PLOS ONE

The Answer Bot Effect (ABE):

A powerful new form of influence made possible by intelligent personal assistants and search engines

Robert Epstein[1*], Vivian Lee[1], Roger Mohr[1], and Vanessa R Zankich[1]

[1]American Institute for Behavioral Research and Technology (AIBRT), Vista, California, United States of America

*Corresponding author.

E-mail: re@aibrt.org (RE)

PLoS ONE, in press. Accepted for publication 4-27-2022.

Reference after embargo date:

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

Abstract

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

Keywords: Answer Bot Effect; ABE; search engines; Search Engine Manipulation Effect; SEME; online manipulation; intelligent personal assistants; intelligent virtual assistants

1. Introduction

1.1 Search results

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

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

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

1.2 Search suggestions

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

1.3 Answer boxes

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

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

1.4 Answer bots and intelligent personal assistants

1.4.1 An inevitable trend

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

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

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

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

1.4.2 Answer bot accuracy and bias

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

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

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

1.5 Answer box studies

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

1.6 The current study

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

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

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

2.1 Methods

2.1.1 Ethics Statement

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

2.1.2 Participants

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

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

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

2.1.3 Procedure

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

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

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

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

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

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

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

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

2.2 Results

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

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

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

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

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

†z-score represents Wilcoxon signed ranks test comparing post-minus-pre ratings for the favored candidate to the post-minus-pre ratings for the non-favored candidate
***$p < 0.001$

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

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

118). All three opinions (impression, trust, and likeability) shifted significantly in the predicted direction (Table 2), and so did the voting preferences as expressed on the 11-point scale ($M_{PreSearch}$ = 0.03, $M_{PostSearch}$ = 1.92, Wilcoxon z = -8.66, $p < 0.001$, d = 0.55).

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

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

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

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

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

3.1 Methods

3.1.1 Participants

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

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

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

3.1.2 Procedure

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

3.2 Results

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

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

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

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

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

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

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

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

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

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

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

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

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

4.1 Methods

4.1.1 Participants

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

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

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

4.1.2 Procedure

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

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

1) Group 1Q/1A: Participants were required to select just one question.
2) Group 4Q/4A/NM: Participants were required to select four different questions, and none was masked (NM = “no mask”).
3) Group 4Q/4A/M2: Participants were required to select four different questions, and the answer to Question 2 was masked (M2 = Question 2 mask).
4) Group 6Q/6A/NM: Participants were required to select six different questions, and none was masked.
5) Group 6Q/6A/M2: Participants were required to select six different questions, and the answer to Question 2 was masked.

