

EMBARGOED until March 14, 2017.

Summary of papers to be presented at the 2nd biennial meeting of the International Convention of Psychological Science, Vienna, Austria, March 24, 2017, and the 97th annual meeting of the Western Psychological Association, Sacramento, CA, April 27, 2017.

A Method for Detecting Bias in Search Rankings, with Evidence of Systematic Bias Related to the 2016 Presidential Election

Robert Epstein (re@aibr.org)

American Institute for Behavioral Research and Technology

Ronald E. Robertson

Northeastern University

Samantha J. Shepherd & Shu Zhang

American Institute for Behavioral Research and Technology

Executive Summary

In a [report](#) published in the *Proceedings of the National Academy of Sciences* summarizing the results of five randomized, controlled experiments with 4,556 participants in two countries, Epstein and Robertson (2015) showed that search results favoring one political candidate can shift the voting preferences of undecided voters substantially—up to 80% in some demographic groups—and that this effect can easily be masked so that few or no people are aware they are viewing biased search results. These and other experiments demonstrating the impact of the Search Engine Manipulation Effect (SEME) beg the question, however, of whether real search rankings are actually biased.

We now describe a method for creating a Nielsen-ratings-type network of confidants whose daily online searches are automatically collected and pooled in a manner that allows for the detection of actual bias in search rankings. The method was implemented in the spring of 2016 in anticipation of the November presidential election. Ninety-five people from 24 U.S. states (mean age 39.9) were recruited, 21 of whom identified themselves as “undecided.” In the months leading up to Election Day, daily election-related searches automatically provided us with the first page of results from a total of 13,207 searches conducted using Google, Bing and Yahoo through the Firefox browser, allowing us to preserve 98,044 election-related web pages. Those pages, in turn, were rated by people recruited from a crowdsourcing website, and the ratings allowed us to compute the average bias per search position, as well as the overall average bias.

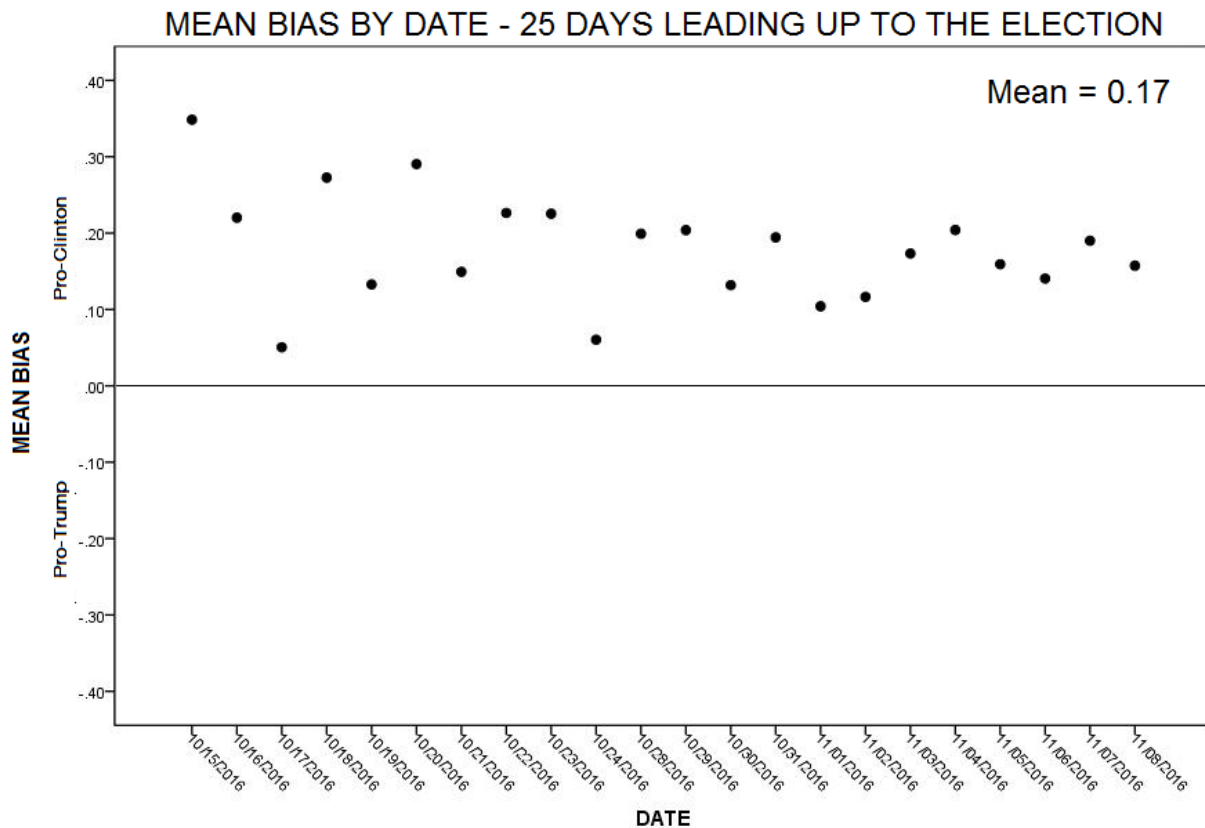
We are still analyzing this wealth of data—all of which will be made public by the Internet Archive (<http://archive.org>) later in 2017—but so far we have found that between May and November 2016, search results displayed in response to a wide range of election-related search terms were, on average, biased in Mrs. Clinton’s favor in all 10 search-result positions. This bias could not be accounted for by the bias in the search terms themselves. We also found different levels of bias in different search engines, as well as evidence of demographically-targeted bias. We don’t know what caused these patterns of bias, but no matter what the cause or causes, given

