

RESEARCH ARTICLE

Can biased search results change people's opinions about anything at all? a close replication of the Search Engine Manipulation Effect (SEME)

Robert Epstein ^{*}, Ji Li

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

^{*} re@aibr.org

Abstract

In previous experiments we have conducted on the Search Engine Manipulation Effect (SEME), we have focused on the ability of biased search results to shift voting preferences. In three new experiments with a total of 1,137 US residents (mean age = 33.2), we sought to determine whether biased search rankings could shift people's opinions on topics that do not involve candidates or elections. Each of the new experiments looked at a different topic, and participants were pre-screened to make sure they didn't have strong opinions about these topics. The topics were: Is artificial intelligence useful or dangerous? Is fracking helpful or dangerous? And: Are people born gay or do they choose to be gay? All participants were first asked various demographic questions, then shown brief summaries of the "pro" and "anti" views on each topic, and then asked their opinions about each topic. Next, participants were allowed to conduct an online search using our mock search engine (Kadoodle) lasting up to 15 minutes. In each experiment, one-third of the participants saw biased search results favoring one perspective; one-third saw biased search results favoring the opposing perspective; and one-third (the control group) saw mixed search results. After completing their search, participants were again asked for their opinions about the topic. Our primary dependent variable was Manipulation Power (MP), the percentage increase in the number of participants favoring one viewpoint after having viewed search rankings favoring that viewpoint. The MPs in the three experiments were 25.0%, 30.9%, and 17.8%, respectively. Corresponding shifts were also found for how persuasive participants found each viewpoint to be and for how much they trusted each viewpoint. We conclude that search rankings favoring one viewpoint on a wide range of topics might be able to cause people who have not yet formulated a strong opinion on such topics to adopt the favored perspective. If our findings prove to be robust, we are exposing what might be considered an unforeseen consequence of the creation of search engines, namely that even without human interference, search algorithms will inevitably alter the thinking and behavior of billions of people worldwide on perhaps any topic for which they have not yet formed strong opinions.

OPEN ACCESS

Citation: Epstein R, Li J (2024) Can biased search results change people's opinions about anything at all? a close replication of the Search Engine Manipulation Effect (SEME). *PLoS ONE* 19(3): e0300727. <https://doi.org/10.1371/journal.pone.0300727>

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

Received: October 13, 2023

Accepted: February 28, 2024

Published: March 26, 2024

Copyright: © 2024 Epstein, Li. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Anonymized raw data are accessible at <https://zenodo.org/records/10790867>. The data have been anonymized to comply with requirements of the sponsoring institution's Institutional Review Board (IRB). Data can also be requested from info@aibr.org. The data have been anonymized to comply with requirements of the sponsoring institution's Institutional Review Board (IRB). The IRB granted exempt status to this study under HHS rules because (a) the anonymity of participants was

preserved and (b) the risk to participants was minimal.

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

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

1. Introduction

Controlled experiments conducted in recent years have shown that bias in search engine results can rapidly shift the opinions and voting preferences of undecided voters—by as much as 80% in some demographic groups [1]. Research has also shown that this effect—the search engine manipulation effect (or SEME, pronounced “seem”)—can easily be masked so that no users are aware of the bias they are seeing [1]. SEME has been partially or fully replicated multiple times since it was first published in 2015 [2–16].

Research has also shown that when people do suspect that they are viewing biased search results, that awareness does not necessarily protect them from the impact of the bias. Epstein and Robertson [1] showed in a controlled experiment with more than 2,000 participants from all 50 US states that the few people (8.6%) who could recognize the bias shifted even farther, on average, than people who did not recognize the bias. Why this occurred is not clear, but it could have been because people have inordinate faith in the validity of computer output, at least in part because they have no idea how computers work [17,18]. Research has also shown that vulnerability to SEME and other new forms of online manipulation varies substantially from one demographic group to another [1,6–12].

At least three other features of search engines make them potentially problematic, at least in the eyes of some experts and public policy makers: First, all the content shown to users on search engines is ephemeral. Search suggestions, answer boxes, and search results are all generated on the fly, impact users, and then disappear, and online ephemeral content can impact people's thinking significantly [19,20]. Because they are not stored anywhere, they leave no paper trail for authorities to trace. If biased search results shifted votes in an election—perhaps, in a close election, so many votes that that bias determined the outcome—there would be no way to go back in time to document such an effect [21,22,cf. 23].

Second, over the years, search engines have based their content on increasingly vast amounts of data they have collected about each user; in other words, they now personalize (or “customize”) content to meet the needs of individual users [24]. Many users like this feature of modern search engines, which they consider to be the digital equivalent of personal shoppers [25,26]. On the downside, a long history of research in the marketing and advertising fields has shown that the more one knows about the customer, the easier it is to manipulate him or her [27,28]. This ability applies as much to voters as it does to shoppers [29,30].

Third, because about 92% of search worldwide—everywhere outside of mainland China (the PROC)—is conducted on just one search engine, with no other search engine attracting more than 4% of search [31,32,cf. 33], if the leading search engine chooses to shift votes or opinions in just one direction, there is no way to counteract that very powerful form of influence. If the bias has been masked, there also may be no way to detect it. This is very different from most forms of influence that affect people every day, especially in the days leading up to elections. Most forms of influence—billboards, television and radio commercials, newspapers ads and editorials, online ads and podcasts—are inherently competitive. If you have the resources, you counter your opponent's brutal attack ad with one or more ads that are even more brutal. But if the dominant search engine chooses—either by deliberate acts of its employees or by unconscious or neglectful management of its algorithms [20,22,34]—to support one political candidate or party, there is no way to counteract that influence.

For these reasons, it is important to understand how SEME works, who it affects, and the magnitude of its ability to alter opinions, beliefs, purchases, behavior, and votes. As we have been arguing elsewhere in recent years, it is also important that we develop permanent monitoring systems that can preserve and analyze ephemeral content on a large scale [35–37]. If we don't preserve ephemeral content, we will never know how and to what

extent existing and emerging tech companies are impacting our minds, our children, and our political systems.

Nearly all the research that has been conducted on SEME has focused on only one of these domains—namely, the ability of biased search results to alter the opinions and voting preferences of undecided voters. We are aware of one conference presentation in which SEME was partially replicated in a context involving people's knowledge about health issues [38]. An earlier study found that search results linked to webpages that contained high-quality information about vaccines communicated more knowledge to people than did search results linked to low-quality webpages [39], but that study did not measure opinion shifts.

We are left with a consequential question that we believe has not yet been answered adequately: Can biased search results shift people's opinions not just about political candidates but about a wide range of different topics—perhaps any topic at all? We acknowledge that this question applies mainly, if not exclusively, to people who have not yet made up their mind about that topic: about where to go on vacation, about what kind of car they should buy, about whether gays should be able to marry or adopt children, and so on. How much power do biased search results have, across a wide range of different topics and issues, to shift the opinions and behavior of people who are vulnerable to being influenced?

2. Experiment 1: Can biased search results shift people's views about artificial intelligence (AI)?

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 below.


