PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES

The narrow search effect and how broadening search promotes belief updating

Eugina Leung[a,1] and Oleg Urminsky[b]

Edited by Robert Bjork, University of California Los Angeles, Los Angeles, CA; received May 7, 2024; accepted January 19, 2025

Information search platforms, from Google to AI-assisted search engines, have transformed information access but may fail to promote a shared factual foundation. We demonstrate that the combination of users' prior beliefs influencing their search terms and the narrow scope of search algorithms can limit belief updating from search. We test this "narrow search effect" across 21 studies (14 preregistered) using various topics (e.g., health, financial, societal, political) and platforms (e.g., Google, ChatGPT, AI-powered Bing, our custom-designed search engine and AI chatbot interfaces). We then test user-based and algorithm-based interventions to counter the "narrow search effect" and promote belief updating. Studies 1 to 5 show that users' prior beliefs influence the direction of the search terms, thereby generating narrow search results that limit belief updating. This effect persists across various domains (e.g., beliefs related to coronavirus, nuclear energy, gas prices, crime rates, bitcoin, caffeine, and general food or beverage health concerns; Studies 1a to 1b, 2a to 2g, 3, 4), platforms (e.g., Google—Studies 1a to 1b, 2a to 2g, 4, 5; ChatGPT, Study 3), and extends to consequential choices (Study 5). Studies 6 and 7 demonstrate the limited efficacy of prompting users to correct for the impact of narrow searches on their beliefs themselves. Using our custom-designed search engine and AI chatbot interfaces, Studies 8 and 9 show that modifying algorithms to provide broader results can encourage belief updating. These findings highlight the need for a behaviorally informed approach to the design of search algorithms.

algorithmic search | artificial intelligence | belief updating | confirmation bias

Significance

In a time of societal polarization, the combination of people's search habits and the search tools they use being optimized for relevance may perpetuate echo chambers. We document this across various diverse studies spanning health, finance, societal, and political topics on platforms like Google, ChatGPT, AI-powered Bing, and our custom-designed search engine and AI chatbot platforms. Users' biased search behaviors and the narrow optimization of search algorithms can combine to reinforce existing beliefs. We find that algorithm-based interventions are more effective than user-based interventions to mitigate these effects. Our findings demonstrate the potential for behaviorally informed search algorithms to be a better tool for retrieving information, promoting the shared factual understanding necessary for social cohesion.

Americans find themselves divided, not merely over policy choices or governmental roles but at times even in their fundamental perceptions of the same factual reality (1), with belief polarization occurring across a variety of different domains. For instance, public opinion during the COVID-19 pandemic showed deep divisions in attitudes toward health measures like vaccination requirements and mask mandates (2). Similarly, over the last five decades, there has been an increase in belief polarization regarding environmental and climate change issues in the United States (3). Moreover, beliefs about social mobility, inequality, and immigration are also highly polarized (1). As belief polarization increasingly spreads across political, health, economic, environmental, and societal domains, it places social cohesion at risk, highlighting the need for interventions that foster a shared evidentiary foundation for societal decision-making.

In principle, search engines have the potential to facilitate social cohesion by providing shared access to broad and diverse perspectives, thereby promoting a common factual understanding among groups with different beliefs. However, as we show, both traditional search algorithms and new emerging search technologies can instead inadvertently maintain belief polarization. At issue is a fundamental dilemma in designing search algorithms: Should the goal be to narrowly optimize relevance of the results or to provide breadth of information?

Search algorithms, such as Google's PageRank, have largely been designed to optimize the relevance of search results (4–7). There are obvious merits to this approach, given the sheer amount of information available online and the fact that some online content is slanted, inaccurate, or incomplete (8). Search algorithms that screen out less relevant and lower-quality content help people navigate a challenging or even overwhelming informational environment. However, when search engines choose to provide highly relevant but narrowly focused content, there is a risk of overprecision—helping users to search for a tree while missing the forest. When users receive narrowly focused information from traditional search engines, they may fail to incorporate a broader perspective on the issues they are exploring. To the degree that next-generation AI-assisted search engines are likewise designed to synthesize and condense information, maximizing relevance and narrowing the scope of information provided, this issue will persist across technologies.

Author affiliations: [a]Marketing, Freeman School of Business, Tulane University, New Orleans, LA 70118; and [b]Marketing, Booth School of Business, University of Chicago, Chicago, IL 60637

Author contributions: E.L. and O.U. designed research; E.L. and O.U. performed research; E.L. analyzed data; and E.L. and O.U. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: wleung@tulane.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2408175122/-/DCSupplemental.

Published March 24, 2025.

Prior research on filter bubbles (9), echo chambers (10), and the search engine manipulation effect (11) has documented how algorithms can fail to promote accurate belief updating, due to the functioning of the search engine itself (e.g., providing targeted search results). In the present research, we focus on a different and complementary problem, that instead originates with the behavioral tendencies of users. We show that even when search providers do not differentially target users, echo chambers can persist because the search terms chosen by users tend to be relatively directionally narrow, reflecting their own preconceptions and beliefs.

Directionally narrow search terms are a modern manifestation of the long-documented psychological phenomenon of *confirmation bias*, the tendency to favor information that reaffirms one's own beliefs (12–14). Confirmation bias in hypothesis generation is the tendency to formulate questions that elicit affirming responses, in line with prior beliefs. Experimental participants tended to devise questions that, if answered correctly, would corroborate rather than invalidate their hypothesis (12, 13). For example, when trying to learn something about another person, people are more likely to ask questions of the other person that would receive affirmative answers that support their initial belief (15). The implication for online search is that people will generate search terms that correspond to their own prior beliefs.

However, there is a second aspect of confirmation bias, with differing implications for online search. Confirmation bias in selective attention involves paying more attention to evidence that aligns with one's beliefs (16, 17). For instance, in a classic study, when participants were exposed to diverse evidence regarding capital punishment, they predominantly acknowledged the data supporting their preconceptions, ignoring or dismissing the contradictory evidence (18). If people engage in selective-attention confirmation bias, then even exposure to a broad set of information via online search would not promote belief updating.

We document a "narrow search effect" in both traditional search-engine and AI-based information search. We propose that people often engage in directionally narrow search (i.e., specific search terms that presume one's own prior beliefs), and the algorithms' prioritization of the most relevant information can in turn amplify this human tendency for confirmation bias in hypothesis generation. When search engines provide directionally narrow search results in response to users' directionally narrow search terms, the results will reflect the users' existing beliefs, instead of promoting belief updating by providing a broad spectrum of related information.

Whether broadening search results, to counter confirmation bias in hypothesis generation, would then promote more belief updating will depend on whether people are receptive to that broader set of information. If people display selective-attention confirmation bias, focusing primarily on the subset of the broader information that supports their views, even broadening search results might fail to facilitate belief updating. This is likely to be the case in topics for which people actively resist changing their beliefs and engage in motivated reasoning. However, we propose that for many topics that people search online, a genuine interest in information may be derailed specifically by unintentional confirmation bias in hypothesis generation. When people use overly narrow search terms that align with their existing beliefs, nudging them toward more inclusive search queries or providing them with broader search results could promote belief updating and reduce the "echo chambers" created by their own search tendencies.