the power of search rankings to shift votes and opinions without people's awareness (<http://bit.ly/1REqzEY>), they are a matter for concern.

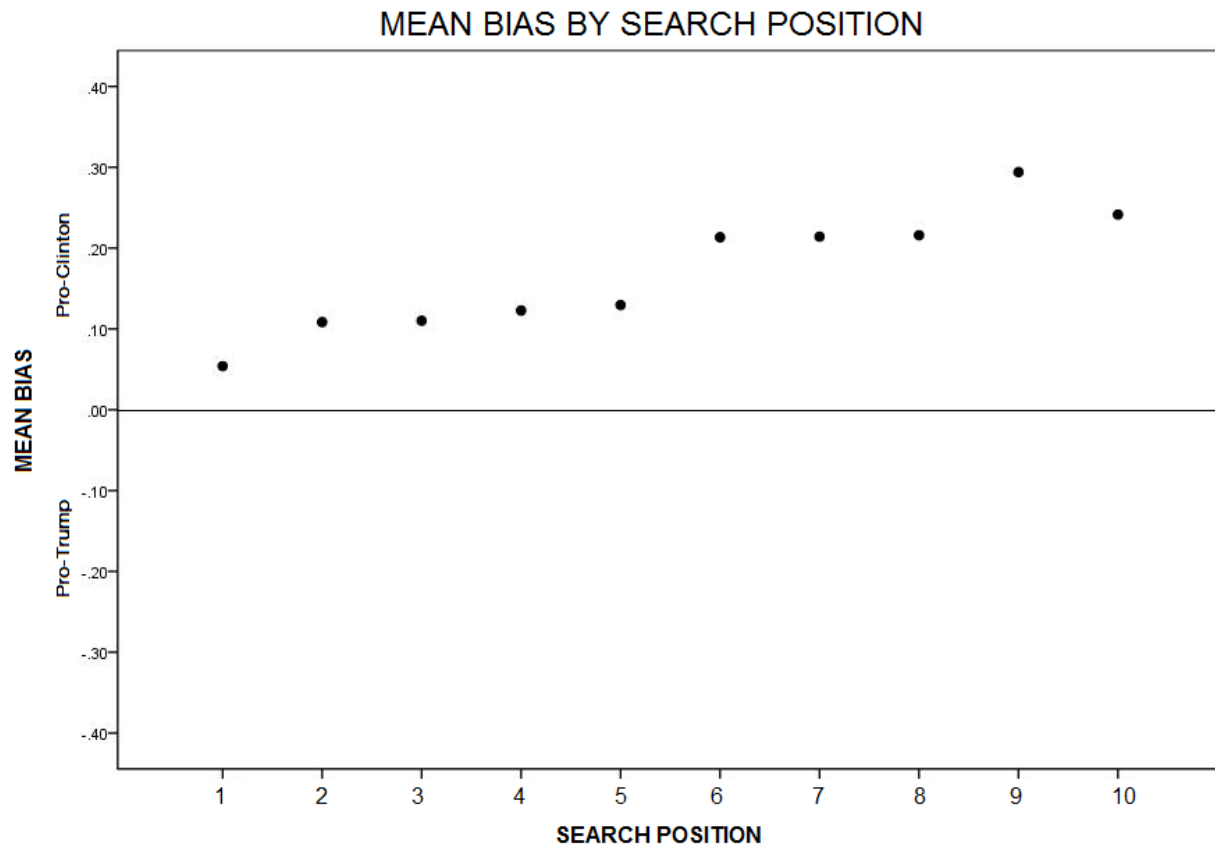
We believe our new tracking system can serve as a prototype for the development of a worldwide ecosystem of software that passively monitors search results, search suggestions, news feeds, advertisements and other normally ephemeral internet content. Such a system might someday play an important role in protecting the integrity of the free-and-fair election. It might also bring, for the first time, some degree of accountability to the Big Tech companies that control ephemeral content.

