer bots is indicated by the increase in the number of people using IPAs. By 2019, there were 157 million smart speakers in American homes [71], and between 2019 and 2021, the number of Americans relying on voice assistants increased by nearly 20% [72]. Worldwide, more than 600 million smart speakers are expected to be in use by 2024 [72].

The spread of IPAs and answer boxes is not the only reason we need to measure and understand ABE. Children's toys are increasingly internet-connected, and many of them answer children's questions [73]. Hello Barbie has been around since 2015 and has been described as the perfect friend that can hold a two-way conversation and impact children's attitudes about gender roles [74]. My Friend Cayla, a conversationally interactive toy released the same year was banned by the German government because of fears that hackers could intercept children's questions and provide disturbing answers [75, 76, cf. 77]. Children are generally more impressionable than adults [78–80], which is why governments have often put restrictions on the kind of advertising that is directed toward young audiences [81]. With children's toys answering questions–much of the time, with no parents around–both the questions children ask and the answers the toys provide can be inappropriate and potentially harmful [74, 82, cf. 83–85]. And, like search engines, these toys don't just facilitate interactions; they also record them [86–88, cf. 89].

Both adults and children are also now conversing by the millions–sometimes knowingly, sometimes not–with chatbots, both through their computers and their mobile devices. When chatbots answer questions or promote viewpoints, they too can shift opinions and behavior [90, cf. 91]. The number of people currently conversing with chatbots is difficult to estimate, but it is certainly a large number that is increasing rapidly [92, 93]. When dating website Ashley Madison was hacked in 2015, the hackers learned, among other things, that "20 million men out of 31 million received bot mail, and about 11 million of them were chatted up by an automated 'engager'" [94, cf. 95]. Even though conversational AIs still perform relatively poorly [96, 97], wishful thinking can keep online suitors talking to chatbots for months [98].

1.4.2 Answer bot accuracy and bias. Do answer boxes, IPAs, conversational toys, and chatbots give users accurate information, and, if not, how are people affected by inaccurate answers? The rate of inaccurate responses varies considerably from one IPA to another: about 48% for Cortana, 30% for Siri, 22% for Alexa, and 13% for the Google Assistant, and these numbers vary from one study to another [99–104, cf. 105]. The level of trust people have for inaccurate answers also varies [106, cf. 107]. For most IPAs, accuracy is determined by the quality of the search engine that the assistant draws from; for Siri and the Google Assistant, that's the Google search engine [108]. Cortana's answers are presumably inferior because they draw from Bing, Microsoft's search engine [109]. Alexa's answers can be spotty because Amazon gets them using crowd sourcing [110, 111].

Needless to say, when people are highly reliant on and trusting of sources–as has becoming increasingly the case with Big Tech answer sources [31, 33, 112, 113]–the impact of inaccurate information can range from inconvenience to serious harm–or at least serious misconceptions. In 2018, a *Mashable* reporter asked Amazon's Alexa to tell him about the vapor trails one often sees following jets flying at high altitudes. Alexa responded with a baseless conspiracy theory: "Trails left by aircraft are actually chemical or biological agents deliberately sprayed at high altitudes for a purpose undisclosed to the general public in clandestine programs directed by government officials" [114, cf. 115].

False information spoken by a smart speaker is highly ephemeral: You hear it, and then it is gone, leaving no trace for authorities to examine. Information in answer boxes is also ephemeral, but it can at least be preserved with a simple screenshot. Among our favorites: In 2017, in response to the query, "presidents in the klan," a Google answer box listed four presidents, even though no U.S. president has ever been a member of the Ku Klux Klan [116] (S2 Fig). In 2018, when people searched for "California Republicans" or "California Republican Party," Google displayed a knowledge panel box listing "Nazism" as the first item under Ideology [117] (S2 Fig). On August 16, 2016, when one of the authors of this paper queried, "when is the election?," a Google answer box correctly showed November 8, 2016, but it also included a photograph of Hillary Clinton inside the answer box–just Clinton, with none of her competitors (S2 Fig).

1.5 Answer box studies

Answer boxes have been studied empirically in a number of different ways in recent years. In a study published in 2017, 12.3% of the 112 million search queries examined produced featured snippets, and the appearance of snippets reduced user clicks to the first search result from 26.0% to 19.6% [118]. A more recent study found that shorter phrases in a search bar are more likely to generate featured snippets [65], and featured snippet sources have been found to vary by location [119]. A 2019 study found significant liberal bias in Google's news boxes [8]. This could occur because of bias in Google's algorithms or simply because left-leaning news stories are more numerous. Whatever the cause, bias in answer boxes is important because it can influence the beliefs and opinions of people who are undecided on an issue. Ludolph and colleagues [5] showed, for example, that participants who received more comprehensible information about vaccinations in a Google knowledge box subsequently proved to be more knowledgeable, less skeptical, and more critical of online information quality compared with participants who were given less comprehensive information.

1.6 The current study

In the three experiments described below, we sought to measure the impact that giving people "the answer" to one or more queries has on the opinions and voting preferences of undecided voters–an important and ever-changing group of people that has long decided the outcomes of close elections worldwide [120–122]. Experiments 1 and 2 look at the impact of answer boxes in a search engine environment, and Experiment 3 looks at the impact of answers provided by a simulation of the Alexa IPA. All three of the experiments were controlled, randomized, counterbalanced, and double-blind.

2. Experiment 1: Biased answer boxes and similarly biased search results

In our first experiment, we sought to determine whether a biased answer box (biased to favor one political candidate) could increase the shift in opinions and voting preferences produced by search results sharing the same bias. In other words, we asked whether a biased answer box could increase the magnitude of SEME [2]. We also sought to determine whether the appearance of an answer box would affect the number of search results people clicked [cf. 118] and the total time people spent searching.

2.1 Methods

2.1.1 Ethics Statement. The federally registered Institutional Review Board (IRB) of the sponsoring institution (American Institute for Behavioral Research and Technology) approved this study with exempt status under HHS rules because (a) the anonymity of participants was preserved and (b) the risk to participants was minimal. The IRB is registered with OHRP under number IRB00009303, and the Federalwide Assurance number for the IRB is FWA00021545. Informed written consent was obtained for all three experiments as specified in the Procedure section of Experiment 1.

2.1.2 Participants. After cleaning, Experiment 1 included 421 eligible voters from 49 US states whom we had recruited from Amazon's Mechanical Turk (MTurk) subject pool [123]. The data had been cleaned to remove participants who had reported an English fluency level below 6 on a 10-point scale, where 1 was labeled "not fluent" and 10 was labeled "highly fluent."

46.3% (n = 195) were male, and 53.7% (n = 226) were female. Participants ranged in age from 18 to 73 (M = 35.3, median = 33.0, SD = 10.8). 7.4% (n = 31) of the participants identified themselves as Asian, 7.4% (n = 31) as Black, 5.7% (n = 24) as Mixed, 2.1% (n = 9) as other, and 77.4% (n = 326) as White (total non-White: n = 95, 22.6%). 61.1% (n = 257) reported having received a bachelor's degree or higher.

90.5% (n = 381) of the participants said that they had previously searched online for information about political candidates, and 92.2% (n = 388) reported that Google was their most used search engine. Participants reported conducting an average of 13.6 (SD = 20.8) internet searches per day. 45.6% (n = 192) of the participants identified themselves as liberal, 27.3% (n = 115) as moderate, 24.5% (n = 103) as conservative, 1.7% (n = 7) as not political, and 1.0% (n = 4) as other.

2.1.3 Procedure. All procedures were conducted online. Participants were first asked two screening questions; sessions were terminated if they said they were not eligible to vote in the US (yes/no question) or if they said they knew a lot about politics in Australia (yes/no question). To assure participants' anonymity (a requirement of the Institutional Review Board of our sponsoring institution), we did not ask for names or email addresses.

People who passed our screening questions were then asked various demographic questions and then given instructions about the experimental procedure. At the end of the instructions page, in compliance with APA and HHS guidelines, participants clicked the continue button to indicate their informed consent to participate in the study, and were given an email address they could contact to report any problems or concerns, or, by providing their MTurk ID, to request that their data be removed from the study. Participants were then asked further questions about their political leanings and voting behavior, along with how familiar they were with the two candidates identified in the political opinion portion of the study.

Participants were randomly assigned to one of four groups: Pro-Candidate-A-with-Answer-Box, Pro-Candidate-B-with-Answer-Box, Pro-Candidate-A-No-Answer-Box, or Pro-Candidate-B-No-Answer-Box. Our candidates were Julia Gillard and Tony Abbott, actual candidates from the 2010 election for prime minister of Australia. We chose this election to assure that our participants would be "undecided" voters. On a 10-point scale from 1 to 10, where 1 was labeled "not at all" and 10 was labeled "quite familiar," our participants reported an average familiarity level of 1.79 [SD = 1.68] for Julia Gillard and 2.33 [2.03] for Tony Abbott.

All of the participants (in each of the four groups) were then shown brief, neutral biographies about each candidate (approximately 150 words each). Participants were then asked six questions about their opinions of the candidates, each on a 10-point Likert scale from "Low" to "High": whether their overall impression of each candidate was positive or negative, how likeable they found each candidate, and how much they trusted each candidate. They were then asked two questions about their voting preferences. First, on a 11-point scale from -5 to +5, with one candidate's name at each end of the scale, and with the order of the names counterbalanced from one participant to another, they were asked which candidate they would most likely vote for if they had to vote today. Finally, they were asked which of the two candidates they would actually vote for today (forced choice).

Participants were then given access to our Google.com simulator, called Kadoodle. They had up to 15 minutes to conduct research on the candidates by viewing and clicking search results, which took them to web pages, exactly as the Google search engine does. All participants had access to five pages of search results, six results per page. All search results were real (from the 2010 Australian election, obtained from Google.com), and so were the web pages to which the search results linked. Links in those web pages had been deactivated.

In the two Box groups, the bias in the answer boxes matched the bias in the search results, with higher-ranking results linking to web pages that made one candidate look better than his or her opponent. Prior to the experiment, all web pages had been rated by five independent judges on an 11-point scale from -5 to +5, with the names of the candidates at each end of the scale, to determine whether a web page favored one candidate or another. See Epstein and Robertson [2] for further procedural details.

Box content contained strongly biased language. The pro-Gillard box, for example, contained language such as: "Julia Gillard is the better candidate. Her opponent, Tony Abbott, uses 'bad language to criticise her,' but she 'has laughed off the comments.'" The pro-Abbott box contained language such as: "Tony Abbott is the better candidate. Julia Gillard, the opposing candidate, is 'clueless about what needs to be done' to improve education. . .. [Her] 'Education Revolution is a failure.'" Each box contained a link to a web page containing the content in quotation marks.