As an initial demonstration, consider US voters searching for information about the presidential election during the period of uncertainty between election day November 4, 2020, and the inauguration of Joe Biden on January 20, 2021. According to our proposed narrow search effect, red-state voters with more of a prior belief that Trump would win and blue-state voters anticipating a Biden win would, on average, search differently. Confirming this prediction, Google Trends data show that the higher the Republican vote share in a state, the more likely Google users in that state were to search "Trump win" or "Trump won" compared to searching "Biden win" or "Biden won" ($r = 0.53$, N = 51, $P < 0.001$; Fig. 1). To the degree that this difference in search terms yielded different results, voters in different states were getting different answers, in line with their different directionally narrow queries.

We systematically test both the narrow search effect and potential interventions across 21 studies (14 preregistered; total N = 9,906;

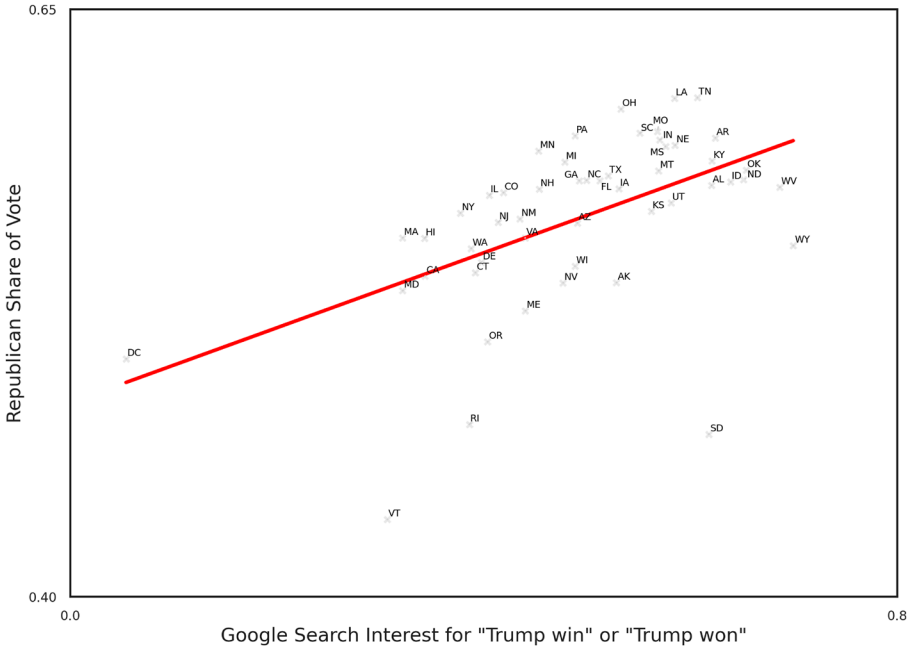

Fig. 1. Scatterplot of politically directional search by vote share across US states.

Table 1. Summary of studies

Study	Domain	Search platform
Set 1: Examine how prior beliefs affect the choice of search terms and how search terms affect belief updating and consequential choices		
1a (Aspredicted #112159)	Caffeine, Gas prices, Crime rates, Nuclear energy, Coronavirus, Bitcoin	N/A
1b (Aspredicted #189224)	Food and beverage	N/A
2a	Caffeine	Google
2b (Aspredicted #111448)	Gas prices	Google
2c (Aspredicted #111449)	Crime rates	Google
2d (Aspredicted #111450)	Nuclear energy	Google
2e	Coronavirus	Google
2f (Aspredicted #111451)	Bitcoin	Google
2g (Aspredicted #189471)	Food and beverage	Custom-designed search engine
3 (Aspredicted #144553)	Caffeine, Gas prices, Crime rates, Nuclear energy	Chat GPT-3.5
4 (Aspredicted #104627)	Caffeine	Google
5 (Aspredicted #8197)	Caffeine	Google
Set 2: Investigate possible interventions to help belief updating		
6	Caffeine	Google
7	Caffeine	Google
8a (Aspredicted #12098)	Caffeine	Custom-designed search engine
8b (Aspredicted #187347)	Thinking abilities and age	Custom-designed search engine
9a	Caffeine	Custom-designed search engine
9b (Aspredicted #198558)	Thinking abilities and age	Custom-designed AI chat bot
Additional studies in *SI Appendix*		
Study presented in the discussion (Aspredicted #143527)	Caffeine	Bing powered by GPT-4
Study presented in the discussion	N/A	Google
Supplementary study	Caffeine	Google

see Table 1 for a summary and *SI Appendix* for full details of methods). We first tested whether people's prior beliefs are reflected in the search terms they choose to use, across different topic domains, including both time-specific topics and generally relevant topics.

Results

Study 1a. 768[*] Prolific participants (M_{age} = 37, 50% female) rated their beliefs on six topics, four consistently relevant over time (the health effects of caffeine, future gas prices, future crime rates, and the impact of nuclear energy) and two that were specific to events at the time (the economic impact of coronavirus and the societal impact of bitcoin in July 2020 and November 2022, respectively). They then generated a search term to use in a Google search to learn more about the topic, and a research assistant unaware of the hypotheses rated the direction of all the search terms (e.g., 1 = More related to health risks, 7 = More related to health benefits). Overall, many of the search terms generated by participants were evaluated to be directionally narrow[†] (caffeine, 25%; gas prices, 16%; crime rates, 9%; nuclear energy, 17%, coronavirus, 34%; bitcoin, 13%). Moreover, the direction of the search term generated by each participant was consistently positively correlated with that participant's prior beliefs toward the topic ($r_{caffeine}$ = 0.13, $r_{gas\ prices}$ = 0.23, $r_{crime\ rates}$ = 0.13, $r_{nuclear\ energy}$ = 0.19, $r_{coronavirus}$ = 0.24, $r_{bitcoin}$ = 0.21; all $ps < 0.001$).

In Study 1a, we assigned participants to a specific topic to search on. However, search is typically motivated by spontaneous concerns, which may vary from person to person. In the next study, we test for narrow search when participants generate their own specific topic on which to search for information.

Study 1b. 713 Prolific participants (M_{age} = 41, 55% female) were asked to name a food or beverage for which they were uncertain about the health effects and rated their beliefs about the health effects of the specific food or beverage item they chose. Participants were then asked to generate a search term they would use to investigate the health effects of their chosen food or beverage. Finally, participants

[*]Across the studies, we use preregistered exclusions, for incompletes, duplicate IP addresses, failed attention check, and those whose search term is irrelevant. The results persist when all participants are included in the analysis; see *Materials and Methods* for details.

[†]"Directionally narrow" is measured as ratings that deviate from the midpoint.

were asked about their motivations for choosing specific search terms.[‡] A research assistant, blind to the study's hypotheses, rated the direction of each search term on a scale from 1 = "more related to health risks" to 7 = "more related to health benefits."