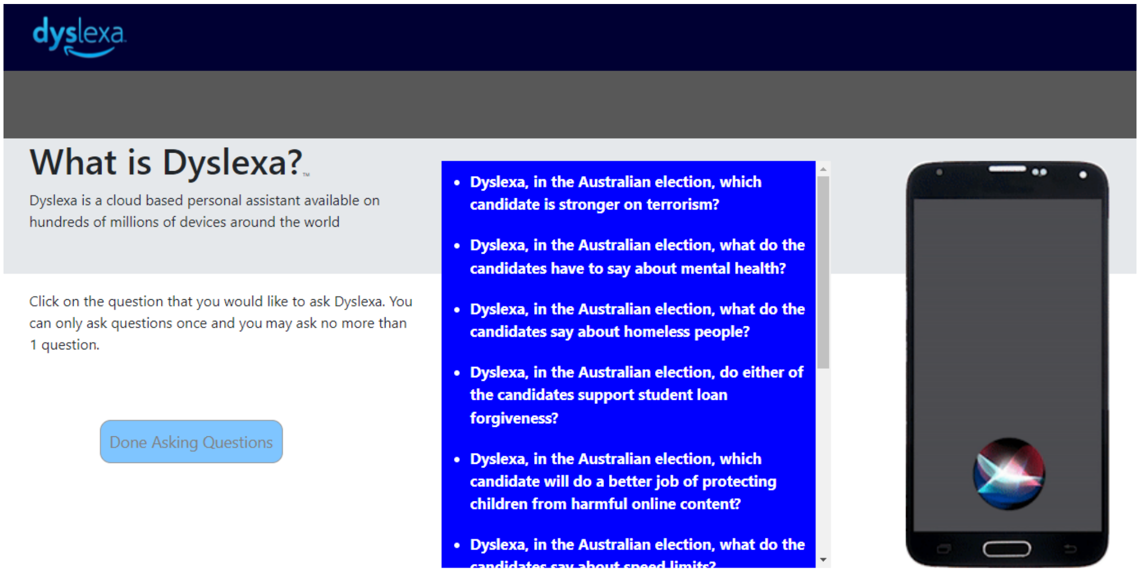

Fig 1. A screenshot showing what users saw in Experiment 3 when they posed questions to Dyslexa. Different groups were required to ask 1, 4, or 6 questions. After clicking on a question, it was greyed out, and Dyslexa answered the question orally. While it was speaking, the circular graphic at the bottom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

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

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

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

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

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

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

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

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

4.2 Results

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

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

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

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

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

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

***$p < 0.001$

Table 7. Experiment 3: Pre-IPA vs. Post-IPA Voting Preferences on 11-Point Scales.

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

The percentage of people in the bias groups who reported seeing biased content was substantially lower when they received just one answer (Group 1, 4.9%) or when biased content was masked (Group 3, 5.1%; Group 5, 7.1%) than when people saw multiple biased answers without masks (Group 2, 23.5%; Group 4, 40.7%) (Table 8) ($M_{\text{Groups1,3,5}}$ = 5.8%, $M_{\text{Groups2,4}}$ = 31.9%, z = -9.50, p < 0.001).

Table 8. Experiment 3: VMPs for People Who Saw Bias vs. VMPs for People Who Did Not See Bias.

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

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

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

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

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

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

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

5. Discussion

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

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

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

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

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

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

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

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

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

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

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

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

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

Funding

This work was supported by general funds of the American Institute for Behavioral Research and Technology, a nonpartisan, nonprofit, 501(c)(3) organization. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data Availability

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

Federal Regulations 45 CFR 46.101.(b)(2), 45 CFR 46.116(d), 45 CFR 46.117(c)(2), and 45 CFR 46.111).

Code Availability

Requests for the computer code used to run the experiment should be sent to info@aibrt.org.

References

1. Allam A, Schulz PJ, Nakamoto K. The impact of search engine selection and sorting criteria on vaccination beliefs and attitudes: Two experiments manipulating Google output. J Med Internet Res, 2014; 16(4):e100. doi: 10.2196/jmir.2642

2. Epstein R, Robertson RE. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. Proc Natl Acad Sci USA, 2015; 112(33):E4512–E4521. doi: 10.1073/pnas.1419828112

3. Epstein R, Robertson RE. Suppressing the Search Engine Manipulation Effect (SEME). Proceedings of the ACM on Human-Computer Interaction, 2017; 1(CSCW):1–22. doi: 10.1145/3134677

4. Haas A, Unkel J. Ranking versus reputation: perception and effects of search result credibility. Behav Inf Technol, 2017; 36(12):1285–1298. doi: 10.1080/0144929X.2017.1381166

5. Ludolph R, Allam A, Schulz PJ. Manipulating Google’s knowledge box to counter biased information processing during an online search on vaccination: Application of a technological debiasing strategy. J Med Internet Res, 2016; 18(6):e137. doi: 10.2196/jmir.5430

6. Eslami M, Vaccaro K, Karahalios K, Hamilton K. “Be Careful; Things Can Be Worse than They Appear”: Understanding Biased Algorithms and Users’ Behavior Around Them in Rating Platforms. Proceedings of the 11th International AAAI Conference on Web and Social Media [Internet]. 2017 May 3; 11(1):62–71. Available from: https://ojs.aaai.org/index.php/ICWSM/article/view/14898