2.1.2 Participants. 378 participants were recruited online from the Amazon Mechanical Turk (MTurk) subject pool during March 2016. The mean age of our participants was 33.8 ($SD = 11.4$). 56.3% ($n = 213$) of our participants identified themselves as female and 43.7% ($n = 165$) as male. For detailed information about the basic demographic characteristics of our participants in all three experiments, see [S1 Table](#).

91.3% ($n = 345$) of our participants reported using Google as their primary search engine; 4.2% ($n = 16$) reported using Bing, 2.6% ($n = 10$) reported using Yahoo, and 1.9% ($n = 7$) reported using some “other” search engine. Participants reported conducting a wide range of number of searches per week—from 1 per week to over 200 per week ($M = 12.8$, $SD = 17.1$). 44.2% ($n = 167$) of participants identified themselves as politically liberal, 30.2% ($n = 114$) as moderate, 17.7% ($n = 67$) as conservative, 6.6% ($n = 25$) as having no political viewpoint, and 1.3% ($n = 5$) as other.


We asked participants how familiar they were with arguments that favored the use of artificial intelligence (AI) and arguments that were critical of AI on a scale from 1 to 10, where 1 represented “Not familiar at all” and 10 represented “Very familiar.” The mean familiarity level with pro-AI arguments was 3.1 ($SD = 2.2$), and the mean familiarity level with anti-AI arguments was 3.1 ($SD = 2.3$).

2.1.3 Procedure. The experiment was conducted online and employed a pre/post design. Participants were first asked, “Do you have strong beliefs about artificial intelligence?” and only people who clicked “No” were allowed to continue. Then participants were given basic

1 Group Favoring Pro AI:



2 Group Favoring Anti AI:



3 Group Favoring Neither:




Fig 1. Ordering of search results for the three groups. Each small square represents a search result, and each group of six squares represents the search results on one page. Dark pink signifies that a search result links to a web page in which the content has been rated (by independent raters) to be pro-AI. Light pink signifies that search results link to web pages that are less favorable to AI. White signifies that the linked web pages are relatively neutral toward AI. Light blue signifies that search results link to web pages that are somewhat anti-AI. Medium blue signifies that search results link to web pages that are strongly anti-AI. In Group 1 (Pro AI), the search results are in order from pro-AI to anti-AI. In Group 2 (Anti-AI), the search results are in the opposite order. In Group 3 (Control), pro- and anti-AI search results alternate.

<https://doi.org/10.1371/journal.pone.0300727.g001>

instructions and asked for their informed consent (S5 Text). As required by the sponsoring institution's IRB, participants were not asked for identifying information such as name, email address, or telephone number. The participants were then asked a series of demographic questions. They were shown brief (about 100 words) paragraphs about AI. The first paragraph presented a point of view favoring AI, and the second presented a point of view opposing AI (see S2 Text for the full content).

Participants were then asked six opinion questions about AI: two regarding their overall impressions, two regarding how persuasive they found the two viewpoints they had read, and two regarding how much they trusted those two viewpoints; for the full text of the questions, which participants answered on 10-point Likert scales, see S1 Fig. Then participants were asked two choice questions: First, on an 11-point scale from 5 to 0 to 5 (S1 Fig), participants indicated which viewpoint they favored, with "Pro AI" and "Anti AI" appearing at each end of the scale, with the positions counterbalanced. Finally, participants were asked to choose which viewpoint they favored in a forced-choice question (S1 Fig)—again, with the positions of the answers counterbalanced. This page of questions comprised the pre-search test.

At the beginning of the experiment, all participants were randomly assigned to one of three groups: Pro-AI, Anti-AI, or the Control Group, in which people saw alternating pro- and anti-AI arguments. The sequences are shown in Fig 1.

All 30 webpages used in this experiment had previously been rated by five independent reviewers on an 11-point scale from 5 to 0 to 5, where "Pro AI" and "Anti AI" appeared at either end of the scale, and their order was counterbalanced. Based on the mean ratings of the reviewers, the search results were ranked from the most Pro AI (referring to the web page to which the search result linked) to the most Anti AI (again, referring to the web page to which the search result linked), with the relatively neutral search results in the middle (Fig 1, Group 1).

After participants answered those eight questions (six opinion questions and two choice questions) about the pro- and anti-AI points of view, they were then given up to 15 min to use

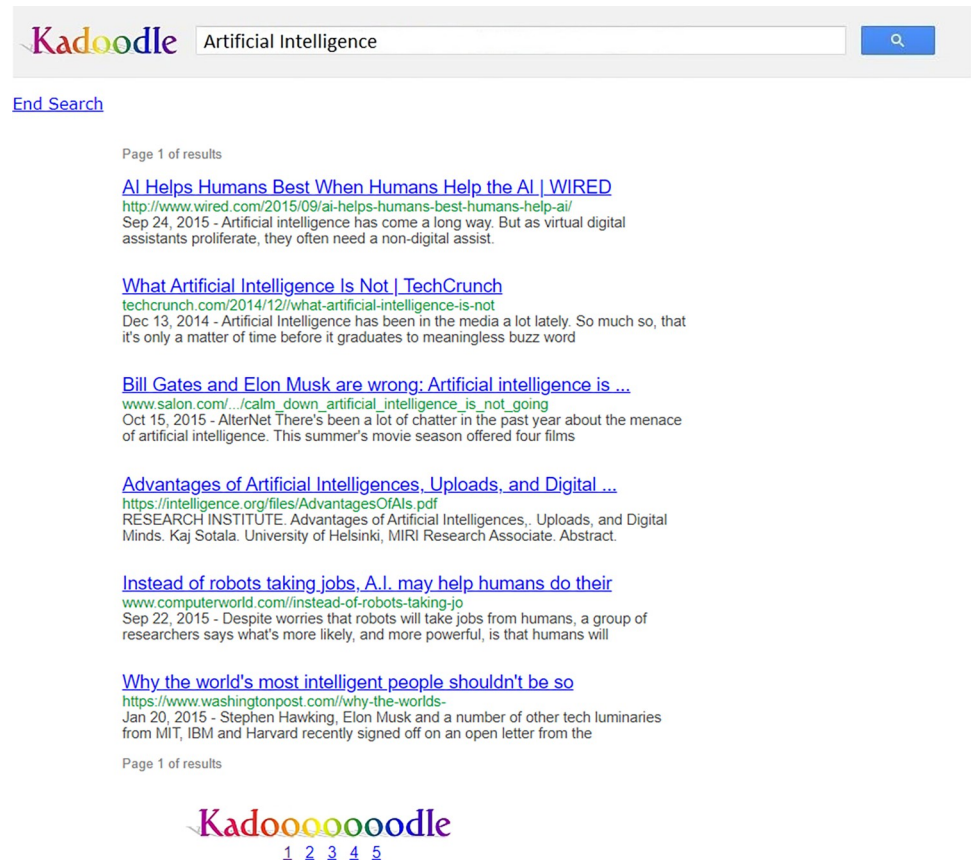


Fig 2. Example of Kadoodle search results page. Each search results webpage contained six different search results. The participant could click on a link to view the corresponding web page, or he or she could click on one of the numbers at the bottom of the page to switch to a different page of search results. The “End Search” shortcut can be seen in the top left corner of the page.