When participants chose to exit the search engine or they timed out after 15 minutes, they were asked the same six opinion questions and two voting-preference questions they had been asked before they began their research. Finally, participants were asked whether anything about the search results "bothered" them. If they answered "yes," participants could type the details of their concerns in an open-ended box. We used this inquiry to detect whether people reported seeing any bias in the search results. Participants were not asked about bias directly because leading questions tend to produce predictable and often invalid answers [124]. To assess bias we searched the textual responses for words such as "bias," "skewed," or "slanted" to identify people in the bias groups who had apparently noticed the favoritism in the search results they had been shown.

2.2 Results

The No-Box condition was, in effect, a standard SEME experiment, and it produced shifts in the direction of the favored candidates consistent with the results of previous SEME experiments [2, 15, 49], and also consistent with the results of other partial or full replications of SEME [1, 4–8]. It produced a VMP (Vote Manipulation Power, a pre-post shift in the proportion of people voting for the favored candidate) of 44.1% (Table 1), and corresponding shifts in the three opinions we measured (Table 2) (see S1 Text for details about how VMP is calculated).

In the No-Box condition, we also looked at the pre-post shift in voting preferences measured on an 11-point scale (see Methods). For this measure, preferences also shifted significantly in the predicted direction, from a mean preference of -0.08 [2.93] for favored candidates pre-search, to a mean preference of 1.88 [3.96] for favored candidates post-search (Wilcoxon $z = -8.36$, $p < 0.001$, $d = 0.56$).

The VMP in the Box condition was higher than the VMP in the No-Box condition, but the VMP increased by only 10.4% (this is a percentage increase, not the additive difference between the VMPs), and the difference was not statistically significant (Table 1). Mean search time also decreased (by 5.5%), but that difference was also not significant. The mean number

Table 1. Experiment 1: VMP, search times, and results clicked by condition.

Condition	n	VMP (%)	Mean Search Time (sec) (SD)	Mean No. of Results Clicked (SD)
No Box	208	44.1	253.9 (259.5)	4.25 (3.6)
Box	213	48.7	239.9 (236.1)	3.35 (3.6)
Change (%)	-	+10.4	-5.5	-21.2
Statistic	-	$z = -0.94$	$t(419) = -0.578$	$t(419) = -2.558$
p	-	= 0.34 NS	= 0.56 NS	< 0.05

https://doi.org/10.1371/journal.pone.0268081.t001

Table 2. Experiment 1: Pre- and post-search opinion ratings of favored and non-favored candidates.

		Favored Candidate Mean (SD)			Non-Favored Candidate Mean (SD)			
		Pre	Post	Diff	Pre	Post	Diff	$z^{\dagger}$
No Box	Impression	7.10 (1.98)	6.90 (2.24)	-0.20	7.07 (2.06)	4.42 (2.23)	-2.65	-8.66***
	Trust	6.33 (2.20)	6.29 (2.51)	-0.04	6.31 (2.25)	3.98 (2.25)	-2.33	-8.33***
	Likeability	6.98 (2.02)	6.84 (2.36)	-0.14	6.83 (2.06)	4.25 (2.30)	-2.58	-8.90***
Box	Impression	7.29 (1.97)	7.25 (2.17)	-0.04	7.24 (2.04)	4.38 (2.23)	-2.86	-9.35***
	Trust	6.31 (2.14)	6.36 (2.46)	0.05	6.27 (2.18)	4.12 (2.27)	-2.15	-8.90***
	Likeability	7.21 (1.97)	7.03 (2.24)	-0.18	7.10 (2.08)	4.34 (2.29)	-2.76	-8.50***

†z-score represents Wilcoxon signed ranks test comparing post-minus-pre ratings for the favored candidate to the post-minus-pre ratings for the non-favored candidate

$^{***}p < 0.001$

https://doi.org/10.1371/journal.pone.0268081.t002

of clicks to search results also decreased, and that difference was highly significant (Table 1, cf. 118). All three opinions (impression, trust, and likeability) shifted significantly in the predicted direction (Table 2), and so did the voting preferences as expressed on the 11-point scale ($M_{Pre\text{-}Search}$ = 0.03, $M_{PostSearch}$ = 1.92, Wilcoxon z = -8.66, $p < 0.001$, d = 0.55).

When users are shown blatantly biased search results, 20 to 30 percent of users can typically spot the bias, but that percentage drops to zero when simple masking procedures are employed [2]. (In the simplest masking procedure, a pro-Candidate-A search result is inserted into position 3 or 4 of a list of pro-Candidate-B search results.) In the present experiment, no masking procedure was employed, and 19.7% of the participants in the No-Box condition reported seeing bias in the search results. In the Box condition, more people reported seeing bias (27.2%) than in the No-Box condition, but the difference between these percentages was not significant (z = 1.82, p = 0.07 NS).

As we noted earlier, when people can spot such bias, they tend to shift even farther in the direction of the bias than people who don't see the bias, presumably because they mistakenly believe that algorithmic output is especially trustworthy. In our No-Box condition, we found the same pattern: The VMP for participants who spotted the bias was significantly larger than the VMP for participants who did not report seeing the bias (VMP_{Bias} = 68.8% [n = 41], VMP_{NoBias} = 39.5% [n = 167], z = 3.37, $p < 0.001$). In the Box condition, we again found this pattern (VMP_{Bias} = 76.9% [n = 58], VMP_{NoBias} = 40.7% [n = 155], z = 4.71, $p < 0.001$).

Demographic analyses of data from Experiment 1 –by educational level, gender, age, and race/ethnicity–are shown in S1–S4 Tables. Demographic effects were relatively small.

3. Experiment 2: Biased answer boxes and unbiased search results

The results of Experiment 1 suggest that a biased answer box can increase the shift in opinions and voting preferences produced by similarly biased search results, but the increases we found were small. Could this be a ceiling effect? In other words, were the biased search results masking the power that biased answer boxes have to change thinking or behavior? To answer this question, we conducted an experiment in which participants saw either no answer boxes or biased answer boxes and in which search results were neutral for all groups. This experiment was controlled, randomized, counterbalanced, and double-blind.

3.1 Methods

3.1.1 Participants. After cleaning, Experiment 2 included 177 eligible US voters from 44 states who had been recruited through the MTurk subject pool. The data had been cleaned to include only participants who had reported an English fluency score of 6 or above on a 10-point scale.

52.0% (n = 92) were male, and 48.0% were female (n = 85). Participants ranged in age from 18 to 67 (M = 34.3, median = 32.0, SD = 10.4). 5.1% (n = 9) of the participants identified themselves as Asian, 9.0% (n = 16) as Black, 4.5% (n = 8) as Mixed, 4.0% (n = 7) as other, and 77.4% (n = 137) as White (total non-White: n = 40, 22.6%). 50.3% (n = 89) reported having received a bachelor's degree or higher.

92.1% (n = 163) of the participants said that they had previously searched online for information about political candidates, and 94.4% (n = 167) reported that Google was their most used search engine. Participants reported conducting an average of 18.1 (SD = 34.1) internet searches per day. 49.2% (n = 87) of the participants identified themselves as liberal, 32.2% (n = 57) as moderate, 14.1% (n = 25) as conservative, 2.3% (n = 4) as not political, and 2.3% (n = 4) as other.

3.1.2 Procedure. Participants were randomly assigned to one of three groups: Pro-Candidate-A-Box, Pro-Candidate-B-Box, or a control group in which the answer box was not present. We used the same candidates and election as we used in Experiment 1, except that search results were unbiased in all three groups. Specifically, pro-Abbott search results alternated with pro-Gillard search results. Our participants reported an average familiarity level of 1.68 [1.64] for Julia Gillard and 2.23 [2.06] for Tony Abbott. The experimental procedure itself was identical in all respects to the procedure in Experiment 1.

3.2 Results

In the No-Box group, the proportions of people voting for each candidate did not change pre-search to post-search ($Pre_{Gillard} = 0.41$, $Post_{Gillard} = 0.52$, $z = -1.19$, $p = 0.23$). The VMP itself could not be computed, because there was no bias condition in this group. Voting preferences expressed on the 11-point scale shifted from -0.02 [3.24] pre-search to 0.24 [3.30] post-search (Wilcoxon's $z = -0.60$, $p = 0.55$ NS, $d = 0.08$), which means that unbiased search results had almost no effect on votes or voting preferences.

In the Box conditions, however, the VMP was 38.6% ($z = -5.50$, $p < 0.001$) (Table 3), and the voting preference expressed on the 11-point scale shifted from 0.08 [3.06] to 0.97 [3.90] (Wilcoxon's $z = -3.57$, $p < 0.001$, $d = 0.26$), which means there was a significant shift toward the favored candidate. Given that there was no bias in the search results, the shift in voting preferences was likely due exclusively to the biased answer boxes. Similarly, more people reported seeing bias in the box condition (12.5%) than in the No-Box condition (0.0%), and the difference between these percentages was significant ($z = -2.20$, $p < 0.05$).

The results in Experiment 2 differ from the results in Experiment 1 in one important respect: The opinions about the candidates (impression, trust, and likeability) did not change

Table 3. Experiment 2: VMP, search times, and results clicked by condition.

Condition	n	VMP (%)	Mean Search Time (sec) (SD)	Mean No. of Results Clicked (SD)
No Box	58	N/A[†]	228.0 (201.2)	4.00 (3.7)
Box	119	38.6	246.1 (265.9)	3.45 (3.2)
Change (%)	-	-	+7.9	-13.8
Statistic	-	-	$t(175) = 0.46$	$t(175) = -1.01$
p	-	-	= 0.65 NS	= 0.31 NS

[†]As noted in the text, since there was no bias in the search results shown in the No-Box condition, VMP could not be calculated.

https://doi.org/10.1371/journal.pone.0268081.t003

Table 4. Experiment 2: Pre- and post-search opinion ratings of favored and non-favored candidates.

		Pre	Post	Diff				
No Box	Impression	7.46 (1.87)	6.34 (2.11)	-1.12				
	Trust	6.29 (2.06)	5.82 (2.22)	-0.47				
	Likeability	7.41 (1.96)	6.47 (2.10)	-0.94				
		Favored Candidate Mean (SD)			Non-Favored Candidate Mean (SD)			
		Pre	Post	Diff	Pre	Post	Diff	$z^{\dagger}$
Box	Impression	7.07 (1.93)	5.93 (2.31)	-1.14	7.31 (1.88)	5.55 (2.28)	-1.76	-2.06 NS
	Trust	6.24 (2.26)	5.60 (2.54)	-0.64	6.38 (2.23)	5.17 (2.29)	-1.15	-2.18 NS
	Likeability	7.03 (2.07)	5.82 (2.34)	-1.21	7.20 (1.88)	5.46 (2.31)	-1.74	-1.61 NS

†z-score represents Wilcoxon signed ranks test comparing post-minus-pre ratings for the favored candidate to the post-minus-pre ratings for the non-favored candidate. This statistic could not be computed for Group 1 because there was no favored candidate.

https://doi.org/10.1371/journal.pone.0268081.t004

significantly (Table 4). This makes sense, given that (a) the answer boxes gave almost no information about the candidates and (b) the search results did not favor either candidate. Differences in opinions did not emerge even though people spent about the same time viewing search results in Experiment 1 as they did in Experiment 2 (M_{E1} = 246.8 s [247.8], M_{E2} = 240.2 s [246.2], $t(596) = 0.30$, $p = 0.77$, $d = 0.03$), and clicked roughly the same number of search results in Experiment 1 as they clicked in Experiment 2 (M_{E1} = 3.80 [3.6], M_{E2} = 3.63 [3.4], $t(596) = 0.51$, $p = 0.61$, $d = 0.05$).