Selected Supporting Data

1) Issue: *Were search results provided by search engines in the U.S. biased toward one candidate or the other?* Yes. Based on a sample of 4,045 election-related searches conducted during a 25-day period from October 15 to November 8 (Election Day) using the Google and Yahoo search engines through the Firefox browser, we found that search results were, on average, biased to favor Hillary Clinton on all of those days. (Note: In the graph below, values above 0 show a Clinton bias, and values below 0 show a Trump bias.)



2) Issue: *On the first page of search results (10 positions), where was the bias expressed?* We found that the bias was expressed in all 10 search results:



3) Issue: *Was the bias the same for all search engines?* No. The level of pro-Clinton bias we found on Google (0.19) was more than twice as high as the level of pro-Clinton bias we found on Yahoo (0.09). The difference between these values was highly statistically significant ($p < 0.001$).

4) Issue: *Did everyone see the same level of bias?* No. We found evidence of demographic targeting. For example:

a. *Decided vs. undecided*. Search results seen by people who said they had decided how to vote were nearly twice as biased in favor of Clinton (0.21) as search results seen by people who said they were undecided (0.11, $p < 0.001$).

b. *Men vs. women*. Search results seen by men were twice as biased in favor of Clinton (0.24) as search results seen by women (0.12, $p < 0.001$).

c. *Blue vs. red vs. swing states*. People in blue states generally saw the highest level of pro-Clinton bias in search rankings (0.24); people in red states saw the next highest level of pro-Clinton bias (0.12); and people in swing states saw the lowest level of pro-Clinton bias (0.10, $p < 0.001$).

d. *Young vs. old*. Search results seen by people under age 35 were more than twice as biased in favor of Clinton (0.21) as search results seen by people 35 and over (0.10, $p < 0.001$).

Interpretation: Our study doesn't look at this issue directly, but the pattern of demographic differences we found is somewhat consistent with the idea that Big Tech companies show people what they want to see. Note, however, that people in blue (that is, pro-Clinton) states weren't the only ones to see pro-Clinton search results; people in red (that is, pro-Trump) and swing states did too.

