10 Ways Big Tech Can Shift Millions of Votes in the November Elections—Without Anyone Knowing

A researcher describes 10 ways Google, Facebook, other companies could shift millions of votes in the US midterms

BY ROBERT EPSTEIN

September 26, 2018 Updated: September 26, 2018

Authorities in the UK have finally figured out that fake news stories and Russian-placed ads are not the real problem. The UK Parliament is about to impose stiff penalties—not on the people who place the ads or write the stories, but on the Big Tech platforms that determine which ads and stories people actually see.

Parliament's plans will almost surely be energized by the latest leak of damning material from inside Google's fortress of secrecy: The Wall Street Journal recently reported on emails exchanged among Google employees in January 2017 in which they strategized about how to alter Google search results and other "ephemeral experiences" to counter President Donald Trump's newly imposed travel ban. The company claims that none of these plans was ever implemented, but who knows?

While U.S. authorities have merely held hearings, EU authorities have taken dramatic steps in recent years to limit the powers of Big Tech, most recently with a comprehensive law that protects user privacy—theGeneral Data Protection Regulation—and a whopping $5.1 billion fine against Google for monopolistic practices in the mobile device market. Last year, the European Union also levied a $2.7 billion fineagainst Google for filtering and ordering search results in a way that favored their own products and services. That filtering and ordering, it turns out, is of crucial importance.

As years of research I've been conducting on online influence has shown, content per se is not the real threat these days; what really matters is (a) which content is selected for users to see, and (b) the way that content is ordered in search results, search suggestions, newsfeeds, message feeds, comment lists, and so on. That's where the power lies to shift opinions, purchases, and votes, and that power is held by a disturbingly small group of people.

I say "these days" because the explosive growth of a handful of massive platforms on the internet—the largest, by far, being Google and the next largest being Facebook—has changed everything. Millions of people and organizations are constantly trying to get their content in front of our eyes, but for more than 2.5 billion people around the world—soon to be more than 4 billion—the responsibility for what algorithms do should always lie with the people who wrote the algorithms and the companies that deployed them.

In randomized, controlled, peer-reviewed research I've conducted with thousands of people, I've shown repeatedly that when people are undecided, I can shift their opinions on just about any topic just by changing how I filter and order the information I show them. I've also shown that when, in multiple searches, I show people more and more information that favors one candidate, I can shift opinions even farther. Even more disturbing, I can do these things in ways that are completely invisible to people and in ways that don't leave paper trails for authorities to trace.

Worse still, these new forms of influence often rely on ephemeral content—information that is generated on the fly by an algorithm and then disappears forever, which means that it would be difficult, if not impossible, for authorities to reconstruct. If, on Election Day this coming November, Mark Zuckerberg decides to broadcast go-out-and-vote reminders mainly to members of one political party, how would we be able to detect such a manipulation? If we can't detect it, how would we be able to reduce its impact? And how, days or weeks later, would we be able to turn back the clock to see what happened?

Of course, companies like Google and Facebook emphatically reject the idea that their search and newsfeed algorithms are being tweaked in ways that could meddle in elections. Doing so would undermine the public's trust in their companies, spokespeople have said. They insist that their algorithms are complicated, constantly changing, and subject to the "organic" activity of users.

This is, of course, sheer nonsense. Google can adjust its algorithms to favor any candidate it chooses no matter what the activity of users might be, just as easily as I do in my experiments. As legal scholar Frank Pasquale noted in his recent book "The Black Box Society," blaming algorithms just doesn't cut it; the responsibility for what an algorithm does should always lie with the people who wrote the algorithm and the companies that deployed the algorithm. Alan Murray, president of Fortune, recently framed the issue this way: "Rule one in the Age of AI: Humans remain accountable for decisions, even when made by machines."

Given that 95 percent of donations from Silicon Valley generally go to Democrats, it's hard to imagine that the algorithms of companies like Facebook and Google don't favor their favorite candidates. A newly leaked video of a 2016 meeting at Google shows without doubt that high-ranking Google executives share a strong political preference, which could easily be expressed in algorithms. The favoritism might be deliberately programmed or occur simply because of unconscious bias. Either way, votes and opinions shift.

It's also hard to imagine how, in any election in the world, with or without intention on the part of company employees, Google search results would fail to tilt toward one candidate. Google's search algorithm certainly has no equal-time rule built into it; we wouldn't want it to! We want it to tell us what's best, and the algorithm will indeed always favor one dog food over another, one music service over another, and one political candidate over another. When the latter happens … votes and opinions shift.

Here are 10 ways—seven of which I am actively studying and quantifying—that Big Tech companies could use to shift millions of votes this coming November with no one the wiser. Let's hope, of course, that these methods are not being used and will never be used, but let's be realistic too; there's generally no limit to what people will do when money and power are on the line.