We also saw a different pattern in the VMPs of the people in the two box groups who detected the bias (23 out of 119 people, 19.3%): When people detect bias in search results (based largely or in part on viewing the web pages to which the search results link), their opinions and voting preferences tend to shift even farther in the direction of the favored candidate than do the opinions and voting preferences of people who do not detect the bias. In Experiment 2, however, we found the opposite pattern. The VMP for people who reported seeing bias in the Box groups was 12.5%; whereas the VMP for people who did not report seeing bias in the Box groups was 44.4% ($z = -2.93$, $p < 0.05$). Bear in mind that each user is seeing only one box; he or she has nothing with which to compare it, and the search results themselves are unbiased. More light is shed on this matter in Experiment 3 (also see Discussion).

The dramatic shift in voting preferences produced by biased answer boxes alone in Experiment 2 raises a disturbing possibility about the power that IPAs might have to impact thinking and behavior. Experiment 2 functioned, after all, like an IPA: A single query produced a single reply (given in the answer box), which appeared above unbiased search results. Could a single biased answer produced by an IPA produce a large shift in opinions and voting preferences? And what if multiple questions produced answers that shared the same bias? Could they produce even larger shifts in opinions and voting preferences? We attempted to answer these questions in Experiment 3.

Demographic analyses of data from Experiment 2 –by educational level, gender, age, and race/ethnicity–are shown in S5–S8 Tables. Demographic effects were relatively small.

4. Experiment 3: Assessing the persuasive power of the intelligent personal assistant (IPA)

4.1 Methods

4.1.1 Participants. After cleaning, our sample for this experiment consisted of 1,138 eligible voters from 48 US states. They were recruited from the MTurk subject pool. The data had been cleaned to remove participants who had reported an English fluency level below 6 on a 10-point scale.

52.3% (n = 595) were male, 46.7% (n = 531) were female, and 1.1% (n = 12) chose not to identify their gender. Participants ranged in age from 18 to 89 ($M = 41.3$, median = 39.0, $SD = 12.9$). 8.3% (n = 94) of the participants identified themselves as Asian, 8.1% (n = 92) as Black, 3.0% (n = 34) as Mixed, 2.3% (n = 26) as other, and 78.4% (n = 892) as White (total non-White: n = 246, 21.6%). 64.1% (n = 729) reported having received a bachelor's degree or higher.

86.6% (n = 986) of the participants reported they had used a virtual assistant like Alexa or Siri. 48.6% (n = 553) of the participants identified themselves as liberal, 27.2% (n = 310) as moderate, 21.4% (n = 244) as conservative, 1.7% (n = 19) as not political, and 1.1% (n = 12) as other.

4.1.2 Procedure. All procedures were run online and were compatible with both desktop and mobile devices. As in the earlier experiments, participants were first asked screening questions and demographic questions and then given instructions about the experimental procedure and asked for their consent to participate in the study.

Participants were randomly assigned to one of five different question/answer (Q/A) groups. Each group was shown the same list of 10 questions, and the order of the questions did not vary. After a participant clicked a question, Dyslexa–our Amazon Alexa IPA simulator–replied vocally with an answer (See S2 Text). The number of questions people were required to ask varied by group, and in two of the groups, the answer to the second question was "masked" in a manner that we will describe below. A screenshot showing how the questions and Dyslexa simulator appeared to users is shown in Fig 1. The five groups were as follows:

1. Group 1Q/1A: Participants were required to select just one question.
2. Group 4Q/4A/NM: Participants were required to select four different questions, and none was masked (NM = "no mask").
3. Group 4Q/4A/M2: Participants were required to select four different questions, and the answer to Question 2 was masked (M2 = Question 2 mask).

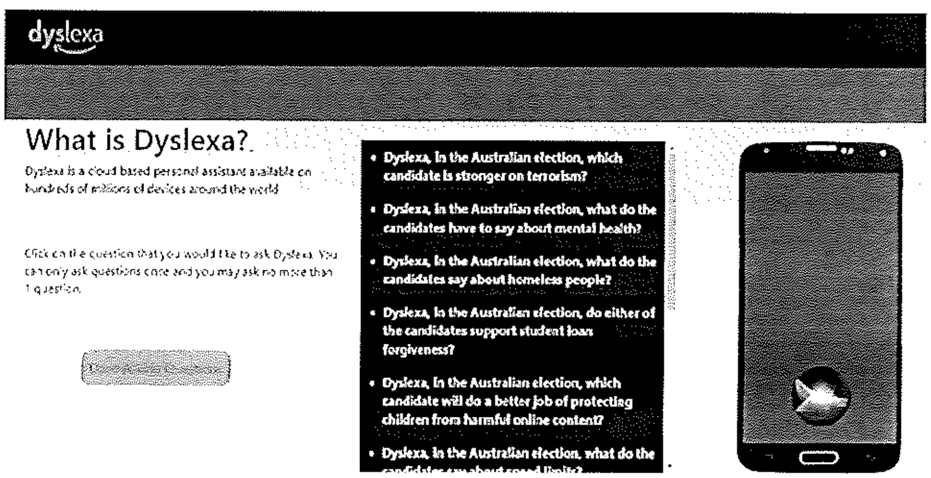

Fig 1. A screenshot showing what users saw in Experiment 3 when they posed questions to Dyslexa. Different groups were required to ask 1, 4, or 6 questions. After clicking on a question, it was greyed out, and Dyslexa answered the question orally. While it was speaking, the circular graphic at the bottom of the phone screen glowed and swirled, just as similar graphics do on most iPhones.

https://doi.org/10.1371/journal.pone.0268081.g001

4. Group 6Q/6A/NM: Participants were required to select six different questions, and none was masked.

5. Group 6Q/6A/M2: Participants were required to select six different questions, and the answer to Question 2 was masked.

Within each of the five groups, participants were randomly assigned to one of three different candidate conditions: Pro-Candidate-A, Pro-Candidate-B, or a control group. Our political candidates were Scott Morrison (Candidate A) and Bill Shorten (Candidate B), actual candidates from the 2019 election for prime minister of Australia. We chose this election to assure that our participants would be "undecided" voters. On a 10-point scale from 1 to 10, where 1 was labeled "not at all" and 10 was labeled "quite familiar," our participants reported an average familiarity level of 1.14 [0.43] for Scott Morrison and 1.05 [0.26] for Bill Shorten.

In the Candidate A condition, the answers were biased in favor of Scott Morrison. For example, when asked, "Dyslexa, in the Australian election, which candidate favors having a stronger relationship with the United States?," Dyslexa replied, "According to recent media reports, Scott Morrison wants to build a stronger relationship with the United States. His opponent, Bill Shorten, wants to continue to increase trade with Russia and China." In the Candidate B condition, the answers were biased in favor of Bill Shorten. In response to the same question, the pro-Shorten reply was "According to recent media reports, Bill Shorten wants to build a stronger relationship with the United States. His opponent, Scott Morrison, wants to continue to increase trade with Russia and China." The answers in each bias group were, in other words, nearly identical; only the names were changed. Mean bias ratings were obtained from five independent raters for each of the 20 answers on an 11-point scale from -5 (pro-Morrison) to +5 (pro-Shorten). The overall bias for Morrison was -3.3 [0.67], and the overall bias for Shorten was 3.4 [0.67] (based on absolute value: $t(18) = -0.07$, $p = 0.98$ NS).

In two of the five groups (Groups 3 and 5), masks were used for the answers to the second question each participant asked. This means that in the pro-Morrison group, a pro-Shorten answer was given in response to the second question asked, and in the pro-Shorten group, a pro-Morrison answer was given in response to the second question asked. This is a standard procedure used in SEME experiments [2] to reduce or eliminate the perception that the content being shown is biased. In SEME experiments, biased search results still produce large shifts in opinions and voting preferences even when aggressive masks are employed that completely eliminate the perception of bias. (See the Results and Discussion sections below for further information about our use of masks.)

In each control group, including Group 1 (1Q/1A), the answer to the first question had a 50/50 chance of supporting either Morrison or Shorten. After that, the bias in the answers alternated between the two candidates with each question asked. In Groups 2 through 5, we used an even number of questions (4 or 6) to ensure that each participant received equal exposure to pro-Morrison and pro-Shorten answers.

Participants were allowed to choose their questions from a list of 10. We provided this relatively long list to increase the likelihood that participants would select questions on topics they cared about. We speculated that allowing people to choose their questions would increase their interest in the answers they were given. We varied the number of questions people could ask to see whether we could have a bigger impact on opinions and voting preferences when people were exposed to a larger number of biased answers. We did not include a two-question group because we would not have been able to use a mask; a mask in the second position would almost certainly have eliminated the bias effect.

Following the demographic questions and instructions, all participants were shown brief, neutral biographies about each candidate (approximately 120 words each–somewhat shorter than the biographies used in Experiments 1 and 2 for the 2010 Australian election). (See S3 Text for the biographies employed in Experiment 3.) Participants were then asked six questions about their candidate preferences (each on a 10-point Likert scale from "Low" to "High"): whether their overall impression of each candidate was positive or negative, how likeable they found each candidate, and how much they trusted each candidate. Then–on an 11-point scale from -5 to +5, with the name of each candidate shown at either end of the scale and with the order of the names counterbalanced from one participant to another–participants were asked which candidate they would most likely vote for if they had to vote today. Finally, they were asked which of the two candidates they would actually vote for today (forced choice). The answers to these two questions had to be consistent; if they weren't, participants were asked to answer them again.

Following these opinion questions, participants were given brief instructions about how to use our IPA, and they then could proceed to ask questions (between one and six questions, according to their group assignment) and hear Dyslexa's answers. Our questions covered a wide range of topics that we thought would be of interest to a US sample (see S2 Text), but we deliberately avoided including hot-button issues such as abortion. If a participant chose to ask, "What are the candidates' positions on abortion?," and Dylexa replied that Morrison wanted to protect abortion rights, the possible partisanship of our participants could have driven them either *toward* or *away from* Morrison–*toward* if they supported abortion rights, *away* if they opposed abortion.

Following the interaction with the IPA, all participants were again asked those six opinion questions and two voting-preference questions. Finally, participants were asked whether anything "bothered" them about the questions they were shown and the answers they heard while interacting with our IPA. As in our previous experiments, this is where participants had an opportunity to express their concerns about content bias or other issues.