The direction of the search term generated by each participant was consistently positively correlated with that participant's prior beliefs about the health effects of their chosen food or beverage ($r = 0.34$; $P < 0.001$). Furthermore, only a minority of participants (8%) reported generating search terms to confirm their beliefs, and the effect persists ($r = 0.31$; $P < 0.001$) after excluding these people. This suggests that the narrow search effect occurs even when people are genuinely searching for information about a health topic of personal interest to them.

Studies 2a to 2f. Next, we investigated how differences in the search results from directionally narrow search terms impact people's subsequent postsearch beliefs, across the same six topics for which we documented that people generate directionally narrow search terms in Study 1. Participants were randomly assigned to one of two narrow search terms, reflecting opposite directions (e.g., either "Nuclear energy is good" or "Nuclear energy is bad" for the topic of nuclear energy). Participants then searched their assigned term on Google, posted a screenshot of their search results for confirmation, and rated their opinion on the topic they had searched.[§] Since participants were randomly assigned to search terms, we can assume similar average prior beliefs and therefore attribute any difference in postsearch beliefs across conditions to the difference in search results arising from the directionally narrow search terms used. Across all six studies (2a to 2f, total N = 1,658), participants' postsearch beliefs significantly differed, in the direction of the search term they had been randomly assigned to use (Cohen's $d_{caffeine} = 1.08$,[¶] $d_{\text{gas prices}} = 0.54$, $d_{\text{crime rates}} = 0.29$, $d_{\text{nuclear energy}} = 1.39$, $d_{\text{coronavirus}} = 0.50$, $d_{\text{bitcoin}} = 0.69$; all $ps < 0.05$).

In a posttest, 251 MTurk participants ($M_{age} = 36$, 50% female) were assigned to one of three caffeine search term conditions from Study 2a (both benefits and risks, benefits only, or risks only) and were shown what they presumed to be the search results for that term. However, in reality, all participants were shown the same Google search results from the broad search term (i.e., both benefits and risks) for all three conditions. We found no difference in beliefs ($P = 0.52$), confirming that our findings from Study 2 are attributable to the actual differences in search results, not to priming or demand effects from assigning the search term to participants.

Study 2g. In this study, participants (801 MTurk participants recruited, yielding 674 valid responses; Final $M_{\text{age}} = 39$, 58% female) were told that they were going to conduct an information search about the health effects of a food or beverage by using a beta-version search engine that worked similarly to Google. Then, participants were asked to name a food or beverage for which they were uncertain about the health effects, as in Study 1b. Afterward, participants were asked to generate a search term to learn about the health effects of this food or beverage, conduct an online search with their search term, write a summary of their findings, and rate their beliefs toward this food or beverage. Participants used the search engine interface we designed to conduct the information search. Unknown to the participants, this search engine interface called the Google API (Application Programming Interface) and displayed the top 10 Google search results for one of the following two randomly assigned narrow search terms: "[food/beverage that the participant named] health benefits" or "[food/beverage that the participant named] health risks." Finally, as in Study 1b, participants were asked about their motivations for choosing specific search terms.[#]

Participants who saw the "benefits" search results for the food/beverage they had selected were significantly more likely to believe that the food or beverage had higher health benefits than those who saw the "risks" search results ($M_{benefits\text{-}search} = 4.30$ vs. $M_{risks\text{-}search} = 3.62$, $t(672) = 4.51$, $P < 0.001$, $d = 1.97$). Participants perceived the search results as similarly useful and relevant in both conditions ($ps < 0.80$). As in Study 1b, few participants (11%) reported that they had written search queries to confirm what they already believed, and the effect persists after excluding these participants [$M_{benefits\text{-}search} = 4.34$ vs. $M_{risks\text{-}search} = 3.68$, $t(601) = 4.23$, $P < 0.001$, $d = 1.93$]. This suggests that the effect of narrow search results on beliefs occurs even when people are genuinely searching for information on a health topic of their choice.

AI tools that employ generative natural language to answer questions, such as ChatGPT, are currently transforming the way that information is being synthesized and provided to users (19). Information obtained from natural language AI may be easier to understand and potentially also more objective than from traditional search (20). As a result, AI language models have been touted as the future of search, and Microsoft has introduced a new Bing as the first AI-based search, with other AI-based search tools under development (21), raising the question of whether the use of AI-based search is also subject to confirmation bias.

Study 3. To test the narrow search effect in the context of AI-based search for information, we conducted Study 3, in which we used each of the four time-stable specific topics (from Study 2a-d), randomly assigning participants to see a ChatGPT 3.5 output (instead of Google search results) for one of the two directionally narrow search terms for that topic used in Study 2, resulting in eight between-subjects conditions.[‖]

Despite the fact that ChatGPT replies included an explicit acknowledgment of the opposing viewpoint,[**] the random assignment to directionally narrow AI queries yielded similar results as for traditional search. As in Study 2a-d, the postsearch beliefs of Prolific participants in Study 3 (N = 774, $M_{age} = 40$, 49% female) significantly differed, in line with the direction of the narrow ChatGPT query term ($d_{caffeine} = 0.1$ $d_{gas\ prices} = 0.53$, $d_{crime\ rates} = 0.53$, $d_{nuclear\ energy} = 0.50$; all $ps < 0.001$, see Fig. 2).

Our results show that people's prior beliefs result in generating directionally narrow search terms (Study 1) and that the differences in results generated by different directionally narrow search terms result in different postsearch beliefs (Studies 2 and 3), across a range of topics. In the next series of studies, we focus on the specific topic of the health consequences of caffeine. We use caffeine health impact as a topic about which people are likely to both have

[‡]In addition to search motivations, we collected data on several individual difference variables. The details of all measures and analyses are reported in *SI Appendix*.

[§]Please refer to *SI Appendix* for detailed methods and measures of all the studies reported.

[¶]Study 2a, the caffeine study, also included a third condition in which participants used a balanced search term, resulting in postsearch attitudes between the two directionally narrow search conditions. Participants assigned to use the benefits search term in their Google search held more positive post-search beliefs towards caffeine than those who used the broad search term [$M_{benefits} = 5.06$ vs. $M_{broad} = 4.43$, $t(179) = -3.02$, $P < 0.01$, $d = -0.45$]. Similarly, participants who were assigned to use the risks search term were less positive postsearch than those who used the broad search term [$M_{risks} = 3.51$ vs. $M_{broad} = 4.43$, $t(185) = 4.13$, $P < 0.001$, $d = 0.60$].

[#]In addition to search motivations, we collected data on several individual difference variables. The details of all measures and analyses are reported in *SI Appendix*.

[‖]The eight conditions are ChatGPT outputs of the following search terms: caffeine health benefits, caffeine health risks, gas prices will go up, gas prices will go down, crime rates will go up, crime rates will go down, nuclear energy is good, and nuclear energy is bad. See *Methods and Materials* for more details.