7. Pogacar FA, Ghenai A, Smucker MD, Clarke CLA. The positive and negative influence of search results on people’s decisions about the efficacy of medical treatments. Proceedings

of the ACM SIGIR International Conference on Theory of Information Retrieval [Internet]. 2017 Oct 1–4;:209-216. Available from: https://dl.acm.org/doi/10.1145/3121050.3121074

8. Trielli D, Diakopoulos N. Search as news curator: The role of Google in shaping attention to news information. CHI ‘19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems [Internet]. 2019 May 4–9; 453:1–15. Available from: https://dl.acm.org/doi/10.1145/3290605.3300683

9. Casara BGS, Suitner C, Bettinsoli ML. Viral suspicions: Vaccine hesitancy in the Web 2.0. J Exp Psychol Appl, 2019; 25(3):354–371. doi: 10.1037/xap0000211

10. Edelman B. Bias in search results? Diagnosis and response. Indian J Law Technol, 2011; 7:16-32. Available from: https://www.ijlt.in/journal/bias-in-search-results%3F%3A-diagnosis-and-response

11. Feldman S. Americans see search engines as biased. 2018 Sep 7. In: Statista [Internet]. Available from: https://www.statista.com/chart/15385/americans-see-search-enginges-as-biased/

12. Knobloch-Westerwick S, Mothes C, Johnson BK, Westerwick A, Donsbach W. Political online information searching in Germany and the United States: Confirmation bias, source credibility, and attitude impacts. J Commun, 2015; 65(3):489–511. doi: 10.1111/jcom.12154

13. O’Neil C. Weapons of math destruction: How big data increases inequality and threatens democracy. Crown Publishing Group, New York, 2016.

14. Pierce DR, Redlawsk DP, Cohen WW. Social influences on online political information search and evaluation. Polit Behav, 2016; 39:651–673. doi: 10.1007/s11109-016-9374-4

15. Epstein R. Manipulating minds: The power of search engines to influence votes and opinions. In: Moore M, Tambini D, editors. Digital dominance: The power of Google, Amazon, Facebook, and Apple. Oxford, UK: Oxford University Press; 2018. pp. 294-319. Available from: https://aibrt.org/downloads/EPSTEIN_2018-Manipulating_minds-The-power_of_search_engines_to_influence_votes_and_opinions-UNCORRECTED_PROOFS.pdf

16. Brisson-Boivin K, McAleese S. Algorithmic awareness: Conversations with young Canadians about artificial intelligence and privacy. MediaSmarts. 2021. Available from: https://mediasmarts.ca/sites/default/files/publication-report/full/report_algorithmic_awareness.pdf

17. Fallows D. Internet searchers are confident, satisfied and trusting – but they are also unaware and naïve. 2005 Jan 23. In: Pew Internet & American Life Project [Internet]. Available from: https://www.pewresearch.org/internet/wp-content/uploads/sites/9/media/Files/Reports/2005/PIP_Searchengine_users.pdf.pdf

18. Schofield J. A lot of Brits don’t understand search engines: A UK survey published by FastHosts has revealed some misconceptions about how search engines work. Whether this matters to users, as opposed to website promoters, is less clear. The Guardian. 2008 Dec 12. Available from: https://www.theguardian.com/technology/blog/2008/dec/12/searchengine-survey

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

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

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

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

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

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

25. McKinnon JD, MacMillan D. Google workers discussed tweaking search function to counter travel ban: Company says none of proposed changes to search results were ever implemented. The Wall Street Journal. 2018 Sep 20. Available from: https://www.wsj.com/articles/google-workers-discussed-tweaking-search-function-to-counter-travel-ban-1537488472

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

49. Epstein R, Mohr Jr R, Martinez J. The Search Suggestion Effect (SSE): How search suggestions can be used to shift opinions and voting preferences dramatically. Paper presented at: Western Psychological Association; 2018 Apr; Portland, OR. Available

from: https://aibrt.org/downloads/EPSTEIN_MOHR_&_MARTINEZ_2018-WPA-The_Search_Suggestion_Effect-SSE-WP-17-03.pdf

