

Chapter 9

Subtle New Forms of Internet Influence Are Putting Democracy at Risk Worldwide

Robert Epstein

9.1 Democracy

The free and fair election is fundamental to democracy, but the internet has made possible new, nearly invisible forms of influence that have likely been affecting the outcomes of close elections worldwide for several years now. Left unchecked, these forms of influence will inevitably grow over time, slowly but surely nullifying the democratic process as we know it.

9.2 Digital Gerrymandering

One possible type of influence has been labeled “digital gerrymandering” by Harvard law Professor Zittrain [4]. In conventional gerrymandering, the boundaries of voting districts are altered to favor one political party, virtually guaranteeing that the majority of voters in the newly-drawn districts will vote for that party. Zittrain has pointed out that social media giants such as Facebook could easily accomplish the same sort of manipulation by sending out multiple prompts to “go out and vote” only to people who are known to favor one candidate or party. In close elections, which are quite common, increasing the number of such people by even a small amount could easily flip an election.

Of greater concern, this kind of manipulation could be accomplished—or perhaps *is* being accomplished—without anyone being the wiser. As it is, social media companies already send out customized advertisements to more than a billion

R. Epstein (✉)

American Institute for Behavioral Research and Technology, Vista, California, USA
e-mail: re@aibr.org

people daily based on gender, age, location, purchase histories, and other factors. They could easily send out—or perhaps are *already* sending out—prompts to vote to select groups of people without anyone knowing that these groups are being singled out.

9.3 Facebook Experiment

In a controlled study conducted in 2010, researchers at the University of California San Diego, working with employees at Facebook, demonstrated that repeatedly flashing “VOTE” ads to over 60 million Facebook users on an election day in the US caused 340,000 more people to vote than otherwise would have [1]—an increase of approximately 0.57 %. The sample of people who received the vote prompts was selected at random for the study, but, given Facebook’s massive database of personal information about its users, the company could easily have targeted people with known preferences for certain candidates or parties; hence, the basis for Zittrain’s speculation. Shifting 0.57 % of voters toward one candidate results in a marginal difference of 1.14 %—enough to flip the outcomes of many close elections (that is, elections with win margins under this percentage).

It is notable that the kinds of advertisements that could mobilize voters are ephemeral in nature. They are flashed briefly to users and then disappear. Unlike the main content of a website, which remains somewhat constant and might even be preserved by tools like the Wayback Machine (see <https://archive.org/web/>), advertisements don’t leave a paper trail. This gives companies such as Facebook complete deniability if it was ever accused of interfering in an election.



Here is the kind of prompt Facebook sent its users in the 2010 experiment

9.4 Search Engine Manipulation Effect (SEME)

Research I have been conducting with Ronald Robertson since early 2013 has identified another problematic source of influence over voter preferences which we call SEME (pronounced “seem”), for Search Engine Manipulation Effect. SEME, we have concluded, is almost certainly already influencing close elections around the world, and it is a much larger effect than digital gerrymandering; in fact, it is proving to be one of the largest behavioral effects ever discovered.

Initially, in a series of five randomized, controlled experiments we completed with 4556 participants in two countries, we demonstrated and repeatedly confirmed that when high ranking search results favor one candidate—that is, make him or her look better than his or her opponent—the proportion of undecided voters supporting that candidate can easily be increased by 20 % or more—up to 80 % in some demographic groups [2]. Perhaps of greater concern, very few participants in our experiments showed any awareness that they were viewing biased search results. In other words, SEME is not only a large behavioral effect, it is also almost entirely invisible.

What’s more, search results, like advertisements, are ephemeral. No records are kept of them, which means that they leave no paper trail. Once again, this gives the company displaying such results complete deniability.

High ranking search results alter opinions because most people mistakenly believe that search rankings are determined by an objective, omniscient, and infallible mechanism that is beyond human control. This is confirmed by a variety of research on consumer behavior which shows that people trust and believe higher-ranked search results more than lower-ranked results. Over 50 % of all clicks go to the top two search results, and more than 90 % of users never leave the first page of results. Research shows that this occurs even when high-ranking items are of poor quality; it is not just for convenience sake that people click on high-ranking items but rather because of those deeply-held beliefs regarding their validity. Because of the enormous value that high-ranking items have for purchases, North American companies are now spending more than 20 billion US dollars per year on SEO in an attempt to push their links to higher positions.

When companies, candidates, or political parties compete in an open marketplace to get people’s attention, fairness is maintained—although, of course, the players with more resources have always have the advantage. What happens, however, when the search engine company *itself* has preferences?

This, of course, is the topic of investigations of Google, Inc. by the U.S. Federal Trade Commission, the European Union, and the government of India, all of which have found that Google unfairly favors its own products and services in its search rankings. Given the power of SEME, one must also wonder: what impact might Google have on elections if its search rankings also favored one candidate over another?

9.5 2014 Lok Sabha Election

The fifth experiment reported by Epstein and Robertson [2], conducted in India during the 2014 Lok Sabha election—the largest democratic election in history—is especially relevant to this question. Our previous experiments had focused on elections that were already completed and on candidates who were unknown to the participants in order to minimize any bias the participants might otherwise have brought to the experiments—in other words, to guarantee that they were truly “undecided” voters.