4.2 Results

We found significant and substantial shifts in both voting preferences (Table 5) and opinions (Table 6) in the direction of the favored candidates in all bias groups. We also found significant shifts in voting preferences in the direction of the favored candidates in all bias groups as expressed on our 11-point voting-preference scale (Table 7). In contrast, in the control groups the proportions of people voting for each candidate before the manipulations changed relatively little or not at all following the manipulations (Group 1, 0.0%; Group 2, 6.6%; Group 3, 2.7%; Group 4, 7.1%; Group 5, 6.8%).

The percentage of people in the bias groups who reported seeing biased content was substantially lower when they received just one answer (Group 1, 4.9%) or when biased content was masked (Group 3, 5.1%; Group 5, 7.1%) than when people saw multiple biased answers

Table 5. Experiment 3: Pre- and Post-IPA VMPs.

Group No.	Group	Total n	Bias Groups n	Bias Groups VMP (%)	McNemar Test X^2	p
1	1Q/1A	222	142	43.8	24.0	< 0.001
2	4Q/4A/NM	229	153	59.5	35.9	< 0.001
3	4Q/4A/M2	230	156	59.2	33.6	< 0.001
4	6Q/6A/NM	230	145	65.8	44.5	< 0.001
5	6Q/6A/M2	227	154	50.0	36.5	< 0.001

https://doi.org/10.1371/journal.pone.0268081.t005

Table 6. Experiment 3: Pre- and post-IPA opinion ratings of favored and non-favored candidates.

		Favored Candidate Mean (SD)			Non-Favored Candidate Mean (SD)			
		Pre	Post	Diff	Pre	Post	Diff	$z^{\dagger}$
Group 1: 1Q1A Condition	Impression	7.13 (1.85)	7.63 (2.00)	+0.50	7.10 (1.73)	6.13 (2.18)	-0.97	-6.32***
	Trust	6.29 (2.20)	6.95 (2.29)	+0.66	6.26 (2.11)	5.65 (2.41)	-0.61	-6.59***
	Likeability	7.15 (1.83)	7.46 (2.00)	+0.31	7.18 (1.72)	6.18 (2.23)	-1.00	-6.43***
Group 2: 4QNM Condition	Impression	6.76 (1.93)	7.73 (2.23)	+0.97	6.89 (1.72)	4.97 (2.04)	-1.92	-8.82***
	Trust	5.88 (2.18)	6.97 (2.51)	+1.09	6.05 (2.05)	4.80 (2.23)	-1.25	-7.80***
	Likeability	6.67 (2.01)	7.41 (2.26)	+0.74	6.93 (1.84)	5.03 (2.13)	-1.90	-7.93***
Group 3: 4QM2 Condition	Impression	6.79 (1.92)	7.28 (1.95)	+0.49	6.96 (1.72)	6.12 (1.85)	-0.84	-5.92***
	Trust	5.81 (2.12)	6.54 (2.27)	+0.73	6.06 (2.07)	5.71 (2.04)	-0.35	-7.50***
	Likeability	6.81 (1.90)	7.13 (2.12)	+0.32	7.04 (1.71)	6.20 (1.99)	-0.84	-5.64***
Group 4: 6QNM Condition	Impression	6.87 (1.75)	7.74 (1.94)	+0.87	6.72 (1.81)	4.83 (2.00)	-1.89	-8.64***
	Trust	5.94 (1.97)	6.90 (2.25)	+0.96	5.99 (2.10)	4.58 (2.11)	-1.41	-7.87***
	Likeability	6.82 (1.87)	7.62 (2.09)	+0.80	6.78 (2.02)	4.96 (2.13)	-1.82	-8.32***
Group 5: 6QM2 Condition	Impression	7.10 (1.65)	7.65 (1.94)	+0.55	7.00 (1.87)	5.34 (2.02)	-1.66	-7.98***
	Trust	6.31 (2.00)	7.09 (2.20)	+0.78	6.18 (2.07)	5.08 (2.29)	-1.10	-7.65***
	Likeability	7.05 (1.70)	7.50 (2.00)	+0.45	6.93 (1.86)	5.42 (2.12)	-1.51	-7.54***

†z-score represents Wilcoxon signed ranks test comparing post-minus-pre ratings for the favored candidate to the post-minus-pre ratings for the non-favored candidate.

$^{***}p < 0.001$

https://doi.org/10.1371/journal.pone.0268081.t006

without masks (Group 2, 23.5%; Group 4, 40.7%) (Table 8) ($M_{\text{Groups1,3,5}}$ = 5.8%, $M_{\text{Groups2,4}}$ = 31.9%, z = -9.50, $p < 0.001$).

The present study sheds new light on the role that bias detection plays in shifting opinions and voting preferences. Previous investigations have shown that the opinions of the few people who are able to detect bias in search results shift even farther in the direction of the bias than the opinions of the people who don't see the bias [2, 15]. This occurs presumably because of the high trust people have in the filtering and ordering of search results, which people mistakenly believe is an objective and impartial process [125, 126]. In the present study, we learned that bias detection *erodes* trust when people are interacting with answers provided by answer boxes (in the absence of biased search results–see Experiment 2) or the vocal answers of an IPA, where search results are entirely absent (Experiment 3). This difference is likely due to the daily regimen of operant conditioning that supports the almost blind trust people have in search results. About 86% of searches are for simple facts, and the correct answers to those queries reliably turn up in the first or second search result. People are learning, over and over again, that what is higher in the list of search results is better and truer than what is lower. When, in a recent experiment, that trust was temporarily broken, the VMP in a SEME procedure was significantly reduced [15].

Table 7. Experiment 3: Pre-IPA vs. Post-IPA voting preferences on 11-point scales.

Group No.	Group	Pre-IPA Voting Preference on 11-Point Scale (SD)	Post-IPA Voting Preference on 11-Point Scale (SD)	z	p	d
1	1Q/1A	0.61 (2.42)	1.70 (2.76)	-5.51	< 0.001	0.42
2	4Q/4A/NM	-0.01 (2.57)	2.41 (2.64)	-8.17	< 0.001	0.93
3	4Q/4A/M2	-0.10 (2.76)	1.38 (2.90)	-5.83	< 0.001	0.52
4	6Q/6A/NM	0.21 (2.46)	2.67 (2.28)	-8.50	< 0.001	1.04
5	6Q/6A/M2	0.20 (2.60)	2.26 (2.62)	-7.99	< 0.001	0.79

https://doi.org/10.1371/journal.pone.0268081.t007

Table 8. Experiment 3: VMPs for people who saw Bias vs. VMPs for people who did not see Bias.

Group No.	Group	*n*	No. Ss in Bias Groups Reporting Bias in IPA Content (%)	No. Ss in Bias Groups Not Reporting Bias in IPA Content (%)	VMP for Ss Who Reported Bias (%)	VMP for Ss Who Did Not Report Bias (%)	*z*	*p*
1	1Q/1A	142	7 (4.9)	135 (95.1)	33.3†	44.3	-0.57	= 0.57 NS
2	4Q/4A/NM	153	36 (23.5)	117 (76.5)	21.7	75.0	-5.78	< 0.001
3	4Q/4A/M2	156	8 (5.1)	148 (94.9)	300.0†	55.7	14.46	< 0.001
4	6Q/6A/NM	145	59 (40.7)	86 (59.3)	63.3	67.4	-0.51	= 0.61 NS
5	6Q/6A/M2	154	11 (7.1)	143 (92.9)	60.0†	49.4	0.68	= 0.50 NS

†The validity of these VMPs is questionable because they are based on a small number of observations. In Groups 1, 3, and 5, respectively, only 7, 8, and 11 people reported seeing bias in the IPA replies.

https://doi.org/10.1371/journal.pone.0268081.t008

So when search results are absent, as they are when people are using IPAs, or when search results are unbiased, as they were in our Experiment 2, people who detect bias do not automatically accept that bias as valid. Accepting that bias as valid seems to occur primarily when people are being influenced by biased search results–again, presumably because of that daily regimen of operant conditioning. That daily regimen of conditioning makes SEME a unique list effect and an especially powerful form of influence [15].

As we noted earlier, we regard the most important measure of change to be the VMP, which indicates the increase or decrease in the proportion of people who indicated in response to a forced-choice question which candidate they would vote for if they had to vote today (see S1 Text). The VMPs in the five groups in Experiment 3 ranged from 43.8% (Group 1) to 65.8% (Group 4). These shifts were all quite high–all higher than the 38.6% shift we found in Experiment 2.

In addition, we found that the more questions people asked (without masks, which tend to lower VMPs), the greater the shift in voting preferences ($VMP_{Q1/A1}$ = 43.8%, $VMP_{Q4/A4/NM}$ = 59.5%, $VMP_{Q6/A6/NM}$ = 65.8%; $X^2 = 6.59$; $p < 0.05$).

A breakdown of VMP data from Experiment 3 based on whether participants had had previous experience with IPAs is shown in S9 Table. Previous experience with IPAs did not appear to impact VMPs in any consistent way.

5. Discussion

Together, the three experiments we have described reveal a dangerous new tool of mass manipulation–one that is, at this writing, controlled worldwide almost entirely by just four large American tech companies: Amazon, Apple, Facebook/Meta, and Google. This new tool, which we call the Answer Bot Effect (ABE), is likely now affecting hundreds of millions of people, and with more and more people coming to rely on electronic devices to give them a single answer to their queries, the number of people affected by ABE will likely swell into the billions within the next few years. ABE should be of concern to every one of us, but especially to parents–whose children are being fed algorithmically-generated answers every day on their computers, mobile phones, tablets, and toys–as well as to public policy makers.

ABE should be of special concern for four reasons: (a) because of the large magnitude of the effect, (b) because it can impact the vast majority of people without their awareness, (c) because it is an ephemeral manipulation, leaving no paper trail for authorities to trace, and (d) because ABE is inherently non-competitive and impossible to counteract. You can counteract a billboard or television commercial, but how can you correct the way a tech platform adjusts its algorithms? Recall that in Experiment 3, a one-question-one-answer interaction on our Alexa simulator produced a 43.8% shift in voting preferences, with only 4.7% of the participants reporting any concerns about bias.

Perhaps the reader thinks we are overstating the seriousness of the problem. Although a full exploration of this issue is beyond the scope of this paper, please consider just two growing bodies of evidence that bring manipulations like ABE into sharper focus: First, in recent years, whistleblowers from Google and Facebook/Meta, along with leaks of emails, documents, and videos from these companies, have shown repeatedly that manipulations like ABE are being deliberately and strategically used by these companies to influence attitudes, beliefs, purchases, voting preferences, and public policy itself [25, 28, 29, 43, 48]. In a leak of emails to the *Wall Street Journal* in 2018, Google employees discuss the possibility of using "ephemeral experiences" to change people's views about Trump's 2017 travel ban [25]. A leaked 8-minute video from Google called "The Selfish Ledger" describes the company's power to "modify behavior" at the "species level" in ways that "reflect Google's values" [127]. In various interviews and the recent documentary film, "The Social Dilemma," former Google insider Tristan Harris spoke about his time working with a large team of Google employees whose job it was to modify "a billion people's attention and thoughts every day" [128].