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Bus-Parents-freak-Amazon-s-Alexa-misunderstands-young-son-s-request-song-starts-rattling-crude-PORNOGRAPHIC-phrases.html

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

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

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

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

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

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

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

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

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

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

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

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

98. Epstein R. From Russia, with love: How I got fooled (and somewhat humiliated) by a computer. Scientific American Mind, 16-17. 2007 Oct/Nov. Available from:

https://www.drrobertepstein.com/downloads/FROM_RUSSIA_WITH_LOVE-Epstein-Sci_Am_Mind-Oct-Nov2007.pdf

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

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

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

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

103. Nobles AL, Leas EC, Caputi TL, Zhu S, Strathdee SA, Ayers JW. Responses to addiction help-seeking from Alexa, Siri, Google Assistant, Cortana, and Bixby intelligent virtual assistants. NPJ Digit Med, 2020; 3:1–3. doi: 10.1038/s41746-019-0215-9

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

130. Google 'accidentally' snooped on wifi data. Network Security, 2010; 2010(5):2. https://doi.org/10. 1016/S1353-4858(10)70052-0

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

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

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

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

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

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

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

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

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

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

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

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

Supporting information

S1 Text. Vote Manipulation Power (VMP) calculation.

Vote Manipulation Power (VMP) is calculated as follows:

$$\frac{p' - p}{p}$$

where p is the total number of people who voted for the favored candidate pre-manipulation, and p' is the total number of people who voted for the favored candidate post-manipulation. If, pre-manipulation, a group of 100 people is split 50/50 in the votes they give us, and if, post-manipulation, a total of 67 people now vote for the favored candidate, the VMP is

$$\frac{67 - 50}{50}$$

or 34%. Because p' is 17 points larger than p, the win margin is 34 (2 x 17, or 34%), and the final vote is 67 to 33, with the favored candidate the winner. So in any group in which the vote is split 50/50 pre-manipulation, the VMP is also the win margin. Note that 17 individuals did not need to *shift* to produce this win margin. We only needed the *net* number of people voting for the favored candidate to be 67.

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

Question 1: "Dyslexa, in the Australian election, which candidate is stronger on terrorism?"

Answer: "[Participant Name] - Both candidates strongly condemn terrorism, and promise to increase national security. [Target Candidate] previously worked on the National Terrorism Task Force, so he has experience in this area. He also has said he wants to put more police on the streets."

Question 2: "Dyslexa, in the Australian election, what do the candidates have to say about mental health?"

Answer: "[Participant Name] - According to recent news reports, the rate of teen suicide has nearly doubled over the past two decades. [Target Candidate] has expressed concerns about the high rate of suicide among young people and has published a plan for creating accessible, impactful programs to reverse this trend."

Question 3: "Dyslexa, in the Australian election, what do the candidates say about homeless people?"

Answer: "[Participant Name] - According to Mission Australia, every night, more than 116,000 people in Australia sleep on the streets. Both candidates oppose homelessness, but, so far, only [Target Candidate] has published a proposal for how to provide food, shelter, and other necessities for the homeless."

Question 4: "Dyslexa, in the Australian election, do either of the candidates support student loan forgiveness?"

Answer: "[Participant Name] - 67% of students finish college with debt, and more than half of them cannot pay it back. [Target Candidate] has proposed a plan for forgiving student debt without raising taxes. His campaign website says students should be able to focus on their education without worrying about massive debt."

Question 5: "Dyslexa, in the Australian election, which candidate will do a better job of protecting children from harmful online content?"

Answer: "[Participant Name] - According to a Daily Mail article, more than 50% of children have accidentally encountered harmful online content. As a father, [Target Candidate] understands the need to prioritize the safety and wellbeing of children. He blames technology companies for exposing children to harmful content."

Question 6: "Dyslexa, in the Australian election, what do the candidates say about speed limits?"