<https://doi.org/10.1371/journal.pone.0300727.g002>

our Kadoodle search engine—a Google simulator—to learn more about AI. Our search engine showed participants five pages of search results, with six search results per page. Participants could click on any of the results and could switch between the pages by clicking on numbers at the bottom of each page (see Fig 2).

All search results were real, scraped from the Google search engine, and all webpages were real, scraped from the internet. The webpages were presented as image files created from the original pages, with no active links.

Participants could end their search by clicking a button on the top left of the screen that read “End Search” (Fig 2). If they failed to click the button, the search session would end when a 15-min time limit was reached.

Now participants were again asked to answer the six opinion questions and two choice questions they had answered prior to search (the post-search test). After participants responded to these questions, they were asked whether anything in the experiment had bothered them. If they answered “yes,” they could then explain what had bothered them in a text box. The purpose of asking participants about what bothered them was to determine whether they detected bias in the search results. We could not ask directly about bias because leading questions of that sort are known to inflate estimates [40].

Finally, participants were thanked for their participation in the experiment and provided with a code number that they could use to be paid by MTurk.

2.2 Results

Most of the experiments we have conducted on online manipulation since we began this type of research in 2013 [1,6–13] have used “vote manipulation power” (VMP) as the most informative metric of change. VMP was defined as the post-manipulation percentage increase in the number of people voting for the candidate favored in the manipulation (see [S1 Text](#) for further details). Because we are now extending our investigation to look at topics that are not election-related, we are introducing a broader variant of VMP, calling it simply “manipulation power” or MP. We define MP as the post-manipulation percentage increase in the number of people choosing the opinion favored in the manipulation (or the belief, candidate, product, perspective, or other categorical content that can be made to look superior to an alternative).

In Experiment 1, the MP was 25.0% (McNemar's Test $X^2 = 22.22, p < 0.001$), which means that in the two bias groups combined, the bias in the search results increased the number of people choosing either a pro-AI or an anti-AI perspective by 25.0%. Specifically, before the search was conducted, the total number of people in the two bias groups who chose the favored perspective was 124. After the search, that number increased by 25.0% to 155.

On the 11-point voting preference scale, pre-manipulation, we found no significant difference between the mean ratings in the three groups ($M_{\text{Pro}} = 1.62, SD = 2.5; M_{\text{Anti}} = 1.78, SD = 2.4; M_{\text{Control}} = 1.31, SD = 2.5$; Kruskal-Wallis $H = 3.13; p = 0.209$ NS). Post manipulation, we found a significant difference between mean ratings in the three groups ($M_{\text{Pro}} = 1.99, SD = 2.7; M_{\text{Anti}} = 0.18, SD = 3.0; M_{\text{Control}} = 1.03, SD = 2.9; H = 24.68; p < 0.001$). Participants in Group 1 shifted 0.37 points toward the favored opinion (Pro AI), and participants in Group 2 condition shifted 1.6 points towards the favored opinion (Anti AI). In addition, the pre-manipulation mean preference for the favored opinion (Groups 1 and 2 combined) was significantly different from the post-manipulation mean preference for the favored opinion (Groups 1 and 2 combined) ($M_{\text{Pre}} = 0.004; SD_{\text{Pre}} = 3.0; M_{\text{Post}} = 0.96; SD_{\text{Post}} = 3.1; M_{\text{Diff}} = 0.956$; Wilcoxon Signed Ranks $z = 6.00; p < 0.001$).

The shift was also indicated by three measures for each of the two opposing opinions: measures of overall impression, persuasiveness, and level of trust ([S1 Fig](#)). Pre to post, the mean favored opinions increased for all three measures, and the non-favored opinions decreased for all three measures. Pre to post, the overall change in opinions was highly significant for all three measures and was in the predicted direction ([Table 1](#)).

In the two bias groups combined (Groups 1 and 2, $n = 246$), the number of people who noticed bias in the search results they saw was 38.2%. This is consistent with the level of bias perception in other SEME experiments when masking has not been employed to disguise the bias [1,4–6].

Table 1. Experiment 1: Pre- and post-search opinion ratings of the favored and non-favored candidate measured on 10-point scales, bias groups only.

	Favored Opinion Mean (SD)			Non-Favored Opinion Mean (SD)			z^\dagger	p
	Pre	Post	Diff	Pre	Post	Diff		
Overall Impression	5.65 (2.6)	6.46 (2.5)	+ 0.81	5.74 (2.6)	4.98 (2.4)	- 0.76	-5.20	< 0.001
Persuasiveness	5.63 (2.4)	6.88 (2.5)	+ 1.25	5.76 (2.3)	5.09 (2.6)	- 0.67	-6.46	< 0.001
Level of Trust	5.61 (2.2)	6.63 (2.3)	+ 1.02	5.85 (2.1)	5.36 (2.4)	- 0.49	-6.15	< 0.001

[†]The z values come from Wilcoxon Signed Ranks Test between the post-search minus pre-search ratings for the favored candidate and the post-search minus pre-search ratings for the non-favored candidate.

<https://doi.org/10.1371/journal.pone.0300727.t001>

3. Experiment 2. Can biased search results shift people's views about fracking?

3.1 Methods

3.1.1 Participants. 394 participants were recruited online from MTurk during March 2016. The mean age of our participants was 32.9 ($SD = 10.2$). 52.0% ($n = 205$) of our participants identified themselves as female and 48.0% ($n = 189$) as male. For detailed information about basic demographic characteristics, see [S1 Table](#).

94.2% ($n = 371$) of the participants selected Google as their primary search engine, 4.1% ($n = 16$) as Bing, 1.5% ($n = 6$) as Yahoo, and 0.3% ($n = 1$) as other. Participants reported the number of searches they conducted per week ranging from 0 to 150 ($M = 14.4$, $SD = 18.1$). 42.4% ($n = 167$) of participants reported being liberal, 33.8% ($n = 133$) as moderate, and 15.5% ($n = 61$) as conservative; 6.1% ($n = 24$) reported having no political views, and 2.3% ($n = 9$) reported their political viewpoint as other. The mean familiarity level of participants with pro-fracking arguments was 3.6 ($SD = 2.3$); for anti-fracking arguments, it was 4.2 ($SD = 2.5$).

3.1.2 Procedure. The procedure for Experiment 2 was the same as for Experiment 1, except that the topic was fracking.

3.2 Results

In Experiment 2, the MP was 30.9% (McNemar's Test $X^2 = 25.14$, $p < 0.001$), which means that in the two bias groups combined, the bias in the search results increased the number of people choosing either a pro-fracking or an anti-fracking perspective by 30.9%. Specifically, before the search was conducted, the total number of people in the two bias groups who chose the favored perspective was 136. After the search, that number increased by 30.9% to 178.

On the 11-point voting preference scale, pre-manipulation, we found no significant difference between the mean ratings in the three groups ($M_{Pro} = -0.79$, $SD = 2.7$; $M_{Anti} = -0.68$, $SD = 2.7$; $M_{Control} = -0.24$, $SD = 2.7$; $H = 2.48$; $p = 0.289$ NS). Post manipulation, we found a significant difference between mean ratings in the three groups ($M_{Pro} = -0.09$, $SD = 3.2$; $M_{Anti} = -2.44$, $SD = 2.8$; $M_{Control} = -0.97$, $SD = 3.7$; $H = 40.35$; $p < 0.001$). Participants in Group 1 shifted 0.7 points toward the favored opinion (Pro Fracking), and participants in Group 2 condition shifted 1.76 points towards the favored opinion (Anti Fracking). In addition, the pre-manipulation mean preference for the favored opinion (Groups 1 and 2 combined) was significantly different from the post-manipulation mean preference for the favored opinion (Groups 1 and 2 combined) ($M_{Pre} = -0.09$; $SD_{Pre} = 2.8$; $M_{Post} = 1.12$; $SD_{Post} = 3.3$; $M_{Diff} = 1.21$; $z = 8.06$; $p < 0.001$).