1. Search Engine Manipulation Effect (SEME)

Ongoing research I began in January 2013 has shown repeatedly that when one candidate is favored over another in search results, voting preferences among undecided voters shift dramatically—by 20 percent or more overall, and by up to 80 percent in some demographic groups. This is partly because people place inordinate trust in algorithmically generated output, thinking, mistakenly, that algorithms are inherently objective and impartial.

But my research also suggests that we are conditioned to believe in high-ranking search results in much the same way that rats are conditioned to press levers in Skinner boxes. Because most searches are for simple facts ("When was Donald Trump born?"), and because correct answers to simple questions inevitably turn up in the first position, we are taught, day after day, that the higher a search result appears in the list, the more true it must be. When we finally search for information to help us make a tough decision ("Who's better for the economy, Trump or Clinton?"), we tend to believe the information on the web pages to which high-ranking search results link.

As The Washington Post reported last year, in 2016, I led a team that developed a system for monitoring the election-related search results Google, Bing, and Yahoo were showing users in the months leading up to the presidential election, and I found pro-Clinton bias in all 10 search positions on the first page of Google's search results. Google responded, as usual, that it has "never re-ranked search results on any topic (including elections) to manipulate political sentiment"—but I never claimed it did. I found what I found, namely that Google's search results favored Hillary Clinton; "re-ranking"—an obtuse term Google seems to have invented to confuse people—is irrelevant.

Because (a) many elections are very close, (b) 90 percent of online searches in most countries are conducted on just one search engine (Google), and (c) internet penetration is high in most countries these days—higher in many countries than it is in the United States—it is possible that the outcomes ofupwards of 25 percent of the world's national elections are now being determined by Google's search algorithm, even without deliberate manipulation on the part of company employees. Because, as I noted earlier, Google's search algorithm is not constrained by equal-time rules, it almost certainly ends up favoring one candidate over another in most political races, and that shifts opinions and votes.

2. Search Suggestion Effect (SSE)

When Google first introduced autocomplete search suggestions—those short lists you see when you start to type an item into the Google search bar—it was supposedly meant to save you some time. Whatever the original rationale, those suggestions soon turned into a powerful means of manipulation that Google appears to use aggressively.

My recent research suggests that (a) Google starts to manipulate your opinions from the very first character you type, and (b) by fiddling with the suggestions it shows you, Google can turn a 50–50 split among undecided voters into a 90–10 split with no one knowing. I call this manipulation the Search Suggestion Effect (SSE), and it is one of the most powerful behavioral manipulations I have ever seen in my nearly 40 years as a behavioral scientist. How will you know whether Google is messing with your election-related search suggestions in the weeks leading up to the election? You won't.

3. The Targeted Messaging Effect (TME)

If, on Nov. 8, 2016, Mr. Zuckerberg had sent go-out-and-vote reminders just to supporters of Mrs. Clinton, that would likely have given her an additional 450,000 votes. I've extrapolated that number from Facebook's own published data.

Because Zuckerberg was overconfident in 2016, I don't believe he sent those messages, but he is surely not overconfident this time around. In fact, it's possible that, at this very moment, Facebook and other companies are sending out targeted register-to-vote reminders, as well as targeted go-out-and-vote reminders in primary races. Targeted go-out-and-vote reminders might also favor one party on Election Day in November.

My associates and I are building systems to monitor such things, but because no systems are currently in place, there is no sure way to tell whether Twitter, Google, and Facebook (or Facebook's influential offshoot, Instagram) are currently tilting their messaging. No law or regulation specifically forbids the practice, and it would be an easy and economical way to serve company needs. Campaign donations cost money, after all, but tilting your messaging to favor one candidate is free.

4. Opinion Matching Effect (OME)

In March 2016, and continuing for more than seven months until Election Day, Tinder's tens of millions of users could not only swipe to find sex partners, they could also swipe to find out whether they should vote for Trump or Clinton. The website iSideWith.com—founded and run by "two friends" with no obvious qualifications—claims to have helped more than 49 million people match their opinions to the right candidate. Both CNN and USA Today have run similar services, currently inactive.

I am still studying and quantifying this type of, um, helpful service, but so far it looks like (a) opinion matching services tend to attract undecided voters—precisely the kinds of voters who are most vulnerable to manipulation, and (b) they can easily produce opinion shifts of 30 percent or more without people's awareness.

At this writing, iSideWith is already helping people decide who they should vote for in the 2018 New York U.S. Senate race, the 2018 New York gubernatorial race, the 2018 race for New York District 10 of the U.S. House of Representatives, and, believe it or not, the 2020 presidential race. Keep your eyes open for other matching services as they turn up, and ask yourself this: Who wrote those algorithms, and how can we know whether they are biased toward one candidate or party?