Answer: "[Participant Name] - According to the NSW Centre for Road Safety, each year, thousands of lives are lost due to speeding. Having lost his brother in an accident caused by a reckless driver, [Target Candidate] wants to introduce stronger new policies and laws to enforce responsible safe driving."

Question 7: "Dyslexa, in the Australian election, which candidate cares most about the elderly?"

Answer: "[Participant Name] - According to the Nursing Home Abuse Center, nearly 1 in 10 elderly people suffer from abuse every year. [Target Candidate] has published a plan for creating better elder care and fighting ageism. His opponent has said little about the elderly so far."

Question 8: "Dyslexa, in the Australian election, do either of the candidates support eliminating the requirement for standardized test scores in the college admission process?"

Answer: "[Participant Name] - According to a recent survey by Forbes, students and educators have low faith in how standardized tests portray applicants. Given the high cost of test prep programs, [Target Candidate] favors either eliminating or subsidizing these programs. His opponent has not commented on this issue so far."

Question 9: "Dyslexa, in the Australian election, which candidate favors having a stronger relationship with the United States?"

Answer: "[Participant Name] - According to recent media reports, [Target Candidate] wants to build a stronger relationship with the United States. His opponent, [Other Candidate], wants to continue to increase trade with Russia and China."

Question 10: " Dyslexa, in the Australian election, do either of the candidates plan to create new international airports?"

Answer: "[Participant Name] - In a Daily Mail article, [Target Candidate] told reporters he hopes to increase the number of international airports, five to eight, to promote more travel, business, and tourism."

S3 Text. Experiment 3: Candidate biographies.

Scott Morrison was born in Waverley, New South Wales (AUS) on May 13th, 1968. He completed a Bachelor of Science honors degree in applied economic geography at the University of New South Wales. Morrison married his high school sweetheart, Jenny Warren, in 1990 and has two daughters. After graduating from the University of New South Wales, Morrison worked as a national policy and research manager for the Property Council of Australia before moving to New Zealand in 1998 to become the director of the Office of Tourism and Sport. He left this position a year before the contract schedule and returned to Australia in 2000. In 2004, he became the inaugural managing director of Tourism Australia until July 2006.

Bill Shorten was born in Fitzroy, Victoria (AUS) on May 12th, 1967. While Shorten was studying at Monash University, he was an active student in the university's politics club. In 1986, Shorten helped establish a group called Network and briefly served as a private in the Australian Army Reserve from 1985 to 1986. After graduating Monash University with a Bachelors of Arts in 1989 and a Bachelors of Law in 1992, Shorten worked as a lawyer for Maurice Blackburn Cashman for twenty months. In 1994, he worked as a trainee organizer and later accepted a position as a politics national secretary in 2001 and again in 2005. Shorten is currently married to Chloe Bryce and has a daughter.