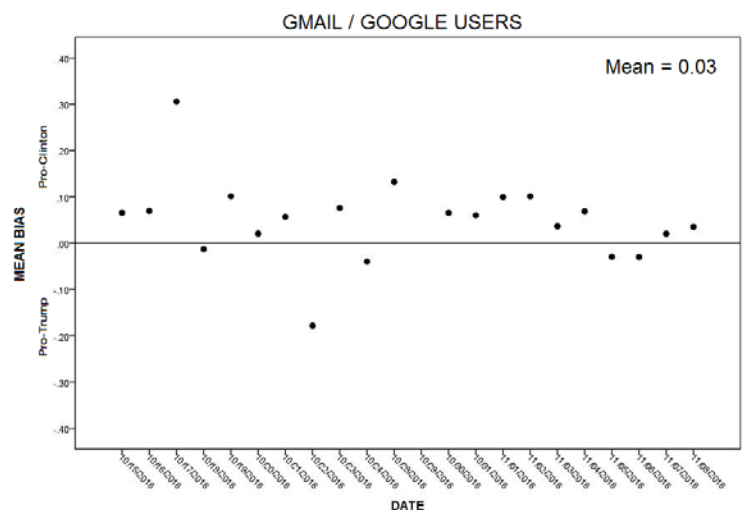
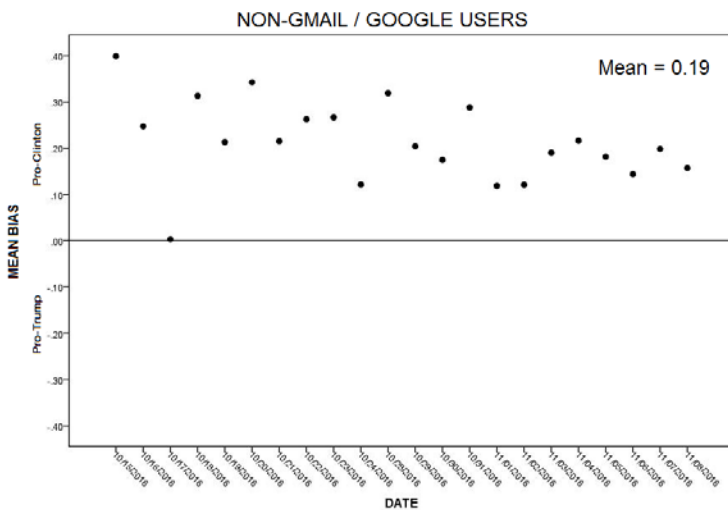
5) Issue: *Could search results have been biased simply because people were selecting biased search terms?* On a scale from -5 (pro-Trump) to +5 (pro-Clinton), the average bias in people's search terms was slightly pro-Trump (-0.08). The search *terms* people used should therefore have yielded a pro-Trump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

6) Issue: *Was there bias before October 15?* Yes, although we typically received data from only a few searches per day before that date, so we are less certain of the numbers. Looking at data from 1,050 searches conducted between May 19 and October 14, 2016, we found, on average, a pro-Clinton bias throughout this period (0.17), as well as a pro-Clinton bias in all 10 search positions on the first page of search results. To put this another way, we found evidence of a pro-Clinton bias in search rankings over a period of nearly six months before the election.

7) Issue: *Did we use all the data we collected?* No, for two reasons:

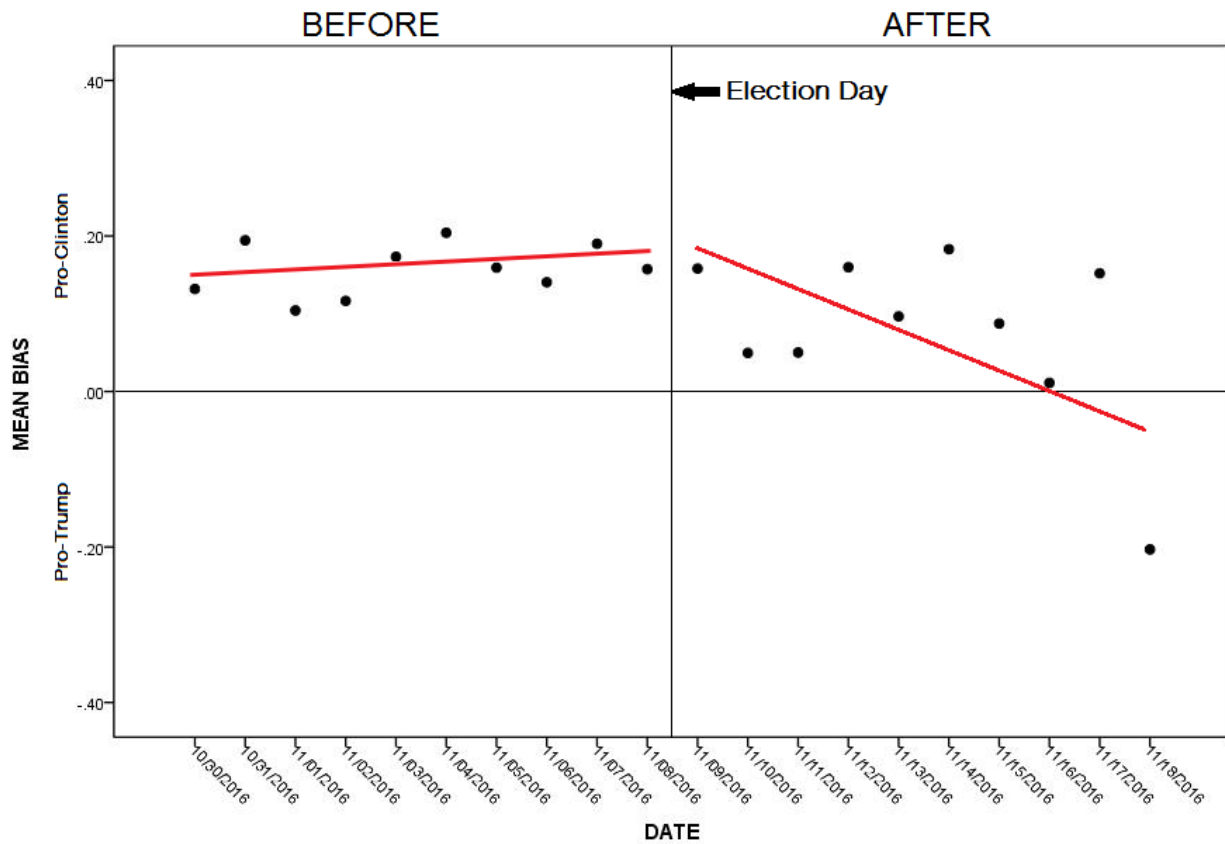
a. First, although we started collecting data on May 19, 2016, it took time for us to recruit confidants for our new network, as well as to refine our new software and procedures. Beginning on October 15, we began reliably to receive more than 50 searches per day, reaching a peak of 483 searches on November 8 (Election Day). We therefore have focused our analysis on this 25-day period.

b. Second, we ultimately had to discard all the data we had received from confidants who communicated with us using gmail, the email service owned by Google. Some of these data were transmitted to us in spurts, with dozens of searches being conducted in a few seconds by the same person; such searches were presumably automated. Some of the statistics these data yielded were also suspect. Compare the pattern of bias values we obtained from *non-gmail* users of the Google search engine (graph on the left) to the pattern of bias we obtained from *gmail* users of the Google search engine (graph on the right):



As you can see, the search results seen by non-gmail users were far more biased (0.19) than the results seen by gmail users (0.03, $p < 0.001$). Perhaps Google identified our confidants through its gmail system and targeted them to receive unbiased results; we have no way to confirm this at present, but it is a plausible explanation for the pattern of results we found. (Note: We are missing data on October 25, 26, and 27 because our server was unable to keep up with the rapidly increasing influx of data. We may be able to recover some of the missing data, but it is unlikely. We made adjustments to our software during this period so our server could accommodate a much faster data flow.)

8) Issue: *What happened to the bias after the election?* It appeared to decrease. The graph below compares the bias observed during the 10 days leading up to the election to the bias observed during the 10 days following the election:



9) Issue: *Could the pro-Clinton bias in search results have shifted votes to Mrs. Clinton?* A comprehensive [study](#) published in 2015 in the *Proceedings of the National Academy of Sciences* found that biased search rankings can easily shift the voting preferences of undecided voters by 20% or more – up to 80% in some demographic groups. Extrapolating from the mathematics introduced in this report, in articles published in [February 2016](#) and thereafter, the lead author of the PNAS study predicted that a pro-Clinton bias in Google’s search results would, over time, shift at least 2.6 million votes to Clinton. She won the popular vote in the November election by [2,864,974 votes](#). Without the pro-Clinton bias in Google’s search results, her win margin in the popular vote might have been negligible.