5. Answer Bot Effect (ABE)

More and more these days, people don't want lists of thousands of search results, they just want the answer, which is being supplied by personal assistants like Google Home devices, the Google Assistant on Android devices, Amazon's Alexa, Apple's Siri, and Google's featured snippets—those answer boxesat the top of Google search results. I call the opinion shift produced by such mechanisms the Answer Bot Effect (ABE).

My research on Google's answer boxes shows three things so far: First, they reduce the time people spend searching for more information. Second, they reduce the number of times people click on search results. And third, they appear to shift opinions 10 to 30 percent more than search results alone do. I don't yet know exactly how many votes can be shifted by answer bots, but in a national election in the United States, the number might be in the low millions.

6. Shadowbanning

Recently, Trump complained that Twitter was preventing conservatives from reaching many of their followers on that platform through shadowbanning, the practice of quietly hiding a user's posts without the user knowing. The validity of Trump's specific accusation is arguable, but the fact remains that any platform on which people have followers or friends can be rigged in a way to suppress the views and influence of certain individuals without people knowing the suppression is taking place. Unfortunately, without aggressive monitoring systems in place, it's hard to know for sure when or even whether shadowbanning is occurring.

7. Programmed Virality and the Digital Bandwagon Effect

Big Tech companies would like us to believe that virality on platforms like YouTube or Instagram is a profoundly mysterious phenomenon, even while acknowledging that their platforms are populated by tens of millions of fake accounts that might affect virality.

In fact, there is an obvious situation in which virality is not mysterious at all, and that is when the tech companies themselves decide to shift high volumes of traffic in ways that suit their needs. And aren't they always doing this? Because Facebook's algorithms are secret, if an executive decided to bestow instant Instagram stardom on a pro-Elizabeth Warren college student, we would have no way of knowing that this was a deliberate act and no way of countering it.

The same can be said of the virality of YouTube videos and Twitter campaigns; they are inherently competitive—except when company employees or executives decide otherwise. Google has an especially powerful and subtle way of creating instant virality using a technique I've dubbed the Digital Bandwagon Effect. Because the popularity of websites drives them higher in search results, and because high-ranking search results increase the popularity of websites (SEME), Google has the ability to engineer a sudden explosion of interest in a candidate or cause with no one—perhaps even people at the companies themselves—having the slightest idea they've done so. In 2015, I published a mathematical model showing how neatly this can work.

8. The Facebook Effect

Because Facebook's ineptness and dishonesty have squeezed it into a digital doghouse from which it might never emerge, it gets its own precinct on my list.

In 2016, I published an article detailing five ways that Facebook could shift millions of votes without people knowing: biasing its trending box, biasing its center newsfeed, encouraging people to look for election-related material in its search bar (which it did that year!), sending out targeted register-to-vote reminders, and sending out targeted go-out-and-vote reminders.

I wrote that article before the news stories broke about Facebook's improper sharing of user data with multiple researchers and companies, not to mention the stories about how the company permitted fake news stories to proliferate on its platform during the critical days just before the November election—problems the company is now trying hard to mitigate. With the revelations mounting, on July 26, 2018, Facebook suffered the largest one-day drop in stock value of any company in history, and now it's facing a shareholder lawsuit and multiple fines and investigations in both the United States and the EU.

Facebook desperately needs new direction, which is why I recently called for Zuckerberg's resignation. The company, in my view, could benefit from the new perspectives that often come with new leadership.

9. Censorship

I am cheating here by labeling one category "censorship," because censorship—the selective and biased suppression of information—can be perpetrated in so many different ways.

Shadowbanning could be considered a type of censorship, for example, and in 2016, a Facebook whistleblower claimed he had been on a company team that was systematically removing conservative news stories from Facebook's newsfeed. Now, because of Facebook's carelessness with user data, the company is openly taking pride in rapidly shutting down accounts that appear to be Russia-connected—even though company representatives sometimes acknowledge that they "don't have all the facts."

Meanwhile, Zuckerberg has crowed about his magnanimity in preserving the accounts of people who deny the Holocaust, never mentioning the fact that provocative content propels traffic that might make him richer. How would you know whether Facebook was selectively suppressing material that favored one candidate or political party? You wouldn't. (For a detailed look at nine ways Google censors content, see my essay "The New Censorship," published in 2016.)