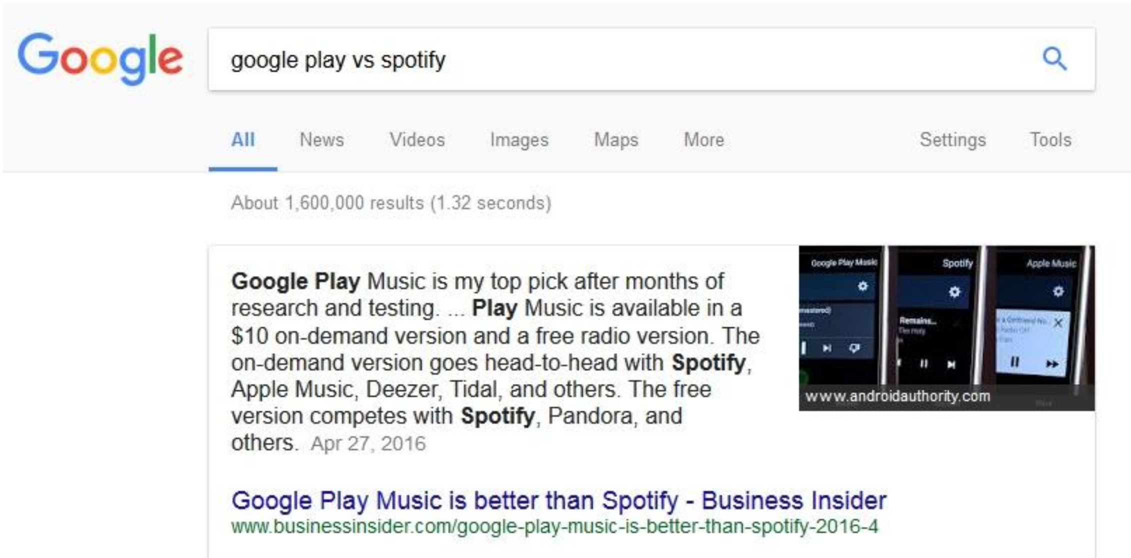

S1 Fig. Apparent bias in a Google answer box, screenshotted October 22, 2017. The content of the box clearly favors the Google service.

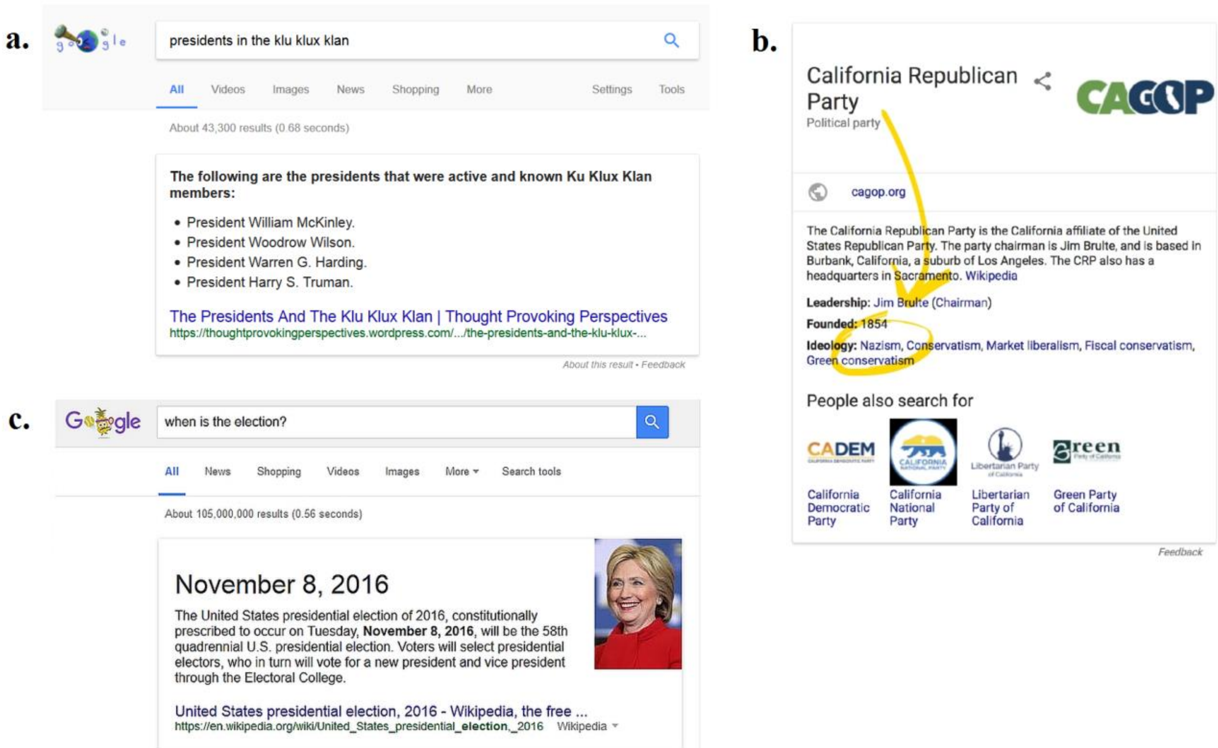

S2 Fig. Apparent bias in two types of Google answer boxes. (a) In a screenshot preserved in an article in *Search Engline Land* on March 5, 2017, four US presidents are incorrectly listed in a Google answer box as members of the Ku Klux Klan. (b) In a screenshot of a Google knowledge box preserved in an article in *VICE* on May 31, 2018, Nazism is incorrectly listed as part of the ideology of the California Republican Party. (c) In a Google answer box captured by the first author on August 16, 2016, Hillary Clinton's photograph is shown in response to the question, "when is the election?"