[**]For example, when asked about caffeine health benefits, ChatGPT also mentioned "However, it's important to note that while moderate caffeine intake can offer some health benefits, excessive intake can lead to adverse effects like insomnia, nervousness, restlessness, stomach upset, fast heartbeat, muscle tremors, and more..."

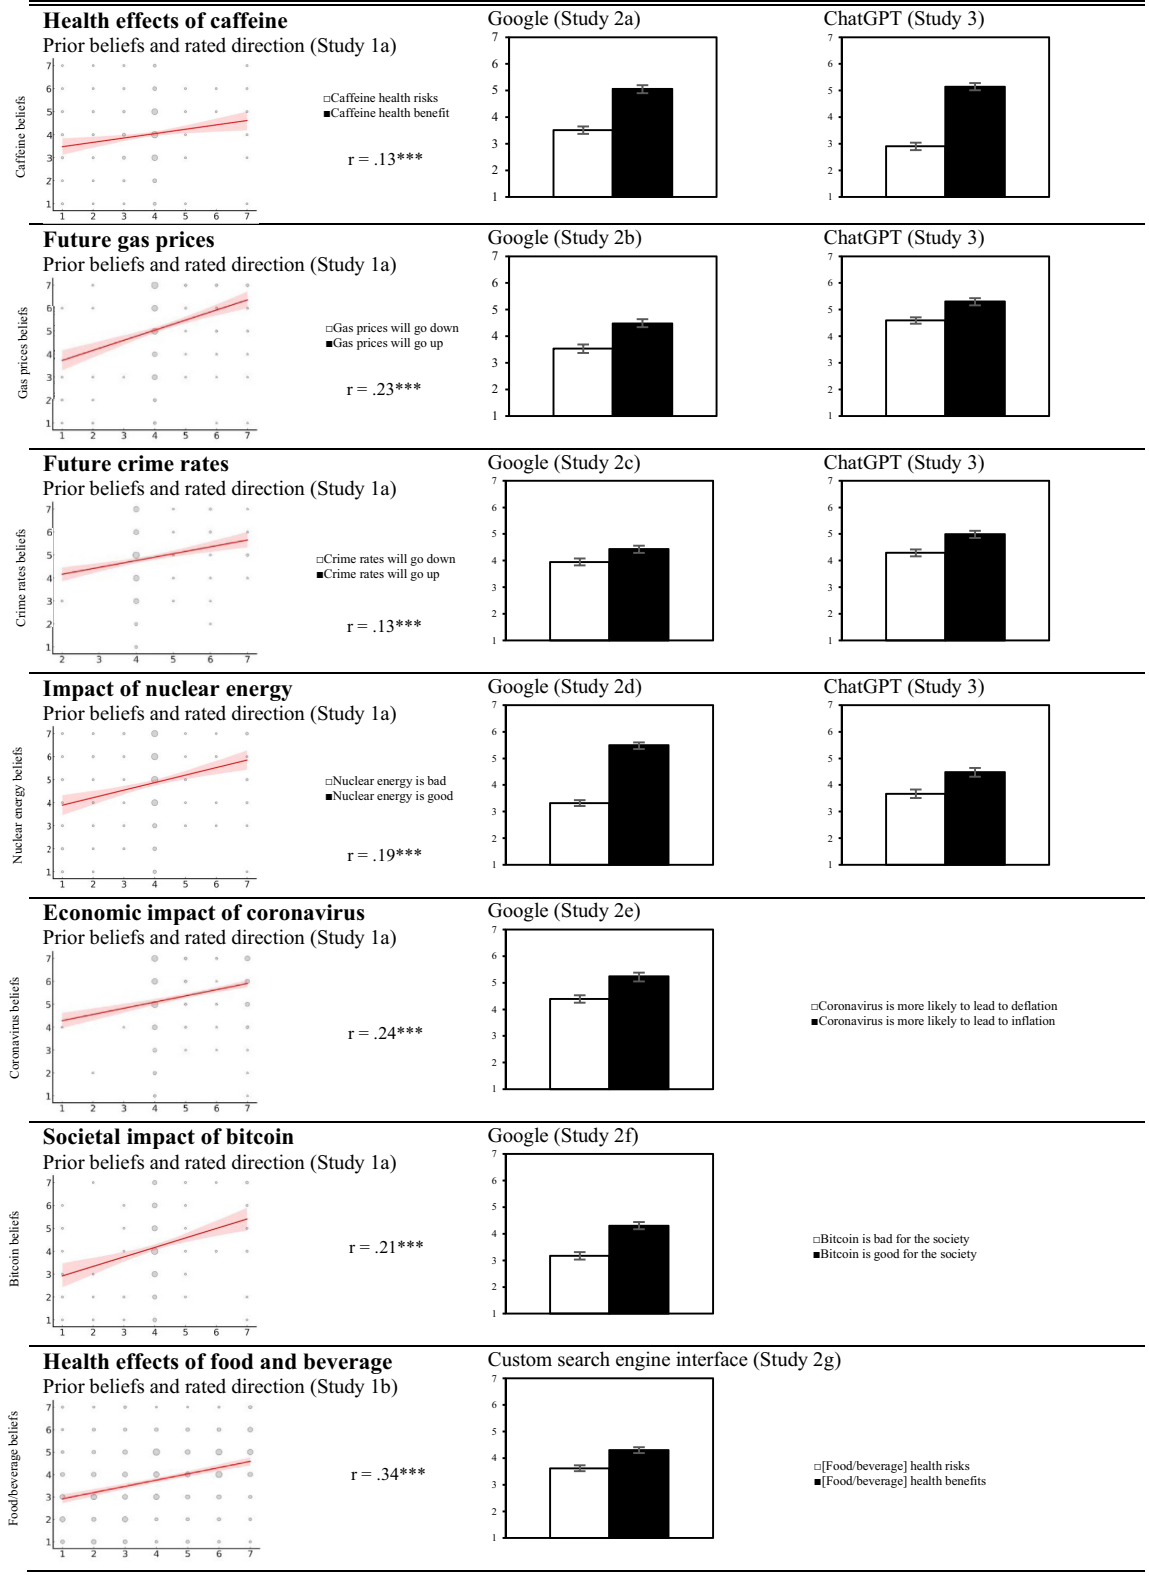

Fig. 2. Summary results of Studies 1a to 1b, 2a to 2g, and 3, ***$P < 0.001$.

malleable preconceived notions and also find either positive or negative caffeine hypotheses valid (e.g., when searching for information on specific medical conditions in a closed research database, 22). Furthermore, public search engines are unlikely to circumvent or adjust the algorithm for this topic, as they might do for more controversial subjects (e.g., US President Obama's place of birth) (23).

To test the relevance of our analysis for spontaneous real-world search, we extracted the most common search terms related to "caffeine" from Google Adwords planner with a threshold of a minimum monthly search volume of 1,000. We found that 26%[††]

[††]The search terms were evaluated by a research assistant blinded to the hypothesis. We conducted an additional study asking MTurk participants to evaluate the search terms and found similar results (*SI Appendix*).

of the 73 search terms were evaluated as directionally narrow by independent coders, consistent with Study 1. An additional content analysis revealed that the search results generated by the caffeine-related terms most commonly searched on Google were more directional, particularly in a negative direction, than the results from a broad search term, consistent with Study 2 (*SI Appendix*).