10. The Digital Customization Effect (DCE)

Any marketer can tell you how important it is to know your customer. Now, think about that simple idea in a world in which Google has likely collected the equivalent of millions of Word pages of information about you. If you randomly display a banner ad on a web page, out of 10,000 people, only five are likely to click on it; that's the CTR—the "clickthrough rate" (0.05 percent). But if you target your ad, displaying it only to people whose interests it matches, you can boost your CTR a hundredfold.

That's why Google, Facebook, and others have become increasingly obsessed with customizing the information they show you: They want you to be happily and mindlessly clicking away on the content they show you.

In the research I conduct, my impact is always larger when I am able to customize information to suit people's backgrounds. Because I know very little about the participants in my experiments, however, I am able to do so in only feeble ways, but the tech giants know everything about you—even things you don't know about yourself. This tells me that the effect sizes I find in my experiments are probably too low. The impact that companies like Google are having on our lives is quite possibly much larger than I think it is. Perhaps that doesn't scare you, but it sure scares me.

The Same Direction

OK, you say, so much for Epstein's list! What about those other shenanigans we've heard about: voter fraud (Trump's explanation for why he lost the popular vote), gerrymandering, rigged voting machines, targeted ads placed by Cambridge Analytica, votes cast over the internet, or, as I mentioned earlier, those millions of bots designed to shift opinions. What about hackers like Andrés Sepúlveda, who spent nearly a decade using computer technology to rig elections in Latin America? What about all the ways new technologies make dirty tricks easier in elections? And what about those darn Russians, anyway?

To all that I say: kid stuff. Dirty tricks have been around since the first election was held millennia ago. But unlike the new manipulative tools controlled by Google and Facebook, the old tricks are competitive—it's your hacker versus my hacker, your bots versus my bots, your fake news stories versus my fake news stories—and sometimes illegal, which is why Sepúlveda's efforts failed many times and why Cambridge Analytica is dust.

"Cyberwar," a new book by political scientist Kathleen Hall Jamieson, reminds us that targeted ads and fake news stories can indeed shift votes, but the numbers are necessarily small. It's hard to overwhelm your competitor when he or she can play the same games you are playing.

Now, take a look at my numbered list. The techniques I've described can shift millions of votes without people's awareness, and because they are controlled by the platforms themselves, they are entirely noncompetitive. If Google or Facebook or Twitter wants to shift votes, there is no way to counteract their manipulations. In fact, at this writing, there is not even a credible way of detecting those manipulations.

And what if the tech giants are all leaning in the same political direction? What if the combined weight of their subtle and untraceable manipulative power favors one political party? If 150 million people vote this November in the United States, with 20 percent still undecided at this writing (that's 30 million people), I estimate that the combined weight of Big Tech manipulations could easily shift upwards of 12 million votes without anyone knowing. That's enough votes to determine the outcomes of hundreds of close local, state, and congressional races throughout the country, which makes the free-and-fair election little more than an illusion.

Full disclosure: I happen to think that the political party currently in favor in Silicon Valley is, by a hair (so to speak), the superior party at the moment. But I also love America and democracy, and I believe that the free-and-fair election is the bedrock of our political system. I don't care how "right" these companies might be; lofty ends do not justify shady means, especially when those means are difficult to see and not well understood by either authorities or the public.

Can new regulations or laws save us from the extraordinary powers of manipulation the Big Tech companies now possess? Maybe, but our leaders seem to be especially regulation-shy these days, and I doubt, in any case, whether laws and regulations will ever be able to keep up with the new kinds of threats that new technologies will almost certainly pose in coming years.

I don't believe we are completely helpless, however. I think that one way to turn Facebook, Google, and the innovative technology companies that will succeed them, into responsible citizens is to set upsophisticated monitoring systems that detect, analyze, and archive what they're showing people—in effect, to fight technology with technology.

As I mentioned earlier, in 2016, I led a team that monitored search results on multiple search engines. That was a start, but we can do much better. These days, I'm working with business associates and academic colleagues on three continents to scale up systems to monitor a wide range of information the Big Tech companies are sharing with their users—even the spoken answers provided by personal assistants. Ultimately, a worldwide ecology of

passive monitoring systems will make these companies accountable to the public, with information bias and online manipulation detectable in real time.
With November drawing near, there is obviously some urgency here. At this writing, it's not clear whether we will be fully operational in time to monitor the midterm elections, but we're determined to be ready for 2020.

Robert Epstein is a senior research psychologist at the American Institute for Behavioral Research and Technology in California. Epstein, who holds a doctorate from Harvard University, is the former editor-in-chief of Psychology Today and has published 15 books and more than 300 articles on internet influence and other topics. He is currently working on a book called "Technoslavery: Invisible Influence in the Internet Age and Beyond." His research is featured in the new documentary "The Creepy Line." You can find him on Twitter @DrREpstein.