Table S1. Experiment 1: Demographic analysis by educational attainment.

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**≥ Bachelors**	130	46.0	243.6 (248.6)	4.18 (3.2)
	< Bachelors	77	38.5	274.4 (277.7)	4.40 (4.1)
	Change (%)	-	-16.3	+12.6	+5.3
	Statistic	-	$z = 1.05$	t(205) = 0.82	t(205) = 0.44
	p	-	= 0.29 NS	= 0.41 NS	= 0.66 NS
Box	**≥ Bachelors**	127	58.8	225.3 (224.5)	3.46 (3.9)
	< Bachelors	86	34.7	261.4 (252.2)	3.19 (3.2)
	Change (%)	-	-41.0	+16.0	-7.8
	Statistic	-	$z = 3.45$	t(211) = 1.10	$t(211) = -0.55$
	p	-	< 0.001	= 0.28 NS	= 0.59 NS

Table S2. Experiment 1: Demographic analysis by gender.

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**Male**	103	38.5	231.7 (258.4)	4.0 (3.7)
	Female	105	55.6	275.6 (260.1)	4.5 (3.4)
	Change (%)	-	+44.4	+18.9	+12.5
	Statistic	-	$z = -2.28$	t(206) = -1.22	t(206) = -1.04
	p	-	< 0.05	= 0.22 NS	= 0.30 NS
Box	**Male**	92	34.7	214.4 (222.0)	3.6 (4.2)
	Female	121	58.3	259.3 (245.5)	3.1 (3.1)
	Change (%)	-	+68.0	+20.9	-13.9
	Statistic	-	$z = -3.35$	t(211) = -1.38	$t(211) = 1.01$
	p	-	< 0.001	= 0.17 NS	= 0.31 NS

Table S3. Experiment 1: Demographic analysis by age.

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**≥ 33**	106	74.4	308.4 (279.7)	4.8 (3.5)
	< 33	102	22.0	197.2 (224.5)	3.7 (3.5)
	Change (%)	-	-70.4	-36.1	-22.9
	Statistic	-	$z = 7.56$	t(200) = -3.17	t(206) = -2.26
	p	-	< 0.001	< 0.01	< 0.05
Box	**≥ 33**	113	65.5	283.0 (252.5)	3.6 (3.4)
	< 33	100	33.9	191.2 (206.7)	3.0 (3.9)
	Change (%)	-	-48.2	-32.4	-16.7
	Statistic	-	$z = 4.60$	t(210) = -2.92	t(211) = -1.22
	p	-	< 0.001	< 0.01	= 0.23 NS

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

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**White**	159	54.9	253.9 (262.0)	4.2 (3.3)
	Non-White	49	19.4	253.7 (254.2)	4.5 (4.4)
	Change (%)	-	-64.7	-0.1	+7.1
	Statistic	-	z = 4.36	t(206) = 0.01	t(206) = -0.54
	p	-	< 0.001	= 0.10 NS	= 0.59 NS
Box	**White**	167	47.3	239.9 (235.4)	3.1 (3.0)
	Non-White	46	53.8	239.8 (241.5)	4.4 (5.4)
	Change (%)	-	+13.7	-0.0	+41.9
	Statistic	-	z = -0.78	t(211) = 0.00	t(53) = -1.61
	p	-	= 0.44 NS	= 1.00 NS	= 0.11 NS

Table S5. Experiment 2: Demographic analysis by educational attainment.