Mr Gandhi, Mr Kejriwal, and Mr Modi, candidates in the 2014 Lok Sabha election in India

In the 2014 experiment, however, we used newspaper ads and online subject pools to recruit 2,096 undecided voters in 27 of India’s 35 states and territories—real voters in the midst of an intense, hotly-contested election campaign. Participants were randomly assigned to groups in which search rankings favored either Mr Modi (the ultimate winner), Mr Kejriwal, or Mr Gandhi. As we found in our previous experiments, exposure to biased search rankings that linked to real web pages (which participants could examine freely) caused voting preferences to shift toward the targeted candidate by 20 % or more. In some demographic groups, such as unemployed males from a certain region of India, the shift was over 70 %. We got this result even though our participants were highly familiar with the candidates.

That we conducted this research in India was especially appropriate given that in March 2014, Google was fined 10 million rupees (USD\$164,000) by the Competition Commission of India for “search engine bias.”

Are there any indications that actual Google search rankings might have favored Mr Modi in the Lok Sabha election? Google's own data—the daily “Google Score” it assigned to the major political candidates based on search volumes—showed that Modi outscored his opponents by at least 25 % for 60 consecutive days prior to the day the polls closed on May 12th.



Google Scores for the major candidates in the 2014 Lok Sabha election in India for the 60 days prior to the close of polls on May 12th. The data were compiled by the author from daily data posted by Google prior to the election. Google has since removed these data from the Internet

Because search volume is one predictor of search ranking, it is reasonable to assume that Mr Modi was also favored in Google's search rankings as more than 430 million votes were cast between April 7 and May 12. If even 10 % of India's voters with internet access were still undecided in the weeks before the polls closed, biased rankings could have driven several million votes toward Mr Modi. What's more, with an increasingly large portion of the world's population having internet access, whatever the impact of biased search rankings is today, it will certainly be much larger in the future. Our own data suggest that more than 85 % of people with internet access are getting at least some of their election information from internet sources—a number that is also likely to increase in coming years.

9.6 2008 and 2012 Presidential Elections

In the United States, Google Scores posted by the company before the 2008 and 2012 presidential elections^{2012 U.S. Presidential Election} showed strong preferences for Barack Obama, the winner of both elections (see <https://plus.google.com/>

photos/104361405143172836769/albums/5795430883215139905/5795430882434931746), and a study published in 2015 on Slate.com confirms that Google search rankings strongly favor Democratic candidates [3].

Couldn't researchers or government agencies simply track and rate search rankings to determine the extent to which they are biased toward one candidate or another? This is not as easy a task as it sounds, and it might even be impossible. Google's revenue model depends on being able to identify users in real time so that it can send them ads targeted to their particular needs; it provides free services so that it can collect relevant information about every user. With the company able to identify individuals and demographic groups with increasing reliability, it is also able to send out customized search rankings to an increasingly large number of users. From a regulatory standpoint, the problem here is that monitors would have no way to look at the customized rankings Google is sending to particular individuals or demographic groups. Rankings that might appear clean on one computer could be highly biased on another.

As I noted earlier, our research also demonstrates that the vast majority of voters are unaware that the search rankings they are viewing are biased toward one candidate; more than 99 % of participants in our India study seemed oblivious. Influence that is invisible to people is the most dangerous kind, because it leaves people falsely believing that they are choosing freely—that they are not being influenced at all.

Models we have developed suggest that opinion shifts of the magnitude we are finding are large enough to flip the outcomes of upwards of 25 % of the world's national elections. As of this writing, we have now replicated SEME nine times with nearly 10,000 participants in multiple countries, and we also been examining ways of suppressing the effect. Among other things, we have discovered that SEME is probably having an enormous impact on a wide range of important decisions people are making every day, not just on voting preferences.

9.7 Power of Google

SEME wouldn't be much of a threat if the online search business were in the hands of a dozen competing companies. Because more than 90 % of search in most European countries and in many other countries around the world is in the hands of a single company, however, no candidate or party has a way of offsetting the influence that Google's search rankings are likely having on elections.

One last point: Although it is reasonable to assume that Google executives are using search rankings to favor candidates that they deem preferable for their business needs (to do otherwise would be imprudent), our data suggest strongly that Google's search rankings are influencing elections even if Google's executives are keeping hands off. This is inevitable because of the fundamental nature of Google's ever-changing search algorithm. So-called "organic" search phenomena will inevitably boost the rankings of some candidates over others; when this happens,

the preferences of undecided voters will shift toward those candidates in a kind of digital bandwagon effect.

Either way, to protect democracy, search rankings related to elections should be strictly regulated.

Robert Epstein is Senior Research Psychologist at the American Institute for Behavioral Research and Technology and recently retired as Professor of Psychology at the University of the South Pacific. The former editor-in-chief of *Psychology Today* magazine and a PhD of Harvard University, Dr Epstein has published 15 books on artificial intelligence and other topics. You can learn more about SEME research at <http://aibr.org> or about Dr Epstein himself at <http://drrobertepstein.com>. You can follow him on Twitter@DrREpstein (<http://twitter.com/DrREpstein>).

References

1. Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415), 295-298.
2. Epstein, R., & Robertson, R. E. (2015, August 4). The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences USA*, 112(33), E4512-E4521. <http://www.pnas.org/content/112/33/E4512.full.pdf?with-ds=yes>
3. Trielli, D., Mussenden, S., & Diakopoulos, N. (2015, December 7) Why Google search results favor Democrats. *Slate*. http://www.slate.com/articles/technology/future_tense/2015/12/why_google_search_results_favor_democrats.single.html
4. Zittrain, J. (2014, June 1). Facebook could decide an election without anyone ever finding out: The scary future of digital gerrymandering—and how to prevent it. *New Republic*, 1. <https://newrepublic.com/article/117878/information-fiduciary-solution-facebook-digital-gerrymandering>