Harris and others have expressed concerns about company policies that are meant to influence people in specific ways, but ABE, SEME, and other new forms of online influence will impact thinking and behavior even without a company policy in place. Algorithms left to their own devices–let's call this practice "algorithmic neglect"–reflect the biases of the people who programmed them [20–23], and the algorithms also quickly learn and reflect the foibles of human users, sometimes magnifying and spreading bigotism, racism, and hatred with frightening rapidity [52, 55, 61, 97, 116, 117]. What's more, a single rogue employee with the right password authority or hacking skills can use a large tech platform like Google to impact reputations, businesses, or elections on a large scale without senior management knowing he or she is doing so [129]. When authorities learned in 2010 that Google's Street View vehicles had been vacuuming up personal Wi-Fi data for 3 years in 30 countries [130], Google blamed the entire operation on a single software engineer, Marius Milner–but they did not fire him, and he remains at the company today [131].

Second, election monitoring projects that have been conducted since 2016 have so far preserved more than 1.5 million politically-related online ephemeral experiences in the weeks preceding national elections in the US. This is actual content–normally lost forever–being displayed on the computer screens of thousands of US voters–the real, personalized content that Big Tech companies are showing politically diverse groups of people as elections approach. The wealth of unusual data preserved in these projects has revealed strong unilateral political bias in ephemeral content, sufficient to have shifted millions of votes in national elections in the US without people's knowledge [132–134].

The experiments we have described build one upon the other. Experiment 1 showed that when the content of an answer box shared the bias of the search results beneath it, it increased the impact that those search results have on thinking and behavior, and it reduced the time people spent searching and significantly reduced the number of search results people clicked. Experiment 2 simulated a situation in which the answer box was biased but the search results were not. The biased answer boxes alone produced a remarkable VMP of 38.6%.

Rounded to the nearest whole number, the VMP in Experiment 2 was 39%. This means that out of 100 undecided voters–people whose vote would normally split 50/50 without having additional information–the votes, on average, of 19.5 people (0.39 x 50) can be shifted by biased answer boxes, yielding a vote of roughly 69 to 30, for a win margin among previously undecided voters of 39% (see S1 Text). In a national election in the US in which 150 million people vote (159 million voted in the 2020 Presidential election), even if only 10% of the voters were undecided and depended on computers for trustworthy answers, if the single-answer-generating algorithms in the days or weeks leading up to Election Day all favored the same candidate, that could conceivably shift more than 2.9 million votes to that candidate (0.10 x 0.39 x 0.5 x 150,000,000). If the other 90% of the voters were split 50/50, that would give the favored candidate a win margin of 5.8 million votes (3.8%).

Unfortunately, the real situation we face is probably worse than the case we just described. At this moment in history, in the US virtually all the single-answer-generating algorithms will likely be supporting the same national and state candidates [135–137], and six months before an election, the percentage of undecided voters might be as high as 60%, not 10% [122, 138, 139].

Bear in mind also that in our experiments we are interacting with our participants only briefly and only once. If undecided voters are subjected to content having the same bias repeatedly over a period of weeks or months, their voting preferences will likely shift even farther than the voting preferences of our participants shifted. Recall that in Experiment 3 the VMP exceeded 65% when people asked six questions–nearly 50% higher than the VMP we found when people asked only one question (Table 5).

What's more, ABE is just one powerful source of influence. When similarly biased content is delivered in search results, search suggestions, YouTube videos, newsfeeds, targeted messages, and so on, the net impact of these manipulations is likely additive, and when Big Tech companies all share the same political bias (or any other type of bias, for that matter), the net impact of their combined influence is also likely additive. Without regulations, laws, and permanent, large-scale monitoring systems to stop them–and none exist at this writing [140]–Big Tech companies indeed have the power to reengineer humanity "at the species level," as Google's "Selfish Ledger" video suggests [127]. At the very least, they can easily tilt the outcomes of close elections worldwide.

In a remarkable and frequently quoted farewell speech delivered by US President Dwight D. Eisenhower just a few days before John F. Kennedy's inauguration in January 1961, Eisenhower–a military insider–not only warned the American people about a rapidly evolving "military-industrial complex," he also spoke of the danger that someday "public policy could itself become the captive of a scientific technological elite" [141]. If ABE, SEME, and other new forms of influence the internet has made possible work anything in the real world like they do in controlled experiments, it is not unreasonable to speculate that while humanity was being distracted by online video games, dating websites, and cat memes, Eisenhower's prediction came true. The technological elite now exist [142],and, if our analyses are correct, they are now very much in control.

Supporting information

S1 Fig. Apparent bias in a Google answer box, screenshotted October 22, 2017. The content of the box clearly favors the Google service.
(TIF)

S2 Fig. Apparent bias in two types of Google answer boxes. (a) In a screenshot preserved in an article in *Search Engine Land* on March 5, 2017, four US presidents are incorrectly listed in a Google answer box as members of the Ku Klux Klan. (b) In a screenshot of a Google knowledge box preserved in an article in *VICE* on May 31, 2018, Nazism is incorrectly listed as part of the ideology of the California Republican Party. (c) In a Google answer box captured by the

first author on August 16, 2016, Hillary Clinton's photograph is shown in response to the question, "when is the election?".
(TIF)

S1 Text. Vote Manipulation Power (VMP) calculation.
(DOCX)

S2 Text. Experiment 3: Alexa simulator, "Dyslexa," questions and answers.
(DOCX)

S3 Text. Experiment 3: Candidate biographies.
(DOCX)

S1 Table. Experiment 1: Demographic analysis by educational attainment.
(DOCX)

S2 Table. Experiment 1: Demographic analysis by gender.
(DOCX)

S3 Table. Experiment 1: Demographic analysis by age.
(DOCX)

S4 Table. Experiment 1: Demographic analysis by race/ethnicity.
(DOCX)

S5 Table Experiment 2: Demographic analysis by educational attainment.
(DOCX)

S6 Table. Experiment 2: Demographic analysis by gender.
(DOCX)

S7 Table. Experiment 2: Demographic analysis by age.
(DOCX)

S8 Table. Experiment 2: Demographic analysis by race/ethnicity.
(DOCX)

S9 Table. Experiment 3: Demographic analysis by previous IPA use.
(DOCX)

Acknowledgments

We thank J. Martinez for assistance in conducting the second experiment and L. Kafader for expert programming assistance. R. Mohr is currently a doctoral candidate at Palo Alto University, Palo Alto, California USA.

Author Contributions

Conceptualization: Robert Epstein.

Formal analysis: Robert Epstein, Vivian Lee, Roger Mohr, Vanessa R. Zankich.

Investigation: Robert Epstein, Roger Mohr.

Methodology: Robert Epstein.

Project administration: Robert Epstein.

Supervision: Robert Epstein.

Writing – original draft: Robert Epstein.

Writing – review & editing: Robert Epstein, Vivian Lee, Vanessa R. Zankich.

References

1. Allam A, Schulz PJ, Nakamoto K. The impact of search engine selection and sorting criteria on vaccination beliefs and attitudes: Two experiments manipulating Google output. J Med Internet Res, 2014; 16(4):e100. https://doi.org/10.2196/jmir.2642 PMID: 24694866
2. Epstein R, Robertson RE. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. Proc Natl Acad Sci USA, 2015; 112(33):E4512– E4521. https://doi.org/10.1073/pnas.1419828112 PMID: 26243876
3. Epstein R, Robertson RE. Suppressing the Search Engine Manipulation Effect (SEME). Proceedings of the ACM on Human-Computer Interaction, 2017; 1(CSCW):1–22. https://doi.org/10.1145/3134677
4. Haas A, Unkel J. Ranking versus reputation: perception and effects of search result credibility. Behav Inf Technol, 2017; 36(12):1285–1298. https://doi.org/10.1080/0144929X.2017.1381166
5. Ludolph R, Allam A, Schulz PJ. Manipulating Google's knowledge box to counter biased information processing during an online search on vaccination: Application of a technological debiasing strategy. J Med Internet Res, 2016; 18(6):e137. https://doi.org/10.2196/jmir.5430 PMID: 27255736
6. Eslami M, Vaccaro K, Karahalios K, Hamilton K. "Be Careful; Things Can Be Worse than They Appear": Understanding Biased Algorithms and Users' Behavior Around Them in Rating Platforms. Proceedings of the 11th International AAAI Conference on Web and Social Media [Internet]. 2017 May 3; 11(1):62–71. Available from: https://ojs.aaai.org/index.php/ICWSM/article/view/14898
7. Pogacar FA, Ghenai A, Smucker MD, Clarke CLA. The positive and negative influence of search results on people's decisions about the efficacy of medical treatments. Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval [Internet]. 2017 Oct 1–4;:209–216. Available from: https://dl.acm.org/doi/10.1145/3121050.3121074
8. Trielli D, Diakopoulos N. Search as news curator: The role of Google in shaping attention to news information. CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems [Internet]. 2019 May 4–9; 453:1–15. Available from: https://dl.acm.org/doi/10.1145/3290605.3300683
9. Casara BGS, Suitner C, Bettinsoli ML. Viral suspicions: Vaccine hesitancy in the Web 2.0. J Exp Psychol Appl, 2019; 25(3):354–371. https://doi.org/10.1037/xap0000211 PMID: 30816729
10. Edelman B. Bias in search results? Diagnosis and response. Indian J Law Technol, 2011; 7:16–32. Available from: https://www.ijlt.in/journal/bias-in-search-results%3F%3A-diagnosis-and-response
11. Feldman S. Americans see search engines as biased. 2018 Sep 7. In: Statista [Internet]. Available from: https://www.statista.com/chart/15385/americans-see-search-enginges-as-biased/
12. Knobloch-Westerwick S, Mothes C, Johnson BK, Westerwick A, Donsbach W. Political online information searching in Germany and the United States: Confirmation bias, source credibility, and attitude impacts. J Commun, 2015; 65(3):489–511. https://doi.org/10.1111/jcom.12154
13. O'Neil C. Weapons of math destruction: How big data increases inequality and threatens democracy. Crown Publishing Group, New York, 2016.
14. Pierce DR, Redlawsk DP, Cohen WW. Social influences on online political information search and evaluation. Polit Behav, 2016; 39:651–673. https://doi.org/10.1007/s11109-016-9374-4
15. Epstein R. Manipulating minds: The power of search engines to influence votes and opinions. In: Moore M, Tambini D, editors. Digital dominance: The power of Google, Amazon, Facebook, and Apple. Oxford, UK: Oxford University Press; 2018. pp. 294–319. Available from: https://aibrt.org/downloads/EPSTEIN_2018-Manipulating_minds-The-power_of_search_engines_to_influence_votes_and_opinions-UNCORRECTED_PROOFS.pdf
16. Brisson-Boivin K, McAleese S. Algorithmic awareness: Conversations with young Canadians about artificial intelligence and privacy. MediaSmarts. 2021. Available from: https://mediasmarts.ca/sites/default/files/publication-report/full/report_algorithmic_awareness.pdf
17. Fallows D. Internet searchers are confident, satisfied and trusting–but they are also unaware and naïve. 2005 Jan 23. In: Pew Internet & American Life Project [Internet]. Available from: https://www.pewresearch.org/internet/wp-content/uploads/sites/9/media/Files/Reports/2005/PIP_Searchengine_users.pdf.pdf
18. Schofield J. A lot of Brits don't understand search engines: A UK survey published by FastHosts has revealed some misconceptions about how search engines work. Whether this matters to users, as opposed to website promoters, is less clear. The Guardian. 2008 Dec 12. Available from: https://www.theguardian.com/technology/blog/2008/dec/12/searchengine-survey

19. Pasquale F. The black box society: The secret algorithms that control money and information. Harvard University Press; 2015.