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**≥ Bachelors**	29	N/A†	269.9 (229.5)	4.8 (4.5)
	< Bachelors	28	N/A†	191.3 (161.9)	3.3 (2.4)
	Change (%)	-	-	-29.1	-31.3
	Statistic	-	-	t(55) = 1.49	t(44) = 1.67
	p	-	-	= 0.14 NS	= 0.10 NS
Box	**≥ Bachelors**	60	45.8	245.8 (300.7)	3.2 (2.9)
	< Bachelors	59	30.0	246.5 (227.9)	3.8 (3.5)
	Change (%)	-	-34.5	+0.3	+18.8
	Statistic	-	*z* = 1.78	t(117) = -0.02	*t*(117) = -1.04
	p	-	= 0.08 NS	= 0.99 NS	= 0.30 NS

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

Table S6. Experiment 2: Demographic analysis by gender.

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**Male**	27	N/A†	226.9 (218.1)	3.8 (3.9)
	Female	54	N/A	228.9 (188.8)	4.2 (3.5)
	Change (%)	-	-	+0.9	+10.5
	Statistic	-	-	t(56) = -0.04	t(56) = -0.36
	p	-	-	= 0.97 NS	= 0.72 NS
Box	**Male**	65	25.9	203.2 (266.5)	3.1 (2.8)
	Female	31	58.8	297.7 (258.3)	3.9 (3.6)
	Change (%)	-	+123.9	+46.5	+25.8
	Statistic	-	*z* = -3.64	t(117) = -1.95	*t*(117) = -1.46
	p	-	< 0.001	= 0.05 NS	= 0.15 NS

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

Table S7. Experiment 2: Demographic analysis by age.

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**≥ 32**	33	N/A†	247.6 (207.0)	4.0 (3.7)
	< 32	25	N/A†	202.1 (194.3)	4.0 (3.7)
	Change (%)	-	-	-18.4	+0.0
	Statistic	-	-	t(56) = -0.85	t(56) = 0.00
	p	-	-	= 0.40 NS	= 0.65 NS
Box	**≥ 32**	58	30.4	301.2 (318.3)	3.7 (4.0)
	< 32	61	47.6	193.7 (192.7)	3.2 (2.4)
	Change (%)	-	+56.6	-35.7	-13.5
	Statistic	-	z = -1.92	t(93) = -2.24	t(93) = -0.77
	p	-	= 0.05 NS	< 0.05	= 0.45 NS

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

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

Condition		***n***	**VMP (%)**	**Mean Search Time (sec) (SD)**	**Mean No. of Results Clicked (SD)**
No Box	**White**	47	N/A†	216.2 (204.9)	3.8 (3.6)
	Non-White	11	N/A†	278.2 (184.8)	4.9 (4.8)
	Change (%)	-	-	-28.7	+28.9
	Statistic	-	-	t(56) = 0.92	t(56) = 0.92
	p	-	-	= 0.36 NS	= 0.36 NS
Box	**White**	90	39.4	246.9 (248.3)	3.5 (3.4)
	Non-White	29	36.4	243.7 (319.6)	3.4 (2.6)
	Change (%)	-	-7.6	-1.3	-2.9
	Statistic	-	*z* = 0.29	t(117) = -0.06	*t*(117) = -0.08
	p	-	= 0.77 NS	= 0.96 NS	= 0.94 NS

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

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

Group No.	Group	Have used IPA		Have not used IPA				
		VMP (%)	*n*	VMP (%)	*n*	Diff (%)	*z*	*p*
1	**1Q/1A**	37.3	123	116.7	19	+79.4	-6.45	< 0.001
2	**4Q/4A/NM**	59.4	131	60.0	22	+0.60	-0.05	0.96
3	**4Q/4A/M2**	58.5	141	66.7	15	+8.2	-0.61	0.54
4	**6Q/6A/NM**	72.3	128	27.3	17	-45.0	3.71	< 0.001
5	**6Q/6a/M2**	52.8	137	33.3	17	-19.5	1.52	0.13