Study 4. We tested the full narrow search effect in the context of caffeine health information. We first asked 751 Prolific participants (M_{age} = 36, 52% female) to rate their initial beliefs regarding the effects of caffeine on health. All participants were then instructed to generate their own Google search terms to learn more about the health effects of caffeine. They were randomly assigned to two between-subjects conditions: the spontaneous-search condition, in which they used self-generated search terms, and the broad-search condition, in which they instead all used the same control search term "caffeine health benefits and risks." After conducting their Google search and uploading an image of the results, they again rated their beliefs regarding caffeine health effects. A separate group of 751 Prolific participants rated the direction of search terms produced by the participants.

Participants' prior belief ratings predicted the direction of the search terms they generated, as evaluated by independent coders (r = 0.22, P < 0.001). In the spontaneous-search condition, consistent with the narrow search effect, prior beliefs strongly predicted postsearch beliefs [b = 0.64, t(366) = 13.88, P < 0.001], mediated by the coder-rated direction of the search term (indirect b = 0.045, 95% CI = [0.019, 0.078]). In contrast, in the broad-search condition, prior beliefs more weakly predicted postsearch beliefs [b = 0.55, t(379) = 11.84, P < 0.001], and this relationship was not mediated by the direction of the (unused) search term they had generated. Imposing broad search therefore reduced the impact of prior beliefs on postsearch beliefs [regression analysis interaction b = −0.14, t(747) = −2.19, P = 0.029]. These results show that because people with different prior beliefs conduct different directionally narrow searches, they are therefore exposed to information more directionally consistent with their beliefs, and as a result they fail to update their beliefs in the way they would have if they had been exposed to broader search results.

Study 5. Next, we use a controlled lab setting to test whether the impact of directional narrowness can even impact consequential decisions based on postsearch beliefs. 346 English-speaking Dutch undergraduates (M_{age} = 21, 52% female) were randomly assigned to conduct a Google search using a search term about either caffeine health benefits or risks. They summarized their findings from the search, rated their beliefs about the effects of caffeine, and chose between a caffeinated or decaffeinated energy drink to take home for their own consumption. Participants who had been randomly assigned to search about benefits (vs. risks) of caffeine held more positive postsearch beliefs about caffeine [$M_{benefits}$ = 4.17 vs. M_{risks} = 2.92, t(344) = 8.32, P < 0.001] and were substantially more likely to choose the caffeinated drink (52% vs. 36%, χ^2 = 8.41, P <- 0.01). Postsearch beliefs mediated the effect of the assigned search term on energy drink choices (indirect b = −0.56, 95% CI = [−0.87, −0.32]).

Our findings thus far demonstrate that people make directionally narrow searches, in line with their prior beliefs and that making directionally narrow searches in turn impacts postsearch belief updating, to the degree that using a different health-related search term can result in different consumption decisions. Next, we explore what can be done to limit the impact of this tendency for directionally narrow search.

Study 6. First, we test whether simply conducting more searches mitigates the narrow search effect, as would be the case if people recognized the directional narrowness when viewing their search results and spontaneously broadened their search terms in subsequent search. In this study (130 MTurk participants, M_{age} = 34, 46% female), we replicated the ratings and self-generated caffeine health searches in the spontaneous search condition of Study 4. Half the participants were randomly assigned to a follow-up search group, were prompted to conduct a second search, and then rated their final beliefs. Overall, final beliefs were correlated with presearch beliefs, and the relationship was not weakened by having participants conduct an additional follow-up search (r = 0.73, P < 0.001), compared to only a single search (r = 0.58, P < 0.001; difference z = 1.39, P = 0.08). This suggests that the postsearch belief updating is not improved by conducting additional searches.

Study 7. Since people do not spontaneously correct for their narrow directional searches, we test whether prompting (but not requiring) individuals to consider how the results from a different search term might have differed would affect their belief updating. 431 MTurk participants (M_{age} = 35, 52% female) were assigned to one of two search terms (caffeine health benefits vs. caffeine health risks) and to one of two counterfactual consideration timing conditions (before search vs. postsearch) in a 2 × 2 between-subjects design. Prior to conducting the assigned search and making judgments, half the participants were asked to consider how their beliefs might have been different if they had used the opposite search term, whereas in the replication conditions, they were asked the same question after making their judgments.

When the prompt was given after they searched and rated their beliefs, we replicated the narrow search effect already shown, such that participants who used the benefits search terms had more positive beliefs [$M_{benefits}$ = 5.16 vs. M_{risks} = 3.49, t(218) = 8.53, P < 0.001]. This effect was significantly reduced [F(1, 427) = 8.13, P < 0.01] when the prompt was given before the search, but was not eliminated [$M_{benefits}$ = 4.76 vs. M_{risks} = 3.89, t(209) = 4.29, P < 0.001; *SI Appendix*, Fig. S1]. Nudging consideration of the directional narrowness prior to search was impactful but ultimately insufficient to eliminate the narrow search effect.

Study 8a. Next, we test the feasibility and effectiveness of using structural changes in search engines to reduce the narrow search effect. We developed a custom search engine platform to investigate interventions implemented at the algorithm level, specifically whether providing broader search results can mitigate the effects of narrow directional search on belief updating. 333 MTurk participants (M_{age} = 36, 59% female) were instructed to use the custom search engine (described in *SI Appendix*) to research caffeine's health effects and generated their own search term. The custom search engine, unknown to participants, displayed one of four randomly assigned sets of 10 Google search results: either for the participant-generated term, or for a directionally narrow term, either "caffeine health risks" or "caffeine health benefits," or for a broad term "caffeine health risks and benefits." Participants reviewed the results, wrote a summary of their findings, and rated their beliefs about caffeine. In addition, to test the possibility that broadened search would generate results that did not fit participants' goals for their search, they then evaluated the usefulness and relevance of the search results.

Postsearch beliefs about caffeine differed across the four types of search results displayed [F(3, 329) = 10.19, P < 0.001, see *SI Appendix*, Fig. S2], with the most positive attitudes among those shown health benefit results, followed by broad-search results in the

middle and the most negative attitudes were observed when shown health risk or their own search term results. These findings suggest that modifying search algorithms can result in differences in people's postsearch beliefs by modifying the composition of search results, holding the searched term constant. Importantly, participants did not find the balanced search results less useful or less relevant than those based on their own search term ($ps > 0.38$), indicating that algorithms may not need to deliver narrow results to meet users' needs and maintain perceived relevance.