20. Bhamore S. Decrypting Google's search engine bias case: Anti-trust enforcement in the digital age. Christ University Law Journal, 2019; 8(1):37–60. https://doi.org/10.12728/culj.14.2

21. Nunez M. Former Facebook workers: We routinely suppressed conservative news. 2016 May 9. In: Gizmodo [Internet]. Available from: https://gizmodo.com/former-facebook-workers-we-routinely-suppressed-conser-1775461006

22. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science, 2019; 366(6464):447–453. https://doi.org/10.1126/science.aax2342 PMID: 31649194

23. Reuters. Amazon scraps secret AI recruiting tool that 'didn't like women'. 2018 Oct 10. In: MailOnline [Internet]. Available from: https://www.dailymail.co.uk/sciencetech/article-6259205/Amazon-scraps-secret-AI-recruiting-tool-showed-bias-against-women.html

24. Sun W, Nasraoui O, Shafto P. Evolution and impact of bias in human and machine learning algorithm interaction. PLoS ONE, 2020; 15(8):e0235502. https://doi.org/10.1371/journal.pone.0235502 PMID: 32790666

25. McKinnon JD, MacMillan D. 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. 2018 Sep 20. Available from: https://www.wsj.com/articles/google-workers-discussed-tweaking-search-function-to-counter-travel-ban-1537488472

26. Meyers PJ. How often does Google update its algorithm? 2019 May 14. In: Moz [Internet]. Available from: https://moz.com/blog/how-often-does-google-update-its-algorithm

27. Epstein R. The new censorship: How did Google become the internet's censor and master manipulator, blocking access to millions of websites? U.S. News & World Report. 2016 Jun 22. Available from: https://www.usnews.com/opinion/articles/2016-06-22/google-is-the-worlds-biggest-censor-and-its-power-must-be-regulated

28. Gallagher R. Google plans to launch censored search engine in China, leaked documents reveal: Search app that will "blacklist sensitive queries" could be launched in six to nine months, according to documents and people familiar with the plans. The Intercept. 2018 Aug 1. Available from: https://theintercept.com/2018/08/01/google-china-search-engine-censorship/

29. Vorhies Z, Heckenlively K. Google leaks: A whistleblower's exposé of Big Tech censorship. Skyhorse Publishing, 2021.

30. Agudo U, Matute H. The influence of algorithms on political and dating decisions. PLoS ONE, 2021; 16(4):e0249454. https://doi.org/10.1371/journal.pone.0249454 PMID: 33882073

31. Ani. Study says people trust computers more than humans. In: Big News Network.com [Internet]. 2021 Apr 15. Available from: https://www.bignewsnetwork.com/news/268700564/study-says-people-trust-computers-more-than-humans

32. Bogert E, Schecter A, Watson RT. Humans rely more on algorithms than social influence as a task becomes more difficult. Sci Rep, 2021; 11. https://doi.org/10.1038/s41598-021-87480-9 PMID: 33850211

33. Edelman. Edelman Trust Barometer 2021. In: Edelman [Internet]. 2021. Available from: https://www.edelman.com/trust/2021-trust-barometer

34. Logg JM, Minson JA, Moore DA. Do people trust algorithms more than companies realize? Harvard Business Review. 2018 Oct 26. Available from: https://hbr.org/2018/10/do-people-trust-algorithms-more-than-companies-realize

35. Howard JJ, Rabbitt LR, Sirotin YB. Human-algorithm teaming in face recognition: How algorithm outcomes cognitively bias human decision-making. PLoS ONE, 2020; 15(8): e0237855. Available from: https://doi.org/10.1371/journal.pone.0237855 PMID: 32822441

36. Curran D. Are you ready? Here is all the data Facebook and Google have on you: The harvesting of our personal details goes far beyond what many of us could imagine. So I braced myself and had a look. The Guardian. 2018 Mar 30. Available from: https://www.theguardian.com/commentisfree/2018/mar/28/all-the-data-facebook-google-has-on-you-privacy

37. Kozyreva A, Lorenz-Spreen P, Hertwig R, Lewandowsky S, Herzog SM. Public attitudes towards algorithmic personalization and use of personal data online: Evidence from Germany, Great Britain, and the United States. Humanit Soc Sci Commun, 2021; 8:1–11. https://doi.org/10.1057/s41599-021-00787-w

38. Fernández-Manzano E, González-Vasco M. Analytic surveillance: Big data business models in the time of privacy awareness. Prof Inform, 2018; 27(2):402–409. https://doi.org/10.3145/epi.2018.mar.19

39. Schneier B. Data and Goliath: The hidden battles to collect your data and control your world. W. W. Norton & Company; 2015.

40. Zuboff S. The Age of surveillance Capitalism: The fight for a human future at the new frontier of power. PublicAffairs; 2019.

41. Bozdag E. Bias in algorithmic filtering and personalization. Ethics Inf Technol, 2013; 15:209–227. https://doi.org/10.1007/s10676-013-9321-6

42. DuckDuckGo. Measuring the "Filter Bubble": How Google is influencing what you click 2018 Dec 4. In: DuckDuckGo [Internet]. Available from: https://spreadprivacy.com/google-filter-bubble-study/

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

44. Epstein R. The new mind control. Aeon. 2016 Feb 18. Available from: https://aeon.co/essays/how-the-internet-flips-elections-and-alters-our-thoughts

45. Feuz M, Fuller M, Stalder F. Personal web searching in the age of semantic capitalism: Diagnosing the mechanisms of personalisation. First Monday, 2011; 16(2). https://doi.org/10.5210/fm.v16i2.3344

46. Pariser E. The filter bubble: How the new personalized web is changing what we read and how we think. Penguin Press, 2011.

47. Lai C, Luczak-Roesch M. You can't see what you can't see: Experimental evidence for how much relevant information may be missed due to Google's web search personalisation. Weber I, Darwish KM, Wagner C, Zagheni E, Nelson L, Samin A, et al. (eds) Social Informatics: 11th International Conference. SocInfo 2019. Lecture Notes in Computer Science, vol 11864. Cham: Springer International Publishing; 2019. Available from: https://link.springer.com/chapter/10.1007/978-3-030-34971-4_17#citeas

48. Newton C. Google's new focus on well-being started five years ago with this presentation. 2018 May 10. In: The Verge [Internet]. Available from: https://www.theverge.com/2018/5/10/17333574/google-android-p-update-tristan-harris-design-ethics

49. Epstein R, Mohr R Jr, Martinez J. The Search Suggestion Effect (SSE): How search suggestions can be used to shift opinions and voting preferences dramatically. Paper presented at: Western Psychological Association; 2018 Apr; Portland, OR. Available from: https://aibrt.org/downloads/EPSTEIN_MOHR_&_MARTINEZ_2018-WPA-The_Search_Suggestion_Effect-SSE-WP-17-03.pdf

50. Al-Abbas LS, Haider AS, Hussein RF. Google autocomplete search algorithms and the Arabs' perspectives on gender: A case study of Google Egypt. GEMA Online Journal of Language Studies, 2020; 20(4):95–112. https://doi.org/10.17576/gema-2020-2004-06

51. Baker P, Potts A. 'Why do white people have thin lips?' Google and the perpetuation of stereotypes via auto-complete search forms. Crit Disc Stud, 2013; 10(2):187–204. https://doi.org/10.1080/17405904.2012.744320

52. Cadwalladr C. Google, democracy and the truth about internet search: Tech-savvy rightwingers have been able to 'game' the algorithms of internet giants and create a new reality where Hitler is a good guy, Jews are evil and. . .Donald Trump becomes president. The Guardian. 2016 Dec 4. Available from: https://www.theguardian.com/technology/2016/dec/04/google-democracy-truth-internet-search-facebook

53. Karapapa S, Borghi M. Search engine liability for autocomplete suggestions: personality, privacy and the power of the algorithm. Int J Law Inf Technol, 2015; 23(3):261–289. https://doi.org/10.1093/ijlit/eav009

54. Keskin B. (2015). What suggestions do Google autocomplete make about children? Black Sea Journal of Public and Social Science, 2015; 7(2):69–78. Available from: https://dergipark.org.tr/en/pub/ksbd/issue/16219/169869

55. Noble SU. Algorithms of oppression: How search engines reinforce racism. New York University Press, 2018.

56. Pradel F. Biased representation of politicians in Google and Wikipedia Search? The joint effect of party identity, gender identity, and elections. Polit Commun, 2020; 38(4):447–478. https://doi.org/10.1080/10584609.2020.1793846

57. Roy S, Ayalon L. Age and gender stereotypes reflected in Google's "autocomplete" function: The portrayal and possible spread of societal stereotypes. The Gerontologist, 2020; 60(6):1020–1028. https://doi.org/10.1093/geront/gnz172 PMID: 31812990

58. Kätsyri J, Kinnunen T, Kusumoto K, Oittinen P, Ravaja N. Negativity bias in media multitasking: The effects of negative social media messages on attention to television news broadcasts. PLoS One, 2016; 11(5):e0153712. https://doi.org/10.1371/journal.pone.0153712 PMID: 27144385

59. Carretié L, Mercado F, Tapia M, Hinojosa JA. Emotion, attention, and the 'negativity bias', studied through event-related potentials. Int J Psychophysiol, 2001; 41(1):75–85. https://doi.org/10.1016/s0167-8760(00)00195-1 PMID: 11239699

60. Arendt F, Fawzi N. Googling for Trump: Investigating online information seeking during the 2016 US presidential election. Inf Commun Soc, 2018; 22(13):1945–1955. https://doi.org/10.1080/1369118X.2018.1473459

61. Solon O, Levin S. How Google's search algorithm spreads false information with a rightwing bias. The Guardian. 2016 Dec 16. Available from: https://www.theguardian.com/technology/2016/dec/16/google-autocomplete-rightwing-bias-algorithm-political-propaganda

62. Gerhart SL. Do web search engines suppress controversy? First Monday, 2004; 9(1–5). https://doi.org/10.5210/fm.v9i1.1111

63. Sullivan D. A reintroduction to Google's featured snippets. 2018 Jan 30. In: Google Blog [Internet]. Available from: https://www.blog.google/products/search/reintroduction-googles-featured-snippets/