The shift was also indicated by three measures for each of the two opposing opinions: measures of overall impression, persuasiveness, and level of trust ([S1 Fig](#)). Pre to post, the mean favored opinions increased for all three measures, and the non-favored opinions decreased for all three measures. Pre to post, the overall change in opinions was highly significant for all three measures and was in the predicted direction ([Table 2](#)).

Table 2. Experiment 2: Pre- and post-search opinion ratings of the favored and non-favored candidate measured on an 11-point scale, bias groups only.

	Favored Opinion Mean (SD)			Non-Favored Opinion Mean (SD)			z^\dagger	p
	Pre	Post	Diff	Pre	Post	Diff		
Overall Impression	5.73 (2.3)	6.68 (2.6)	+ 0.95	5.78 (2.3)	4.80 (2.6)	- 0.98	-7.66	< 0.001
Persuasiveness	5.77 (2.2)	6.95 (2.5)	+ 1.18	5.66 (2.3)	4.75 (2.7)	- 0.91	-7.47	< 0.001
Level of Trust	5.44 (2.0)	6.26 (2.6)	+ 0.82	5.49 (2.1)	4.77 (2.6)	- 0.72	-6.56	< 0.001

[†] The z values come from Wilcoxon signed ranks test between post-search minus pre-search ratings for the favored candidate and the post-search minus pre-search ratings for the non-favored opinion.

In the two bias groups combined (Groups 1 and 2, $n = 286$), the number of people who noticed bias in the search results they saw was 50.3%. This is higher than the typical level of bias perception we have found in other SEME experiments when masking has not been employed to disguise the bias [1,4–6].

4. Experiment 3. Can biased search results shift people's views about sexual orientation?

4.1 Methods

4.1.1 Participants. 365 participants were recruited online from MTurk during March 2016. The mean age of our participants was 32.8 ($SD = 10.6$). 55.9% ($n = 204$) of our participants identified themselves as female and 44.1% ($n = 161$) as male. For detailed information about basic demographic characteristics, see [S1 Table](#).

93.7% ($n = 342$) of our participants reported Google as their primary search engine, 2.5% ($n = 9$) as Bing, 2.2% ($n = 8$) as Yahoo, and 1.6% ($n = 6$) as other. Participants reported the number of searches they conducted per week ranging from 0 to 250 ($M = 14.1$, $SD = 24.6$). 42.7% ($n = 156$) of participants reported being liberal, 36.7% ($n = 134$) as moderate, and 12.1% ($n = 44$) as conservative; 5.8% ($n = 21$) reported having no political views, and 2.7% ($n = 10$) reported their political viewpoint as other. The mean familiarity level of participants with born-gay arguments was 7.3 ($SD = 2.6$); for choose-to-be-gay arguments, it was 7.3 ($SD = 2.5$).

4.1.2 Procedure. The procedure for Experiment 3 was the same as for Experiment 1, except that the topic was sexual orientation—specifically, whether people are born gay or whether they choose to be gay.

4.2 Results

In Experiment 3, the MP was 17.8% (McNemar's Test $X^2 = 11.81$, $p < 0.001$), which means that in the two bias groups combined, the bias in the search results increased the number of people choosing either a pro-AI or an anti-AI perspective by 17.8%. Specifically, before the search was conducted, the total number of people in the two bias groups who chose the favored perspective was 135. After the search, that number increased by 17.8% to 159.

On the 11-point voting preference scale, pre-manipulation, we found no significant difference between the mean ratings in the three groups ($M_{\text{Born}} = 1.54$, $SD = 2.9$; $M_{\text{Choose}} = 1.27$, $SD = 3.1$; $M_{\text{Control}} = 1.86$, $SD = 2.9$; $H = 2.25$; $p = 0.325$ NS). Post manipulation, we found a significant difference between mean ratings in the three groups ($M_{\text{Born}} = 2.56$, $SD = 3.0$; $M_{\text{Choose}} = 0.41$, $SD = 3.6$; $M_{\text{Control}} = 2.01$, $SD = 3.2$; $H = 26.87$; $p < 0.001$). Participants in Group 1 shifted 1.02 points toward the favored opinion (Born Gay), and participants in Group 2 condition shifted 0.86 points towards the favored opinion (Choose to be Gay). In addition, the pre-manipulation mean preference for the favored opinion (Groups 1 and 2 combined) was significantly different from the post-manipulation mean preference for the favored opinion (Groups 1 and 2 combined) ($M_{\text{Pre}} = 0.12$; $SD_{\text{Pre}} = 3.3$; $M_{\text{Post}} = 1.05$; $SD_{\text{Post}} = 3.6$; $M_{\text{Diff}} = 0.93$; $z = 7.13$; $p < 0.001$).

The shift was also indicated by three measures for each of the two opposing opinions: measures of overall impression, persuasiveness, and level of trust ([S1 Fig](#)). Pre to post, the mean favored opinions increased for all three measures, and the non-favored opinions decreased for all three measures. Pre to post, the overall change in opinions was highly significant for all three measures and was in the predicted direction ([Table 3](#)).

In the two bias groups combined (Groups 1 and 2, $n = 252$), the percentage of people who noticed bias in the search results they saw was 28.6%. This is similar to the typical level of bias

Table 3. Experiment 3: Pre- and post-search opinion ratings of the favored and non-favored candidate measured on an 11-point scale, bias groups only.

	Favored Opinion Mean (SD)			Non-Favored Opinion Mean (SD)			z^\dagger	p
	Pre	Post	Diff	Pre	Post	Diff		
Overall Impression	6.00 (2.6)	6.65 (2.9)	+ 0.65	5.93 (2.7)	5.24 (3.0)	- 0.69	-5.97	< 0.001
Persuasiveness	5.70 (2.8)	6.63 (3.1)	+ 0.93	5.69 (2.8)	4.87 (3.1)	- 0.82	-7.91	< 0.001
Level of Trust	5.77 (2.7)	6.52 (3.0)	+ 0.75	5.82 (2.8)	5.08 (3.1)	- 0.74	-6.90	< 0.001

[†]The z values come from Wilcoxon Signed Ranks Test between the post-search minus pre-search ratings for the favored candidate and the post-search minus pre-search ratings for the non-favored candidate.

<https://doi.org/10.1371/journal.pone.0300727.t003>

perception we have found in other SEME experiments when masking has not been employed to disguise the bias [1,4–6].