Study 8b. Next, we again tested the effectiveness of using structural changes in search engines to reduce the narrow search effect, but in a different context, the relationship between age and thinking abilities. Similar to Study 8a, we developed a custom search engine to examine whether broader search results could mitigate the effects of a narrow search. A total of 1,002 participants (final valid responses = 770; M_{age} = 40, 61% female) first read a short prompt explaining that there has been recent debate about whether younger or older people make better leaders, with one of the key considerations being thinking abilities. This prompt was based on age-related criticism of then US President Biden at the time of the study, which ultimately lead to his withdrawal from the presidential race, but we did not specifically mention him, to avoid partisanship motives among participants.

Participants were then instructed to use this search engine to explore the relationship between thinking abilities and age by generating their own search terms. The custom search engine, unknown to participants, displayed one of two randomly assigned sets of 10 Google search results: either based on the participant-generated term (spontaneous condition) or a broad set combining search results reflecting both positive and negative perspectives on how thinking abilities change with age (broad condition). Participants reviewed the results, rated their beliefs about the relationship between thinking abilities and age, and evaluated the usefulness and relevance of the search results. Finally, participants were asked about their motivations for choosing specific search terms.[‡‡]

Participants who were shown balanced search results had more positive beliefs about the relationship between thinking abilities and age than those shown results based on their own search terms [$M_{spontaneous}$ = 3.28 vs. $M_{balanced}$ = 3.88, $t(768) = -5.83$, $P < 0.001$].

Only a minority of participants (5%) reported that they generated their search term to find evidence confirming their belief, and the results remained significant when excluding these participants [$M_{spontaneous}$ = 3.26 vs. $M_{balanced}$ = 3.91, $t(739) = -6.19$, $P < 0.001$]. This suggests that the narrow search effect occurs even when people are genuinely searching to find information.

Similar to Study 8a, there were no significant differences in participants' ratings of the usefulness ($P = 0.38$) or relevance ($P = 0.29$) of the search results between the two conditions. Overall, the results of Studies 8a and 8b demonstrate across two different search topics that changing the search algorithm to show either positive or negative information, or to show narrow or broad search results, impacts beliefs on health and policy-related topics.

Study 9a. Finally, we contrast the results from current relevance-maximizing search algorithms with an alternative algorithmic strategy, presenting a combination of narrow and broader search results. Similar to Study 8, we used our custom-designed search engine interface (leveraging the Google API) for this study. Specifically, 193 MTurk participants (M_{age} = 38, 40% female) were informed that they would be using a beta version of a search engine, similar to Google, to research the effects of caffeine. They were asked to generate a search term, conduct an online search, summarize their findings, and rate their beliefs about caffeine. In the control condition, the search engine displayed the top 10 Google results for the participant-generated search term. However, in the broadened-results condition, the search engine instead displayed a list of 10 results, alternating between the top Google results for the participant-generated search term and for "caffeine health risks and benefits." Participants also rated the search results' usefulness and relevance.

Participants in the control condition, who saw narrower search results, held less positive postsearch beliefs toward caffeine than those in the broadened-results condition [$M_{spontaneous}$ = 4.22 vs. M_{mixed} = 4.72, $t(191) = -2.11$, $P = 0.036$]. Participants perceived the search results as similarly useful and relevant in both conditions ($ps < 0.57$).

Study 9b. This study extended the results of Study 9a to a new context and a new technological tool. We used ChatGPT and a custom interface to simulate two short-answer AI chatbots, one that provides narrow answers to queries and one that provides broad answers. We tested how the use of one vs. the other chatbot impacts participants' beliefs regarding the influence of age on mental abilities.

A total of 793 participants (M_{age} = 38, 59% female) were initially instructed to generate a query aimed at learning about the impact of age on mental abilities. They then interacted with a custom-designed AI chatbot interface that, without their awareness, was randomly assigned to return responses to their query from either the "narrow response" version or from the "broad response" version of the chatbot.

In the narrow condition, the AI chatbot displayed a ChatGPT response to their query with a prompt to provide the most relevant and accurate answer, structured in bullet points and limited to approximately 250 words. In the broad condition, the chatbot instead displayed a ChatGPT response to the query with a prompt to provide a balanced viewpoint, covering pros and cons, multiple perspectives, and additional details, also within a 250-word limit. It is important to note that both conditions provided answers based on participants' specific searches—the responses varied only in their breadth.

After receiving the AI-generated response, participants were asked to rate their beliefs about the impact of mental abilities and age. Additionally, they assessed the usefulness and relevance of the AI chatbot's answers. Finally, participants were asked about their motivations for choosing specific search terms.[§§]

Participants in the broad condition reported more positive beliefs about how mental abilities change with age compared to those in the narrow condition [M_{broad} = 3.98 vs. M_{narrow} = 3.40, $t(791) = 5.46$, $P < 0.001$]. Only a minority of participants (7.1%) reported that they generated their search term to find evidence confirming their belief, and the results remained significant when excluding these participants [M_{broad} = 3.97 vs. M_{narrow} = 3.41, $t(735) = 5.12$, $P < 0.001$].

Participants rated the chatbot's responses similarly in both conditions in terms of usefulness and relevance ($ps > 0.20$). The results highlight the potential for an AI-generated informational chatbot designed for broad responses to provide a more comprehensive understanding, fostering belief updating without diminishing the perceived quality of information.

The results of Studies 9a and 9b indicate that informational technology can be configured to adjust for the human tendency for

[‡‡]In addition to search motivations, we collected data on several individual difference variables. The details of all measures and analyses are reported in *SI Appendix*.

[§§]In addition to search motivations, we collected data on several individual difference variables. The details of all measures and analyses are reported in *SI Appendix*.

confirmation bias, by broadening the information provided beyond the most relevant narrow-search results. This "broadened search" approach could foster more comprehensive exposure to information, promote belief updating, and counter the narrow search effect, without necessarily sacrificing perceived relevance or usability.

Discussion

This research documents the profound impact that confirmation bias can have on how people search for information, both on their search results and on their subsequent belief updating and decisions. Across 21 studies, we observed that prior beliefs tend to persist after spontaneous search, due to people's tendency to search narrowly, generating narrow search results that are in line with their prior beliefs. However, when the information technology is designed to instead provide them with broader information, people update their beliefs more after search. While nudging people to consider the directional narrowness of their search results may help to some degree, it is also essential that the development of search algorithms takes into account how the tendency for people to generate belief-confirming search terms can undermine an algorithm's ability to provide broader information. Our results suggest that structural changes which broaden search and AI algorithms can mitigate confirmation bias, promote belief updating, and potentially foster more broadly informed decision-making.

Our analysis has focused on one type of confirmation bias—the tendency to formulate questions that solicit affirming responses. Prior research has also documented another aspect of confirmation bias, a tendency to selectively incorporate evidence that aligns with preexisting beliefs (16, 18). This would suggest that even when people are shown broader search results, they may discount the information that challenges their own views. Our research suggests that, at least for the kinds of topics we have studied here, confirmation bias in evidence incorporation does not overrule the belief-updating benefits of broader search results. While these effects may be more limited for more politically polarized topics (due to motivated reasoning bolstering evidence-incorporation confirmation bias, e.g., ref. 24), our findings suggest that targeting specifically formulation-based confirmation bias by broadening search can effectively promote belief updating.