64. Singhal A. Introducing the knowledge graph: things, not strings. 2012 May 16. In: Google [Internet]. Available from: https://blog.google/products/search/introducing-knowledge-graph-things-not/

65. Strzelecki A, Rutecka P. Direct answers in Google search results. IEEE Access, 2020; 8:103642–103654. https://doi.org/10.1109/ACCESS.2020.2999160

66. Google Search Help. How Google's featured snippets work. 2022. In: Google Search Help [Internet]. Available from: https://support.google.com/websearch/answer/9351707?hl=en

67. Wright M. Who will be the next president? Google says it's Hillary Clinton. 2015 Jun 25. In: The Next Web [Internet]. Available from: http://thenextweb.com/shareables/2015/06/23/poor-old-jeb

68. Fishkin R. Less than half of Google searches now result in a click. 2019 Aug 13. In: SparkToro [Internet]. Available from: https://sparktoro.com/blog/less-than-half-of-google-searches-now-result-in-a-click/

69. Fishkin R. In 2020, two thirds of Google searches ended without a click. 2021 Mar 22. In: SparkToro [Internet]. Available from: https://sparktoro.com/blog/in-2020-two-thirds-of-google-searches-ended-without-a-click/

70. Wu D, Dong J, Shi L, Liu C, Ding J. Credibility assessment of good abandonment results in mobile search. Inf Process Manag, 2020; 57(6). https://doi.org/10.1016/j.ipm.2020.102350

71. Richter F. Smart speaker adoption continues to rise. 2020 Jan 9. In: Statista [Internet]. Available from: https://www.statista.com/chart/16597/smart-speaker-ownership-in-the-united-states/

72. Liu S. Virtual assistant technology in the U.S.–Statistics & facts. 2021 Aug 18. In: Statista [Internet]. 2021. Available from: https://www.statista.com/topics/7022/virtual-assistants-in-the-us/#topicHeader__wrapper

73. Maras M. From secretly watching your kids to tracking their every move: the terrifying ways smart toys can be used by hackers. Daily Mail Online. 2018 May 10. Available from: https://www.dailymail.co.uk/sciencetech/article-5714587/From-secret-spying-kids-chatting-terrifying-ways-smart-toys-hacked.html

74. Vlahos J. Barbie wants to get to know your child. The New York Times. 2015 Sep 16. Available from: https://www.nytimes.com/2015/09/20/magazine/barbie-wants-to-get-to-know-your-child.html

75. Cox K. These toys don't just listen to your kid; they send what they hear to a defense contractor. 2016 Dec 6. In: The Consumerist [Internet]. Available from: https://consumerist.com/2016/12/06/these-toys-dont-just-listen-to-your-kid-they-send-what-they-hear-to-a-defense-contractor/

76. Walsh M. My Friend Cayla doll banned in Germany over surveillance concerns. ABC News. 2017 Feb 17. Available from: https://www.abc.net.au/news/2017-02-18/my-friend-cayla-doll-banned-germany-over-surveillance-concerns/8282508

77. Munr K. Hacking a talking toy parrot. 2017 Dec 1. In: Pen Test Partners [Internet]. Available from: https://www.pentestpartners.com/security-blog/hacking-a-talking-toy-parrot/

78. Jaswal VK, Croft AC, Setia AR, Cole CA. Young children have a specific, highly robust bias to trust testimony. Psychol Sci, 2010; 21(10):1541–1547. https://doi.org/10.1177/0956797610383438 PMID: 20855905

79. Jaswal VK, Pérez-Edgar K, Kondrad RL, Palmquist CM, Cole CA, Cole CE. Can't stop believing: Inhibitory control and resistance to misleading testimony. Dev Sci, 2014; 17(6):965–976. https://doi.org/10.1111/desc.12187 PMID: 24806881

80. Ma L, Ganea PA. Dealing with conflicting information: Young children's reliance on what they see versus what they are told. Dev Sci, 2010; 13:151–160. https://doi.org/10.1111/j.1467-7687.2009.00878.x PMID: 20121871

81. Children's Online Privacy Protection Act of 1998, 15 U.S.C. §§ 6501–6506 (1998). Available from: https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title15-chapter91&edition=prelim

82. Steeves V. A dialogic analysis of Hello Barbie's conversations with children. Big Data Soc, 2020; 7 (1):1–12. https://doi.org/10.1177/2053951720919151

83. Danovitch JH, Alzahabi R. Children show selective trust in technological informants. J Cogn Dev, 2013; 14:499–513. https://doi.org/10.1080/15248372.2012.689391

84. Murray GW. Who is more trustworthy, Alexa or Mom?: Children's selective trust in a digital age. Technol Mind Behav, 2021; 2(3). https://doi.org/10.1037/tmb0000050

85. Tempesta E. That doesn't sound like wheels on the bus! Parents freak out when Amazon's Alexa misunderstands young son's request for a song and starts rattling off crude pornographic phrases. The Daily Mail Online. 2016 Dec 30. Available from: https://www.dailymail.co.uk/femail/article-4076568/That-doesn-t-sound-like-Wheels-Bus-Parents-freak-Amazon-s-Alexa-misunderstands-young-son-s-request-song-starts-rattling-crude-PORNOGRAPHIC-phrases.html

86. Hern A. CloudPets stuffed toys leak details of half a million users: Company's data compromised, leaking information including email addresses, passwords, and voice recordings. The Guardian. 2017 Feb 28. Available from: https://www.theguardian.com/technology/2017/feb/28/cloudpets-data-breach-leaks-details-of-500000-children-and-adults

87. Jones ML, Meurer K. Can (and should) Hello Barbie keep a secret? IEEE International Symposium on Ethics in Engineering, Science, and Technology (ETHICS), 2016;:1–6. https://doi.org/10.1109/ETHICS.2016.7560047

88. McReynolds E, Hubbard S, Lau T, Saraf A, Cakmak M, Roesner F. Toys that listen: A study of parents, children, and internet-connected toys. Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, 2017;: 5197–5207. https://doi.org/10.1145/3025453.3025735

89. Murnane K. Amazon does the unthinkable and sends Alexa recordings to the wrong person. 2018 Dec 20. In: Forbes [Internet]. Available from: https://www.forbes.com/sites/kevinmurnane/2018/12/20/amazon-does-the-unthinkable-and-sends-alexa-recordings-to-the-wrong-person/?sh=7b4d69bd3ca5

90. Crutzen R, Peters GY, Portugal SD, Fisser EM, Grolleman JJ. An artificially intelligent chat agent that answers adolescents' questions related to sex, drugs, and alcohol: An exploratory study. J Adol Health, 2011; 48(5):514–519. https://doi.org/10.1016/j.jadohealth.2010.09.002 PMID: 21501812

91. Miner AS, Laranjo L, Kocaballi AB. Chatbots in the fight against the COVID-19 pandemic. NPJ Digit Med, 2020; 3(65). https://doi.org/10.1038/s41746-020-0280-0 PMID: 32377576

92. World Health Organization. WHO Health Alert brings COVID-19 facts to billions via WhatsApp. In: World Health Organization [Internet]. Available from: https://www.who.int/news-room/feature-stories/detail/who-health-alert-brings-covid-19-facts-to-billions-via-whatsapp

93. Leggett K. 2018 customer service trends: How operations become faster, cheaper–and yet, more human. Forrester; 2018 Jan 24. p. 4, Fig 2. Available from: https://imistatic.blob.core.windows.net/document-library/boldchat/pdf/en/forrester-2018-customer-service-trends.pdf

94. Newitz A. Ashley Madison code shows more women, and more bots. 2015 Aug 31. In: Gizmodo [Internet]. Available from: https://gizmodo.com/ashley-madison-code-shows-more-women-and-more-bots-1727613924

95. Ishowo-Oloko F, Bonnefon J, Soroye Z, Crandall J, Rahwan I, Rahwan T. Behavioural evidence for a transparency-efficiency tradeoff in human-machine cooperation. Nat Mach Intell, 2019; 1:517–521. https://doi.org/10.1038/s42256-019-0113-5

96. Christian B. The most human human: What artificial intelligence teaches us about being alive. New York: Anchor Books; 2012.

97. Stuart-Ulin CS. Microsoft's politically correct chatbot is even worse than its racist one. 2018 Jul 31. In: Quartz [Internet]. Available from: https://qz.com/1340990/microsofts-politically-correct-chat-bot-is-even-worse-than-its-racist-one/

98. Epstein R. From Russia, with love: How I got fooled (and somewhat humiliated) by a computer. Scientific American Mind, 16–17. 2007 Oct/Nov. Available from: https://www.drrobertepstein.com/downloads/FROM_RUSSIA_WITH_LOVE-Epstein-Sci_Am_Mind-Oct-Nov2007.pdf

99. Brandt M. Smart assistants are getting smarter. 2018 Apr 25. In: Statista [Internet]. Available from: https://www.statista.com/chart/9580/how-smart-are-smart-assistants/

100. Enge E. Rating the smarts of the digital personal assistants in 2019. 2019 Oct 24. In: Perficient [Internet]. Available from: https://www.perficient.com/insights/research-hub/digital-personal-assistants-study

101. Ferrand J, Hockensmith R, Houghton RF, Walsh-Buhi ER. Evaluating smart assistant responses for accuracy and misinformation regarding Human Papillomavirus vaccination: Content analysis study. J Med Internet Res, 2020; 22(8):e19018. https://doi.org/10.2196/19018 PMID: 32744508

102. Lovato SB, Piper AM, Wartella EA. Hey Google, do unicorns exist? Conversational agents as a path to answers to children's questions. Proceedings of the 18th ACM International Conference on Interaction Design & Children. 2019 Jun;: 301–313. https://doi.org/10.1145/3311927.3323150

103. Nobles AL, Leas EC, Caputi TL, Zhu S, Strathdee SA, Ayers JW. Responses to addiction help-seeking from Alexa, Siri, Google Assistant, Cortana, and Bixby intelligent virtual assistants. NPJ Digit Med, 2020; 3:1–3. https://doi.org/10.1038/s41746-019-0211-0 PMID: 31934645

104. Palanica A, Fossat Y. Medication name comprehension of intelligent virtual assistants: A comparison of Amazon Alexa, Google Assistant, and Apple Siri between 2019 and 2021. Front Digit Health, 2021; 3:669971. https://doi.org/10.3389/fdgth.2021.669971 PMID: 34713143

105. Boyd M, Wilson N. Just ask Siri? A pilot study comparing smartphone digital assistants and laptop Google searches for smoking cessation advice. PLoS ONE, 2018; 13(3):e0194811. https://doi.org/10.1371/journal.pone.0194811 PMID: 29590168

106. Toader D, Boca G, Toader R, Măcelaru M, Toader C, Ighian D, et al. The effect of social presence and chatbot errors on trust. Sustainability, 2019; 12(1):256. https://doi.org/10.3390/su12010256