5. Discussion

In the three experiments we have described above, we produced shifts in preferences of 25.0%, 30.9%, and 17.8%, respectively, after our participants conducted just one search on our Kadoo search engine. These numbers are based on shifts in the two bias groups combined. The fact that participants were randomly assigned to one or the other of those two groups means we were able to shift people's thinking for or against a particular perspective arbitrarily. Mean voting preferences on an 11-point scale and mean opinion ratings also shifted predictably in the direction of the bias. These results support our conjecture that bias in the online search results displayed to users by search engine companies have the potential to change people's thinking about—well, perhaps anything at all.

Key to this finding is the fact that we deliberately worked with people who did not, to begin with, already have strong opinions about the three topics we explored. Presumably, people with strong opinions about such matters would be difficult to influence with biased search results [e.g., 41]. That is a matter we are continuing to investigate in ongoing research. If indeed people who are undecided on some issue are the most vulnerable to such a manipulation, it is notable that search engine companies are not only in a unique position to employ bias in search results to impact people's thinking, they are also in a unique position to identify people who are most vulnerable to this type of manipulation—that is, people who have not yet made up their minds. A company such as Google, which openly tracks people through their emails [42] (using Gmail, the most widely used email system in the world [43]), online searches (using the Google search engine, which handles 92% of online search in most countries [31,32]), Chrome (the most widely used browser in the world [44]), Android (the most widely used mobile operating system in the world [45]), and many other platforms and applications [46,47], can easily identify people who are undecided or uncommitted on some issue.

5.1 Conclusions

Even without intent by employees or executives at Google and other companies that operate search engines, the power that biased search results appear to have to shift opinions about a wide range of topics should be a matter of great concern to legislators, regulators, and public policy makers. We make this strong assertion because, by definition, search algorithms always do three things: they *filter* content (by selecting a small amount of content to display while setting aside a vast amount of other content), they *order* content (by ranking the content they will display), and they *customize* content (by adjusting both the filtering and ordering to best match the interests and needs of the user). In other words, in some sense *all* search results are

biased, and therein lies their value. For a given user, search results will always favor one dog food over another, and we wouldn't want it any other way. The problem is that for people sitting on a fence, that customized, filtered, and ranked content the search engine shows them appears to be effective as a tool for pushing people to tumble off one side of that fence.

To put this another way, the search engine is not only the most powerful tool ever invented for providing factual answers to simple questions, it might also be the most powerful tool ever invented for influencing people's opinions, even if influence was never the intent. In a separate study [48] we examined this issue from the perspective of operant conditioning. About 86% of the searches people conduct on major search engines are for simple facts [49], and those facts almost invariably turn up in the top position of search results. Like rats in a Skinner box, we thus learn, over and over again, that the most valuable and accurate search results are the ones at the top of the list. When that day comes when we pose an open-ended query—"best restaurant in Atlanta," "is fracking safe?," or "how to solve the immigration problem"—we again tend to attend to and trust those high-ranking links, which will bring us to web pages that likely favor one perspective. That should surprise no one; there are no equal-time rules in search algorithms, after all. They are designed to find the "best" results, not to show a series of pro- and anti- results in alternating order (like the results we showed in our Control Group). Google does so by examining link patterns [50], but no matter what technique is used, a search algorithm will always, or nearly always, tend to favor one perspective over another. That favoritism might occur because one perspective is dominant on the internet, because of the conscious or unconscious biases of the programmers who created and maintain the search algorithm [51–53], or because of company policies that elevate or suppress content deliberately through white listing or black listing [34,54]. The present study extends previous research only in helping to shed light on one issue: Could search engine bias shift people's views about a wide range of different topics? The answer appears to be yes.

To put this issue yet another way: SEME is a list effect with a difference. Unlike other list effects researchers have studied over the past century, beginning with the serial position effect [55–57], SEME is supported by a daily regimen of operant conditioning that will never stop. Simple factual searches will continue to teach people *ad nauseum* that high-ranking search results are truer and more valid than lower ranking search results. Presumably this is why companies worldwide spend vast sums each year trying to push their products a notch or two higher in Google search results; a single increment can increase clicks by 32.3% [58].

In research we are currently conducting on what we call the "digital personalization effect" (DPE), we are learning that personalization—for example, showing people content from sources we know they trust—can dramatically increase the impact of SEME and other new forms of influence the internet has made possible [59]. When you combine three causal factors—(1) bias in search results, which is an essential and important feature of good search results, (2) customization in search results, which Google in particular has long taken pride in providing [24,60], and (3) a company's ability to identify just those users who are especially vulnerable to influence—a troubling picture emerges. The picture becomes even more alarming when one recognizes that both search suggestions [11] and answer boxes [10]—both of which are commonly shown by Google search—also have the power to shift opinions. What if all of these factors align to push opinions in the same direction? And what if these types of influence are similarly biased in online experiences people are having day after day on multiple platforms? We are currently exploring these questions in experiments on what we call the "multiple exposure effect" (MEE) and the "multiple platforms effect" (MPE).

5.2 Limitations and future research

Our conclusions are subject to a number of constraints, two of which we believe are obvious and nontrivial. First, our subjects were drawn from the MTurk subject pool. In recent years, that subject pool has been tainted by bots [61,62], and concerns have been raised about just how representative the US portion of that subject pool is of the general population [63,64]. Fortunately, we conducted the present experiments in early 2016, well before most of the substantive concerns about MTurk were expressed [61,62]. Nevertheless, we acknowledge that the subjects in our experiments are not necessarily representative of the general population, a matter that can only be explored with replications using other sampling methods. On the bright side, our participants were demographically diverse (S1 Table)—far more so than the small group of sophomores at a single college or university who have so often been utilized in social science studies [65–67].

Second, we made no attempt to measure the long-term effects of the opinion shifts we measured. The impact we had on participants in our six bias groups might have been as ephemeral as search results typically are. Although our procedures can shed no light on this issue, we would be remiss in not pointing out that a search engine company such as Google could easily expose users to similarly biased content dozens or even hundreds of times over a period of a few months. If such exposures are additive in their impact, as our most recent research suggests [68], it is not unreasonable to believe that our experiments might be underestimating the power that biased search results might have on people's thinking about virtually any topic (as long as the users have not already formed strong opinions).

Our Kadoodle simulator also differed from Google's search engine in some respects. Google typically shows many pages of search results with about 10 results per page (on desktop and laptop computers). We showed only five pages of search results with only six results per page. We also did not show people search suggestions or answer boxes, which have become common on Google search pages in recent years. When answer boxes are added to search results, people spend less time examining search results and click on fewer search results [10]; if the answer boxes share the bias of the search results, however, opinions and votes shift even farther in the direction of the bias than they would have had only search results been shown [10]. Again, if search suggestions share the same bias as the search results, they too will increase the impact of those results [11]. So although our simulator differs from Google's home page, it does so mainly in ways that make it less powerful as a source of influence.

Our "pro" and "anti" design could also be viewed as simplistic, although our 11-point scale—the first of our two choice questions—does allow participants to indicate intermediate degrees of support. We also acknowledge that we only evaluated three topics; this necessarily limits the generalizability of our results. The fact that we found different levels of opinion shifts for each of these three topics, ranging from 17.8% to 30.9%, suggests that bias in search results might, in general, have more power to influence people's views about some topics than other topics. There are also undoubtedly individual differences in people's vulnerability to such influence. Finally, we note that our data were collected in 2016. Given that both the form and function of search engines has changed little since then, we have no reason to believe that our findings would be difficult to replicate today. That said, it is important that our findings be replicated with contemporary groups.