We did not find any evidence that broadening search conflicted with users' goals or reduced the perceived relevance of the results. That said, in general, the benefits of broadening search will depend on the informativeness and validity of the information excluded by directionally narrow search. When people are searching for specific factual information (e.g., the opening hours for visiting the Eiffel Tower), broader search may not be beneficial and could even be more confusing if it leads to displaying irrelevant information (e.g., the height of Eiffel Tower). Even more problematic, the prevalence of online misinformation on a topic (e.g., where US President Barack Obama was born) can result in mixed effects of broadening search, increasing exposure to misinformation for some users, while diluting the misinformation shown to others, depending on their search terms. Notably, these are fundamental problems for search in general, whether or not search is broadened, and search engines have at times reportedly been engineered to avoid solely relying on page-rank algorithms to address these issues (25, 26).

In sum, the narrow search effect (and the benefit of broadening search results) is most likely to occur under the specific conditions implied by our theoretical framework—i) when users hold biased beliefs that shape the search terms they use, ii) when information-provision technology yields different results depending on the directionality of the request, and iii) when users' beliefs are sufficiently malleable to update based on the information received. These findings may not generalize to contexts where shared, prominent cues determine spontaneous searches regardless of people's beliefs, information is sufficiently scarce or well-integrated such that balanced results are provided regardless of the query, or when people hold strong views that they are resistant to updating.

Nevertheless, our results suggest that the design of information-provision tools should take into account peoples' tendency for directionally narrow search and the potential for an algorithmic focus on relevance to narrow search results, hindering belief updating. Recently, Microsoft (27) introduced a new Bing, an AI-assisted search engine featuring the Prometheus OpenAI language model, for improved search relevance, and natural language interaction capabilities. Users can pose complex questions conversationally, and AI-assisted Bing offers summarized responses, based on AI-synthesis of a range of web sources. Notably, AI-assisted Bing seems to also involve a separate "prompt engineering" phase, generating the prompts instead of literally repeating the user's query. For example, we found that AI-assisted Bing reformulated the narrow queries "nuclear energy is good" or "gas prices will go up" to the broader versions "nuclear energy pros and cons" and "gas prices prediction", respectively. However, AI-assisted Bing does not seem to consistently apply this across queries. As a result, we replicated the narrow search effect regarding health effects of caffeine using AI-assisted Bing (*SI Appendix*).

Additionally, Study 9b demonstrates the potential to leverage AI capabilities to develop a next generation of information-provision tools that promote belief updating by offering broader, more balanced responses. Adopting a hybrid approach that leverages both narrow and broad responses could enable AI systems to dynamically adjust to users' informational needs, promoting both belief updating and informed decision-making. Notably, implementing a broad-answer design in AI chatbots via prompt engineering may be more feasible at scale than updating the search algorithms used in a traditional search engine to achieve a similar effect. By incorporating diverse viewpoints within a single AI response, such systems could provide users with the specific relevant information requested alongside information from perspectives that the user may have omitted, promoting belief updating without necessarily reducing perceived relevance or usefulness. Our findings provide initial proof-of-concept and highlight the need for research that more fully tests psychologically informed prompt-engineering approaches, an emerging question that bridges the psychology of decision-making and human–computer interaction.

A large research literature has found that defaults consistently impact behavior (28), and our findings suggest that optimizing relevance in search algorithms has had the effect of defaulting people to narrow search, thereby impeding belief updating. We suggest that this default should be questioned and that there are likely to be benefits to instead explicitly providing opportunities for people to receive broader search results. In fact, in an additional study (N = 101, M_{age} = 35, 49% female), we found that the majority (84%) of participants indicated interest in using a "Search Broadly" button (e.g., the opposite of the "I'm Feeling Lucky" button on the Google homepage that loads the top search result) if it were provided by search engines as a way to receive broader results.[¶¶]

As AI evolves and information provision becomes only more algorithmically mediated, it is crucial to recognize and mitigate the risks of the narrow search effect. The wise words of Victor

[¶¶]Note that although AI-assisted Bing offers a "more balanced" option, this refers to a balance in conversation style (i.e., balancing between being friendly and informative) rather than broadening of content.

Hugo, "Wide horizons lead the soul to broad ideas; circumscribed horizons engender narrow ideas," capture an essential challenge embodied in both existing and emerging information-provision technologies. When information-provision technology not only focuses on relevance but also broadens horizons, individuals will access more thorough and broader information, fostering more broadly informed beliefs and decisions. By creating an environment with more shared factual understanding, broader search algorithms can play a significant role in mitigating belief polarization, thereby potentially contributing to a more cohesive society.

Materials and Methods

All study materials, data, preregistrations, deviations from preregistration, and analysis code can also be found online: https://osf.io/86xeb/?view_only=b610eea785494d0190ee46d3be2ddda4. Analytic and methodological details for all studies are included in *SI Appendix*. The findings reported in this manuscript were significant even when we did not exclude data from participants (see *SI Appendix* for ancillary analyses). This research was approved by the Chicago IRB (Protocol ID: IRB24-0383) and the Tulane University Social-Behavioral IRB (Protocol ID: 2023-1395).

Exclusion Criteria.

Studies. In all studies, we excluded records with duplicate IP addresses and failed attention checks prior to analysis. In studies where participants were asked to come up with a search term, we excluded participants who came up with an irrelevant search term. In the studies where we used our updated search engine interfaces (Studies 2g, 8b), we were able to collect participants' IDs as well as the search terms they entered into our custom generative AI or search engine. Thus, in those studies, we also excluded participants who did not use the search engine as instructed (e.g., participants who did not enter something close to their entered search term in the search engine) or those whose search terms were not consistent between what they entered in Qualtrics and what they entered into the custom search engine, as rated by an independent coder. The table below lists the number of participants who completed the study and the final number of participants after excluding the participants mentioned above.

	Number of participants recruited	Number of participants included
Study 1a to 1b	1,603	1,481
Study 2a to 2g	2,524	2,332
Study 3	801	774
Study 4	803	751
Study 5	354	346
Study 6	131	130
Study 7	454	431
Study 8a to 8b	1,420	1,103
Study 9a to 9b	1,266	986
Study presented in the discussion (New Bing)	301	296
Study presented in the discussion (Search Broadly)	101	101
Supplementary study	148	146

Posttests. Prior to conducting posttests, we excluded search terms that do not contain any caffeine-related words (i.e., caffeine, coffee, tea, soda). We also excluded search terms in a paragraph format copied from other websites (e.g., we excluded search terms like: "The Mayo Clinic state that consuming more than 500 to 600 mg of caffeine a day may lead to insomnia, nervousness, restlessness, irritability, an upset stomach, a fast heartbeat, and even muscle tremors. However, previous research has linked even moderate amounts of caffeine to negative health effects…") because the participants generating the search terms had not complied with the instructions.

Preregistration. The following table contains a list of the 14 preregistered studies along with their respective preregistration numbers.