107. Shank DB, Gott A. Exposed by AIs! People personally witness artificial intelligence exposing personal information and exposing people to undesirable content. International Journal of Human-Computer Interaction, 2020; 36(17):1636–1645. https://doi.org/10.1080/10447318.2020.1768674

108. Panzarino M. Apple switches from Bing to Google for Siri web search results on iOS and Spotlight on Mac. 2017 Sep 25. In: TechCrunch [Internet]. Available from: https://techcrunch.com/2017/09/25/apple-switches-from-bing-to-google-for-siri-web-search-results-on-ios-and-spotlight-on-mac/

109. Gavin R. Delivering personalized search experiences in Windows 10 through Cortana. 2016 Apr 28. In: Windows Blog [Internet]. Available from: https://blogs.windows.com/windowsexperience/2016/04/28/delivering-personalized-search-experiences-in-windows-10-through-cortana/

110. Baca MC. Amazon starts crowdsourcing Alexa responses from the public. What could possibly go wrong? Amazon opens its Alexa Answers project to the public and gambles on the goodwill of the Internet. The Washington Post. 2019 Sep 13. Available from: https://www.washingtonpost.com/technology/2019/09/13/amazon-starts-crowdsourcing-alexa-responses-public-what-could-possibly-go-wrong/

111. Wiggers K. Amazon is poorly vetting Alexa's user-submitted answers. 2019 Nov 1. In: VentureBeat [Internet]. Available from: https://venturebeat.com/2019/11/01/amazon-alexa-answers-vetting-user-questions/

112. Newton C. The Verge Tech Survey 2020. The Verge. 2020 Mar 2. Available from: https://www.theverge.com/2020/3/2/21144680/verge-tech-survey-2020-trust-privacy-security-facebook-amazon-google-apple

113. Wang Y, Wu L, Luo L, Zhang Y, Dong G. Short-term internet search using makes people rely on search engines when facing unknown issues. PLoS ONE, 2017; 12(4):e0176325. https://doi.org/10.1371/journal.pone.0176325 PMID: 28441408

114. Koerber B. Amazon Alexa spouted conspiracy theory when asked about chemtrails: 'For a purpose undisclosed to the general public in clandestine programs directed by government officials. 2018 Apr 11. In: Mashable [Internet]. Available from: https://mashable.com/article/amazon-alexa-chemtrails-conspiracy-theory-echo

115. Chung H, Iorga M, Voas J, Lee S. Alexa, can I trust you? Computer, 2017; 50(9):100–104. https://doi.org/10.1109/MC.2017.3571053 PMID: 29213147

116. Sullivan D. Google's "One True Answer" problem–when featured snippets go bad: Obama's planning a coup? Women are evil? Several presidents were in the KKK? Republicans are Nazis? Google can go spectacularly wrong with some of its direct answers. 2017 Mar 5. In: Search Engine Land [Internet]. Available from: https://searchengineland.com/googles-one-true-answer-problem-featured-snippets-270549

117. Thompson A. Google listed "Nazism" as the ideology of the California Republican Party: "Nazism" appears as the first "ideology" of California Republicans in the "knowledge box" of search results. Vice News. 2018 May 31. Available from: https://www.vice.com/en/article/vbq38d/google-is-listing-nazism-as-the-first-ideology-of-the-california-republican-party

118. Soulo T. Ahref's study of 2 million featured snippets: 10 important takeaways. 2017 May 29. In: Ahrefs-blog [Internet]. Available from: https://ahrefs.com/blog/featured-snippets-study/

119. Scull A. Dr Google will see you now: Google health information previews and implications for consumer health. Med Ref Serv Q, 2020; 39(2):165–173. https://doi.org/10.1080/02763869.2020.1726151 PMID: 32329674

120. Hillygus DS, Shields TG. The persuadable voter. Princeton: Princeton University Press; 2009.

121. Liu Y, Ye C, Sun J, Jiang Y, Wang H. (2021). Modeling undecided voters to forecast elections: From bandwagon behavior and the spiral of silence perspective. Int J Forecast, 2021; 37(2):461–483. https://doi.org/10.1016/j.ijforecast.2020.06.011

122. Mayer WG. The swing voter in American politics. Washington, D.C.: Brookings Institutional Press; 2008.

123. Sheehan KB. Crowdsourcing research: Data collection with Amazon's Mechanical Turk. Commun Monogr, 2017; 85(1):140–156. https://doi.org/10.1080/03637751.2017.1342043

124. Loftus EF. Leading questions and the eyewitness report. *Cogn Psychol*, 1975; 7(4):560–572. https://doi.org/10.1016/0010-0285(75)90023-7

125. Ray L. 2020 Google search survey: How much do users trust their search results? 2020 Mar 2. In: Moz [Internet]. Available from: https://moz.com/blog/2020-google-search-survey

126. Robertson A. It's time to stop trusting Google search already. 2017 Nov 10. In: The Verge [Internet]. Available from: https://www.theverge.com/2017/11/10/16633574/stop-trusting-google-search-texas-shooting-twitter-misinformation

127. Epstein R. "The Selfish Ledger" [Transcript]. 2018 May 25. Available from: https://aibrt.org/downloads/GOOGLE-Selfish_Ledger-TRANSCRIPT.pdf

128. Kulwin N. The Internet Apologizes. . . Even those who designed our digital world are aghast at what they created. A breakdown of what went wrong—from the architects who built it. New York Magazine. 2018 Apr 16. Available from: http://nymag.com/intelligencer/2018/04/an-apology-for-the-internet-from-the-people-who-built-it.html

129. Epstein R. Google's snoops: Mining our private data for profit and pleasure. Dissent. 2014 May. Available from: https://www.dissentmagazine.org/online_articles/googles-snoops-mining-our-data-for-profit-and-pleasure

130. Google 'accidentally' snooped on wifi data. Network Security, 2010; 2010(5):2. https://doi.org/10.1016/S1353-4858(10)70052-0

131. Hill K. Wi-spy Google engineer outed as 'hacker' 'god' Marius Milner. 2012 May 1. In: Forbes [Internet]. Available from: https://www.forbes.com/sites/kashmirhill/2012/05/01/wi-spy-google-engineer-outed-as-hacker-god-marius-milner/?sh=63fe94011d74

132. Epstein R, Bock S, Peirson L, Wang H. Large-scale monitoring of Big Tech political manipulations in the 2020 Presidential election and 2021 Senate runoffs, and why monitoring is essential for democracy. Paper presented at: 24th annual meeting of the American Association of Behavioral and Social Sciences. 2021 Jun 14. Available from: https://aibrt.org/downloads/EPSTEIN_et_al_2021-Large-Scale_Monitoring_of_Big_Tech_Political_Manipulations-FINAL_w_AUDIO.mp4

133. Epstein R. Taming Big Tech: The case for monitoring. 2018 May 17. In: Hacker Noon [Internet]. Available from: https://hackernoon.com/taming-big-tech-5fef0df0f00d

134. Why Google poses a serious threat to democracy, and how to end that threat, United States Senate Judiciary Subcommittee on the Constitution. (2019 Jul 16) (testimony of Robert Epstein). Available from: https://www.judiciary.senate.gov/imo/media/doc/Epstein%20Testimony.pdf

135. Caralle K. 90 percent of political donations from Google-related companies go to Democrats: Study. Washington Examiner. 2018 Sep 7. Available from: https://www.washingtonexaminer.com/policy/technology/90-percent-of-political-donations-from-google-related-companies-go-to-democrats-study

136. Dunn J. The tech industry's major players are firmly behind Hillary Clinton. Business Insider. 2016 Nov 1. Available from: https://www.businessinsider.com/tech-company-donations-clinton-vs-trump-chart-2016-11?r=US&IR=T

137. Pearlstein J. Techies donate to Clinton in droves. To Trump? Not so much. Silicon Valley employees are Hillary Clinton's top funders. Wired. 2016 Aug 31. Available from: https://www.wired.com/2016/08/techies-donate-clinton-droves-trump-not-much/

138. American National Election Studies. The ANES guide to public opinion and electoral behavior: Time of presidential election vote decision 1948–2004. ANES. 2021 Aug 16. Available from: https://electionstudies.org/resources/anes-guide/top-tables/?id=109

139. Annenberg Public Policy Center. 2008 National Annenberg Election Survey Telephone and Online Data Sets (account required to access data). 2010 Dec 8. Available from: https://www.annenbergpublicpolicycenter.org/2008-naes-telephone-and-online-data-sets/

140. Epstein R, Zankich VR. The surprising power of a click requirement: How click requirements and warningns affect users' willingness to disclose personal information. PLoS ONE, 2022; 17(2):e0263097. https://doi.org/10.1371/journal.pone.0263097 PMID: 35180222

141. Eisenhower DD. 'Military-industrial complex speech', transcript, *Yale Law School*, 1961 Jan 17. Available from: https://avalon.law.yale.edu/20th_century/eisenhower001.asp

142. Brockmann H, Drews W, Torpey J. A class for itself? On the worldviews of the new tech elite. PLoS ONE, 2021; 16(1):e0244071. https://doi.org/10.1371/journal.pone.0244071 PMID: 33471828

[redacted]

From: Allie Bones
Sent: Wednesday, August 3, 2022 5:22 PM
To: misinformation@cisecurity.org
Subject: Election misinformation

We would like to report this tweet as disinformation:
https://twitter.com/AZGOP/status/1554949821026291712?s=20&t=RgfXSZFsSvDmVm3qW-xfZQ.

Counties are responsible for calculating how many ballots to order. Under Arizona law, counties should review their voter registration statistics by precinct and party (for the Primary) to determine how many ballots to order. See A.R.S. 16-401 & 16-508.

Thank you,

Allie Bones

Allie Bones (She | Her | Hers)
Assistant Secretary of State
Arizona Secretary of State

Email: abones@azsos.gov
Office: 602-542-4919
Cell: 602-540-5348

1700 W. Washington St., 7th Fl. | Phoenix, AZ | 85007

Attachment 5

From:	C.Murphy Hebert
To:	Maria Benson NASS
Subject:	FW: Arizona: Misinformation regarding Sharpies
Date:	Wednesday, November 4, 2020 3:03:00 PM
Attachments:	image001.png

Hi – FyI I just sent a this message and a similar one to Facebook because the "sharpie" posts are out of hand and we can't stamp them out one at a time.

If you can provide any help or need some more background information about the situation here, please let me know.

Thanks!
murphy

From: C.Murphy Hebert
Sent: Wednesday, November 4, 2020 3:02 PM
To: Twitter Government & Politics <gov@twitter.com>
Subject: Arizona: Misinformation regarding Sharpies

We have a huge mis/disinformation issue happening right now with people providing false information about the use of Sharpies on ballots is so widespread that we are not able to flag individually. Is there a way that we can approach this more holistically?

Thanks!

Murphy

C. Murphy Hebert
Director of Communications
Arizona Secretary of State

Email: [redacted]@azsos.gov
Office: 602-542-2228
Cell: 602[redacted]

1700 W. Washington St., 7th Fl. | Phoenix, AZ | 85007