Regarding future research, we have already mentioned three projects we have in progress that will shed more light on new forms of manipulation that the internet has made possible: MEE, MPE, and DPE. Regarding the range of opinions that might be influenced by biased (or by biased and personalized) search results, determining and understanding that range can be accomplished by varying topics in systematic ways. We have already learned that different

demographic groups vary in how vulnerable they are to the manipulation we employed in the present experiments (S2 to S5 Tables), and we also have found demographic effects in other studies of online influence [1,9–11,13,59]. Further research might show predictable patterns in how vulnerable different demographic groups are (and, for that matter, in how vulnerable different individuals are) to having their opinions altered on different topics by biased search results. An extensive literature on influence and decision-making has already shown how demographic characteristics interact both with types of influence and the topics being considered [1,11,69–71].

Future research should also explore an odd feature of the search engine—one that we alluded to earlier and that might be considered a disturbing unforeseen consequence [72] of the creation of search engines—a consequence that is unavoidable, given the way search engines must work. Search results are useful precisely because they order information from best to worst; an equal-time rule would make them worthless, although perhaps—as a way of protecting the free-and-fair election from undue influence—an exception could be made someday for links to information about political candidates. Generally speaking, however, search results will *always* train people to value high-ranking results over lower ones, which means perforce that search results shown in response to open-ended queries will *always* shift people's thinking and behavior, sometimes in trivial ways and sometimes in profound ones. This will almost always occur, moreover, without people's awareness [1]. If shifts of this sort are deliberately programmed by software engineers, humanity might always be unduly influenced by such people. Our guess, though, is that only an infinitesimally small portion of open-ended queries are of interest to programmers or executives at tech companies. That means that the vast majority of shifts in opinions and behavior being produced by search engines 24 hours a day in people around the world are currently being determined by algorithms without human guidance.

Where algorithms are being left to their own devices (so to speak) by their human creators, they are currently determining what content goes viral or gets suppressed, what many people buy, what many people believe, and whom many people vote for. As generative AI systems are increasingly incorporated into the algorithms that currently dominate our lives, will the growing power of these systems be used in humanity's interest? Will we even understand what is happening to us?

Supporting information

S1 Fig. Pre-search impression questions.

(DOCX)

S1 Table. Experiment 1, 2, & 3 demographics.

(DOCX)

S2 Table. Demographics analysis by gender.

(DOCX)

S3 Table. Demographics analysis by age.

(DOCX)

S4 Table. Demographics analysis by education level.

(DOCX)

S5 Table. Demographics analysis by ethnicity.

(DOCX)

S1 Text. Manipulation Power (MP) calculation.

(DOCX)

S2 Text. Artificial intelligence summary.
(DOCX)

S3 Text. Fracking summary.
(DOCX)

S4 Text. Sexual orientation summary.
(DOCX)

S5 Text. Request for informed consent.
(DOCX)

Acknowledgments

We thank R. Robertson for assistance in conducting the experiments and C. Mourani, E. Olson, and V. R. Zankich for assistance in data analysis and manuscript preparation.

Author Contributions

Conceptualization: Robert Epstein.

Formal analysis: Robert Epstein, Ji Li.

Methodology: Robert Epstein.

Project administration: Robert Epstein.

Supervision: Robert Epstein.

Writing – original draft: Robert Epstein.

Writing – review & editing: Ji Li.

References

1. 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):E4512–21. Available from: <https://doi.org/10.1073/pnas.1419828112> PMID: 26243876
2. Ludolph R, Allam A, Schulz PJ. Manipulating Google's knowledge graph 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. Available from: <https://doi.org/10.2196/jmir.5430> PMID: 27255736
3. Haas A, Unkel J. Ranking versus reputation: perception and effects of search result credibility. *Behav Inf Technol*. 2017; 36(12):1285–98. Available from: <http://dx.doi.org/10.1080/0144929x.2017.1381166>.
4. 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. In: *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*. New York, NY, USA: ACM; 2017. Available from: <https://dl.acm.org/doi/10.1145/3121050.3121074>.
5. Agudo U, Matute H. The influence of algorithms on political and dating decisions. *PLOS One*. 2021; 16(4):e0249454. Available from: <https://doi.org/10.1371/journal.pone.0249454> PMID: 33882073
6. Epstein R, Robertson RE. Why is the Search Engine Manipulation Effect (SEME) so large? Testing an operant conditioning hypothesis. Paper presented at the 96th annual meeting of the Western Psychological Association; 2016; Long Beach, CA.
7. Epstein R, Robertson RE. Suppressing the Search Engine Manipulation Effect (SEME). *Proceedings of the ACM: Human-Computer Interaction*. 2017; 1(2). Available from: <https://doi.org/10.1145/3134677>.
8. Epstein R, Zankich VR. The surprising power of a click requirement: How click requirements and warnings affect users' willingness to disclose personal information. *PLoS One*. 2022; 17(2):e0263097. Available from: <https://doi.org/10.1371/journal.pone.0263097> PMID: 35180222