Study	Domain	Search engine	Aspredicted number
1a	Caffeine, Gas prices, Crime rates, Nuclear energy, Coronavirus, Bitcoin		112159
1b	Food and beverage		189224
2b	Gas prices	Google	111448
2c	Crime rates	Google	111449
2d	Nuclear energy	Google	111450
2f	Bitcoin	Google	111451
2g	Food and beverage	Custom search engine interface	189471
3	Caffeine, Gas prices, Crime rates, Nuclear energy	Chat GPT-3.5	144553
4	Caffeine	Google	104627
5	Caffeine	Google	8197
8a	Caffeine	Custom search engine interface	12098
8b	Thinking abilities and age	Custom search engine interface	187347
9b	Thinking abilities and age	Custom AI chatbot interface	198558
Study presented in the discussion	Caffeine	Bing powered by GPT-4	143527

We had initially preregistered Study 9a (AsPredicted #39777), but due to technical issues with the link to the Google API, the platform stopped pulling Google API search results during data collection. As a result, we could not secure the number of participants stated on the preregistration form, so we no longer consider this study as preregistered.

Data, Materials, and Software Availability. Data have been deposited in the Open Science Framework (29).

ACKNOWLEDGMENTS. We thank Anyi Ma, Gabriele Paolacci, and Stefano Puntoni for comments and Yuhan Chen and Chih-Yu Chiang for technical support for the custom search engine. We would also like to thank Anastasiya Apalkova, Jieyi Chen, Nicholas Herzog, and Min Thapa for their assistance with data collection on Mturk or Prolific.

1. A. Alesina, A. Miano, S. Stantcheva, The polarization of reality. *AEA Papers Proc.* **110**, 324–328 (2020).
2. M. G. Findling, R. J. Blendon, J. M. Benson, Polarized public opinion about public health during the COVID-19 pandemic: Political divides and future implications. *JAMA Health Forum* **3**, e220016 (2022).
3. E. K. Smith, M. J. Bognar, A. P. Mayer, Polarisation of climate and environmental attitudes in the United States, 1973–2022. *npj Clim. Action* **3**, 1–14 (2024).
4. D. F. Gleich, PageRank beyond the web. *SIAM Rev.* **57**, 321–363 (2015).
5. Google, How results are automatically generated (2023). https://www.google.com/search/howsearchworks/how-search-works/ranking-results/. Accessed 5 April 2024.
6. N. Langville, C. D. Meyer, *Google's PageRank and Beyond: The Science of Search Engine Rankings* (Princeton University Press, 2011).
7. L. Page, S. Brin, R. Motwani, T. Winograd, The PageRank citation ranking: Bringing order to the web (Stanford infolab, 1999). http://ilpubs.stanford.edu:8090/422/1/1999-66.pdf. Accessed 9 March 2025.
8. S. Altay, M. Berriche, A. Acerbi, Misinformation on misinformation: Conceptual and methodological challenges. *Social Media + Society* **9**, 20563051221150412 (2023).
9. E. Pariser, *The Filter Bubble: What the Internet Is Hiding from You* (Penguin Books Limited, 2011).
10. K. H. Jamieson, J. N. Cappella, *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment* (Oxford University Press, 2008).
11. R. Epstein, R. E. Robertson, The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proc. Natl. Acad. Sci. U.S.A.* **112**, E4512–E4521 (2015).
12. C. R. Mynatt, M. E. Doherty, R. D. Tweney, Confirmation bias in a simulated research environment: An experimental study of scientific inference. *Q. J. Exp. Psychol.* **29**, 85–95 (1977).
13. H. Shaklee, B. Fischhoff, Strategies of information search in causal analysis. *Mem. Cognit.* **10**, 520–530 (1982).
14. M. Snyder, W. B. Swann, Hypothesis-testing processes in social interaction. *J. Pers. Soc. Psychol.* **36**, 1202–1212 (1978).
15. M. Zuckerman, C. R. Knee, H. S. Hodgins, K. Miyake, Hypothesis confirmation: The joint effect of positive test strategy and acquiescence response set. *J. Pers. Soc. Psychol.* **68**, 52–60 (1995).
16. J. Klayman, "Varieties of confirmation bias" in *Psychology of Learning and Motivation*, J. Busemeyer, R. Hastie, D. L. Medin, Eds. (Academic Press, 1995), pp. 385–418.
17. D. Kuhn, J. Lao, Effects of evidence on attitudes: Is polarization the norm? *Psychol. Sci.* **7**, 115–120 (1996).
18. C. G. Lord, L. Ross, M. R. Lepper, Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *J. Pers. Soc. Psychol.* **37**, 2098–2109 (1979).
19. N. Grant, C. Metz, A new chat bot is a 'code red' for Google's search business. *The New York Times*, 21 December 2022. https://www.nytimes.com/2022/12/21/technology/ai-chatgpt-google-search.html. Accessed 9 March 2025.
20. V. Hristidis, N. Ruggiano, E. L. Brown, S. R. R. Ganta, S. Stewart, ChatGPT vs Google for queries related to dementia and other cognitive decline: Comparison of results. *J. Med. Internet Res.* **25**, e48966 (2023).
21. N. Grant, Google devising radical search changes to beat back A.I. Rivals. *The New York Times*, 16 April 2023. https://www.nytimes.com/2023/04/16/technology/google-search-engine-ai.html. Accessed 9 March 2025.
22. V. O. Kayhan, Seeking health information on the web: Positive hypothesis testing. *Int. J. Med. Inform.* **82**, 268–275 (2013).
23. V. N. Pham, Our foreign president Barack Obama: The racial logics of birther discourses. *J. Int. Intercult. Commun.* **8**, 86–107 (2015).
24. S. Su, Updating politicized beliefs: How motivated reasoning contributes to polarization. *J. Behav. Exp. Econ.* **96**, 101799 (2022).
25. J. Purtill, Google's hidden search algorithms are being investigated by researchers. Here's what they've found. *Australian Broadcasting Corporation News*, 4 August 2021. https://www.abc.net.au/news/science/2021-08-05/google-curating-covid-search-results-algorithm-project-finds/100343284. Accessed 9 March 2025.
26. K. Grind, S. Schechner, R. McMillan, J. West, How Google interferes with its search algorithms and changes your results. *Wall Street Journal*, 15 November 2019. https://www.wsj.com/articles/how-google-interferes-with-its-search-algorithms-and-changes-your-results-11573823753. Accessed 9 March 2025.
27. Microsoft, Introducing the new Bing (2023). https://www.microsoft.com/en-us/bing. Accessed 30 April 2024.
28. J. M. Jachimowicz, S. Duncan, E. U. Weber, E. J. Johnson, When and why defaults influence decisions: A meta-analysis of default effects. *Behav. Public Policy* **3**, 159–186 (2019).
29. E. Leung, O. Urminsky, Data and Materials for The Narrow Search Effect. Open Science Framework. https://osf.io/86xeb/?view_only=b610eea785494d0190ee46d3be2ddda4. Accessed 9 March 2025.