9. Epstein R, Flores A. The YouTube manipulation effect (YME): A quantification of the possible impact that the ordering of YouTube videos might have on opinions and voting preferences. SSRN [Preprint]. 2023 [cited 2023 Jul 31]. Available from: <http://dx.doi.org/10.2139/ssrn.4527207>.
10. Epstein R, Lee V, Mohr R, Zankich VR. The Answer Bot Effect (ABE): A powerful new form of influence made possible by intelligent personal assistants and search engines. PLOS ONE. 2022; 17(6): e0268081. Available from: <https://doi.org/10.1371/journal.pone.0268081> PMID: 35648736
11. Epstein R, Aries S, Grebbien K, Salcedo AM, Zankich VR. The search suggestion effect (SSE): How autocomplete search suggestions can be used to impact opinions and votes. SSRN [Preprint]. 2023 [cited 2023 Aug 8]. Available from: <http://dx.doi.org/10.2139/ssrn.4535163>.
12. Epstein R, Tyagi C, Wang H. What would happen if twitter sent consequential messages to only a strategically important subset of users? A quantification of the Targeted Messaging Effect (TME). PLOS ONE. 2023; 18(7):e0284495. Available from: <https://doi.org/10.1371/journal.pone.0284495> PMID: 37498911
13. Epstein R, Huang Y, Megerdooian M. The Opinion Matching Effect (OME): A subtle but powerful new form of influence that is apparently being used on the internet. SSRN [Preprint]. 2023 [cited 2023 Aug 4]. Available from: <http://dx.doi.org/10.2139/ssrn.4532141>.
14. 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 International AAAI Conference on Web and Social Media. 2017; 11(1):62–71. Available from: <http://dx.doi.org/10.1609/icwsm.v11i1.14898>.
15. Trielli D, Diakopoulos N. Search as news curator: The role of Google in shaping attention to news information. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. New York, NY, USA: ACM; 2019. Available from: <http://dx.doi.org/10.1145/3290605.3300683>.
16. Draws T, Tintarev N, Gadiraju U, Bozzon A, Timmermans B. This is not what we ordered: Exploring why biased search result rankings affect user attitudes on debated topics. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. New York, NY, USA: ACM; 2021. Available from: <http://dx.doi.org/10.1145/3404835.3462851>.
17. Smith A. Public attitudes toward computer algorithms. Pew Research Center: Internet, Science & Tech. 2018 [cited 2023 Sep 7]. Available from: <https://www.pewresearch.org/internet/2018/11/16/public-attitudes-toward-computer-algorithms/>.
18. Bogert E, Schechter A, Watson RT. Humans rely more on algorithms than social influence as a task becomes more difficult. Sci Rep. 2021; 11(1). Available from: <https://doi.org/10.1038/s41598-021-87480-9> PMID: 33850211
19. Chen K-J, Cheung HL. Unlocking the power of ephemeral content: The roles of motivations, gratification, need for closure, and engagement. Comput Hum Behav. 2019; 97:67–74. Available from: <http://dx.doi.org/10.1016/j.chb.2019.03.007>.
20. McKinnon JD, MacMillan D. Google workers discussed tweaking search function to counter travel ban. Wall Street journal [Internet]. 2018 Sep 21 [cited 2023 Sep 7]; Available from: <https://www.wsj.com/articles/google-workers-discussed-tweaking-search-function-to-counter-travel-ban-1537488472>.
21. Cadwalladr C. Google is not 'just' a platform. It frames, shapes and distorts how we see the world. The Guardian [Internet]. 2016 Dec 11 [cited 2023 Sep 7]; Available from: <https://www.theguardian.com/commentisfree/2016/dec/11/google-frames-shapes-and-distorts-how-we-see-world>.
22. Grind K, Schechner S, McMillan R, West J. How Google interferes with its search algorithms and changes your results. Wall Street journal [Internet]. 2019 Nov 15 [cited 2023 Sep 7]; Available from: <https://www.wsj.com/articles/how-google-interferes-with-its-search-algorithms-and-changes-your-results-11573823753>.
23. Walker K. Supporting election integrity through greater advertising transparency [Internet]. Google. 2018 [cited 2023 Sep 17]. Available from: <https://blog.google/outreach-initiatives/public-policy/supporting-election-integrity-through-greater-advertising-transparency/>.
24. Horling B, Kulick M. Personalized Search for everyone [Internet]. Google Official Blog. 2009 [cited 2023 Sep 20]. Available from: <https://googleblog.blogspot.com/2009/12/personalized-search-for-everyone.html>.
25. Arora N, Ensslen D, Fiedler L, Liu WW, Robinson K, Stein E, et al. The value of getting personalization right—or wrong—is multiplying [Internet]. Mckinsey. 2021 [cited 2023 Sep 21]. Available from: https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-value-of-getting-personalization-right-or-wrong-is-multiplying#.
26. Shanahan T, Tran TP, Taylor EC. Getting to know you: Social media personalization as a means of enhancing brand loyalty and perceived quality. J Retail Consum Serv. 2019; 47:57–65. Available from: <http://dx.doi.org/10.1016/j.jretconser.2018.10.007>.

27. Erevelles S, Fukawa N, Swayne L. Big Data consumer analytics and the transformation of marketing. *J Bus Res.* 2016; 69(2):897–904. Available from: <http://dx.doi.org/10.1016/j.jbusres.2015.07.001>.
28. van Reijmersdal EA, Boerman SC, van Noort G. Effects of online behaviorally targeted native advertising on persuasion: A test of two competing mechanisms. *Comput Hum Behav Rep.* 2022; 7 (100221):100221. Available from: <http://dx.doi.org/10.1016/j.chbr.2022.100221>.
29. Rahat G, Kenig O. From party politics to personalized politics?: Party change and political personalization in democracies. Oxford University Press; 2018. Available from: <http://dx.doi.org/10.1093/oso/9780198808008.001.0001>.
30. Robertson RE, Lazer D, Wilson C. Auditing the personalization and composition of politically-related search engine results pages. In: Proceedings of the 2018 World Wide Web Conference on World Wide Web [Internet]. ACM Press; 2018. Available from: <http://dx.doi.org/10.1145/3178876>.
31. Bleu N. 19+ crucial search engine statistics for 2023 (market share, usage, & demographics) [Internet]. Startup Bonsai. 2022 [cited 2023 Sep 27]. Available from: <https://startupbonsai.com/search-engine-statistics/>.
32. Search engine market share worldwide [Internet]. StatCounter Global Stats. [cited 2023 Sep 27]. Available from: <https://gs.statcounter.com/search-engine-market-share>.
33. Nadler J, Cicilline DN. Investigation of Competition in Digital Markets [Internet]. 2020. Available from: https://fm.cnb.com/applications/cnbc.com/resources/editorialfiles/2020/10/06/investigation_of_competition_in_digital_markets_majority_staff_report_and_recommendations.pdf.
34. Meyers L. Google Manipulates Search Results, Former Engineer Shockingly Confirms [Internet]. CCN. 2023 [cited 2023 Sep 27]. Available from: <https://www.ccn.com/google-manipulates-search-results-former-engineer-shockingly-confirms/>.
35. Epstein R. Taming Big Tech: The case for monitoring [Internet]. HackerNoon. 2018 [cited 2023 Sep 27]. Available from: <https://hackernoon.com/taming-big-tech-5fef0df0f00d>.
36. Epstein R, Bock S, Peirson L, Voillot M, Wang H. How we preserved more than 1.5 million online “ephemeral experiences” in the recent US elections, and what this content revealed about online election bias. In: 102nd meeting of the Western Psychological Association. Portland, OR; 2022 [cited 2023 Sep 27]. Available from: https://aibr.org/downloads/EPSTEIN_et_al_2022-WPA-How_We_Preserved_More_Than_1.5_Million_Online_Ephemeral_Experiences_in_Recent_US_Elections...pdf.
37. Epstein R. America's Digital Shield: A new online monitoring system will make Google and other tech companies accountable to the public [Written testimony, 480 pp.]. Congressional Record. 2023 Dec 13. Available from: https://www.judiciary.senate.gov/imo/media/doc/2023-12-13_pm_-_testimony_-_epstein.pdf.
38. Prinz R, Brighton H, Luan S, Gigerenzer G. Can biased search engine results influence healthcare decisions? In: International Convention of Psychological Science. Vienna, Austria; 2017. Available from: https://aibr.org/downloads/ICPS-Symposium-Subtle_New_Forms_of_Internet_Influence-24March2017-Vienna.pdf.
39. 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. Available from: <https://doi.org/10.2196/jmir.2642> PMID: 24694866
40. Loftus EF. Leading questions and the eyewitness report. *Cogn Psychol.* 1975; 7(4):560–72. Available from: [http://dx.doi.org/10.1016/0010-0285\(75\)90023-7](http://dx.doi.org/10.1016/0010-0285(75)90023-7).
41. Nyhan B, Reifler J. When corrections fail: The persistence of political misperceptions. *Polit Behav.* 2010; 32(2):303–30. Available from: <http://dx.doi.org/10.1007/s11109-010-9112-2>.
42. Nield D. All the ways Google tracks you—and how to stop it. *Wired* [Internet]. 2019 May 27 [cited 2023 Oct 9]; Available from: <https://www.wired.com/story/google-tracks-you-privacy/>.
43. Hunag K. Email Client Market Share in February 2022 [Internet]. Litmus. 2022 [cited 2023 Oct 7]. Available from: <https://www.litmus.com/blog/email-client-market-share-february-2022/>.
44. StatCounter. Browser version (partially combined) market share worldwide [Internet]. StatCounter Global Stats. 2023 [cited 2023 Oct 7]. Available from: <https://gs.statcounter.com/browser-version-partially-combined-market-share>.
45. StatCounter. Mobile operating system market share worldwide [Internet]. StatCounter Global Stats. 2023 [cited 2023 Oct 7]. Available from: <https://gs.statcounter.com/os-market-share/mobile/worldwide/>.
46. Editor of IndexRise. List of Google products and services [Internet]. Indexrise. [cited 2023 Oct 9]. Available from: <https://www.indexrise.com/google-services/>.
47. Nakashima R. AP Exclusive: Google tracks your movements, like it or not [Internet]. AP News. 2018 [cited 2023 Oct 9]. Available from: <https://apnews.com/article/828aefab64d4411bac257a07c1af0ecb>.

48. Epstein R, Lothringer M, Zankich VR. How a daily regimen of operant conditioning might explain the power of the search engine manipulation effect (SEME) [Internet]. *Behav Soc Iss* 2024. Available from: <https://doi.org/10.1007/s42822-023-00155-0>.
49. Ong SQ, Hardwick J, Soulo T. Top Google searches [Internet]. Ahrefs Blog. Ahrefs; 2023 [cited 2023 Oct 9]. Available from: <https://ahrefs.com/blog/top-google-searches/>.
50. Langville AN, Meyer CD. *Google's PageRank and beyond: The science of search engine rankings*. Princeton, NJ: Princeton University Press; 2012.
51. Lee NT, Resnick P, Barton G. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms [Internet]. Brookings. 2019 [cited 2023 Oct 9]. Available from: <https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>.
52. Nunez M. Former Facebook workers: We routinely suppressed conservative news [Internet]. Gizmodo. 2016 [cited 2023 Oct 9]. Available from: <https://gizmodo.com/former-facebook-workers-we-routinely-suppressed-conser-1775461006>.
53. Rainie L, Anderson J. Code-dependent: Pros and cons of the algorithm age [Internet]. Pew Research Center: Internet, Science & Tech. 2017 [cited 2023 Oct 9]. Available from: <https://www.pewresearch.org/internet/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age/>.
54. Arthur Bloom J. EXCLUSIVE: Documents detailing Google's 'news blacklist' show manual manipulation of special search results [Internet]. The Daily Caller. 2019 [cited 2023 Oct 9]. Available from: <https://dailycaller.com/2019/04/09/google-news-blacklist-search-manipulation/>.
55. Ebbinghaus H. *Memory: A contribution to experimental psychology*. New York: Teachers College Press; 1913.
56. Murre JMJ, Dros J. Replication and analysis of Ebbinghaus' forgetting curve. *PLoS One*. 2015; 10(7): e0120644. Available from: <https://doi.org/10.1371/journal.pone.0120644> PMID: 26148023
57. Mack CC, Cinel C, Davies N, Harding M, Ward G. Serial position, output order, and list length effects for words presented on smartphones over very long intervals. *J Mem Lang*. 2017; 97:61–80. Available from: <https://doi.org/10.1016/j.jml.2017.07.009> PMID: 29200611
58. Dean B. We analyzed 4 million Google search results. Here's what we learned about organic click through rate [Internet]. Backlinko. 2019 [cited 2023 Oct 9]. Available from: <https://backlinko.com/google-ctr-stats>.
59. Epstein R, Newland A, Tang L. The Digital Personalization Effect (DPE): How personalization of online content can dramatically increase the impact of new forms of online influence. In: Paper to be presented at the 104th annual meeting of the Western Psychological Association. San Francisco, CA; 2024.
60. Krishnan P. Extend the reach of your site personalization in Google Optimize [Internet]. Google. 2019 [cited 2023 Oct 9]. Available from: <https://blog.google/products/marketingplatform/analytics/extend-reach-your-site-personalization-google-optimize/>.
61. Bai H. Evidence that A Large Amount of Low Quality Responses on MTurk Can Be Detected with Repeated GPS Coordinates [Internet]. 2018 [cited 2023 Oct 9]. Available from: <https://www.maxhuibai.com/blog/evidence-that-responses-from-repeating-gps-are-random>.
62. Dreyfuss E. A bot panic hits amazon's mechanical Turk. *Wired* [Internet]. 2018 Aug 17 [cited 2023 Oct 9]; Available from: <https://www.wired.com/story/amazon-mechanical-turk-bot-panic/>.
63. Litman L. Strengths and limitations of Mechanical Turk [Internet]. CloudResearch. 2017 [cited 2023 Oct 9]. Available from: <https://www.cloudresearch.com/resources/blog/strengths-and-limitations-of-mechanical-turk/>.
64. Moss A., & Litman L. (2020, June 12). *Demographics of people on Amazon Mechanical Turk*. CloudResearch. <https://www.cloudresearch.com/resources/blog/who-uses-amazon-mturk-2020-demographics/>
65. Huff C, Tingley D. "Who are these people?" Evaluating the demographic characteristics and political preferences of MTurk survey respondents. *Res Politics*. 2015; 2(3):205316801560464. Available from: <http://dx.doi.org/10.1177/2053168015604648>.
66. Casler K, Bickel L, Hackett E. Separate but equal? A comparison of participants and data gathered via Amazon's MTurk, social media, and face-to-face behavioral testing. *Comput Hum Behav*. 2013; 29(6):2156–60. Available from: <http://dx.doi.org/10.1016/j.chb.2013.05.009>.
67. Ramsey SR, Thompson KL, McKenzie M, Rosenbaum A. Psychological research in the internet age: The quality of web-based data. *Comput Hum Behav*. 2016; 58:354–60. Available from: <http://dx.doi.org/10.1016/j.chb.2015.12.049>.
68. Epstein R, Ding M, Mourani C, Newland A, Olson E, Tran F. Multiple searches increase the impact of similarly biased search results: An example of the Multiple Exposure Effect (MEE). SSRN [Preprint]. 2023 [cited 2023 Nov 17]. Available from: <https://dx.doi.org/10.2139/ssrn.4636728>.

69. Hu J, Zeng H-J, Li H, Niu C, Chen Z. Demographic prediction based on user's browsing behavior. In: Proceedings of the 16th international conference on World Wide Web. New York, NY, USA: ACM; 2007.
70. Yom-Tov E. Demographic differences in search engine use with implications for cohort selection. *Inf Retr Boston*. 2019; 22(6):570–80. Available from: <http://dx.doi.org/10.1007/s10791-018-09349-2>.
71. Chen L, Ma R, Hannák A, Wilson C. Investigating the impact of gender on rank in resume search engines. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems [Internet]. New York, NY, USA: ACM; 2018. Available from: <http://dx.doi.org/10.1145/3173574.3174225>.
72. Merton RK. The unanticipated consequences of purposive social action. *Am Sociol Rev*. 1936; 1(6):894–904. Available from: <http://dx.doi.org/10.2307/2084615>.