How Do User Opinions Influence Their Interaction With Web Search Results?

Luyan Xu
Renmin University of China
Beijing, China
xuluyan@ruc.edu.cn

Mengdie Zhuang
University College London
London, UK
m.zhuang@ucl.ac.uk

Ujwal Gadiraju
Delft University of Technology
Delft, The Netherlands
u.k.gadiraju@tudelft.nl

ABSTRACT

Understanding the influence of users' opinions on their search behavior together with their inherent biases in web search has garnered widespread interest in recent times. This is largely due to the implications of promoting critical thinking, explaining phenomena such as *political polarization*, or the manifestation of *echo chambers*. It is important to understand how personal opinions can bias users' interaction with search results. Moreover, there is a lack of understanding of the impact of user search intents, namely non-purposeful browsing versus searching with a pre-defined goal, on users' interactions with search results. We take a step towards bridging this knowledge gap through an empirical study in this paper. To do so, we select two controversial topics in abortion and gun control, and invite users to learn about them through *'Purposeless'* and *'Purposeful'* web searching. Our findings suggest that users with strong personal opinions exhibit biased interactions with the search results. However, the effect of users' opinions on their interactions with search results can differ depending on whether users search purposelessly or with a purpose. Our findings advance the current understanding of the effect of users' opinions in web search sessions, and show that users' search intents shape their interaction with search results. This work has broad design implications on dealing with bias in interactive information retrieval systems.

CCS CONCEPTS

• **Human-centered computing** → *User studies*; **HCI design and evaluation methods**.

KEYWORDS

Web Search, Cognitive Bias, SERPs, User Behavior

ACM Reference Format:

Luyan Xu, Mengdie Zhuang, and Ujwal Gadiraju. 2021. How Do User Opinions Influence Their Interaction With Web Search Results?. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21), June 21–25, 2021, Utrecht, Netherlands.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3450613.3456824

UMAP '21, June 21–25, 2021, Utrecht, Netherlands

ACM ISBN 978-1-4503-8366-0/21/06…$15.00
https://doi.org/10.1145/3450613.3456824

1 INTRODUCTION

The inherent tendency for people to prefer one thing over another, or have diverse and conflicting opinions on practically any given topic can be reflected in nearly every aspect of our lives [3]. This has been embedded, spread and retrieved through different forms of information dissemination (e.g. books, presentations, online comments, in-person). With the rise in digital data and the ubiquity of web search, information can now spread and converge much faster, thereby influencing many more people than ever before.

In the context of information retrieval using search engines, where searchers either seek or are presented with controversal information, confirmation bias can be observed wherein users search for information to confirm their own hypothesis/beliefs [1, 18]. Many methods have been developed to understand users' opinions towards a topic during web search and help them cope with the entailing bias. Prior works have also established that the results returned by search engines can significantly influence users' opinion [3, 21]. Besides, it has been found that the correctness of search results can change depending on the formulation of users' search queries [1]. Others have argued and shown that a remedy for bias starts with creating an awareness of its existence [3, 11].

Previous works have studied the dynamics and effect of users' opinions and bias in a complete topical search process, containing multiple query reformulations and the corresponding post-query browsing sessions. However, little is known about how the influence of users' opinions varies with respect to a users' search intent or the lack thereof - e.g. an exploratory search session without a well-defined search intent versus one with a well-defined search intent.

To address this gap, we investigate whether users' opinions changed through interacting with the search results, and how the users' opinions affect their behaviors when they interact with the search results. We recruited 200 distinct users from a crowdsourcing platform, who interacted with web search results in two scenarios – purposeless browsing and browsing with a clear intent, on two controversial topics, 'Abortion' and 'Gun control'. We found that users' opinions could significantly affect their interaction with search results. However, this varied based on their search was purposeful or purposeless. Our findings enhance and enrich the existing understanding of users' opinions through post-query browsing sessions. For the benefit of the community, we publicly release a dataset containing all the statements we used in this work and the gathered user behavior logs (anonymized) [1].

[1]https://osf.io/5k3wx/?view_only=c2e5cae8ea9a43ecbd09eae811957dcd

2 RELATED LITERATURE

2.1 Confirmation Bias in Web Search

Studying people's opinion and bias in web search has always attracted attention due to its implications in encouraging users to consume trustworthy information [22], promoting critical thinking [23], providing a socio-psychological explanation for certain phenomena (e.g., political polarization [4], and the development of echo chambers [15]).

Confirmation bias [17] is a type of cognitive bias which characterizes a person's inability to remain impartial in the act of acquiring new information despite their preconceptions, prior beliefs and hypotheses. Such a person would favor the acquisition of information that *confirms* (hence the nomenclature) their beliefs. We concentrate on the topic of confirmation bias in the context of search, and in particular, online search, where prior research has broken down the psychological patterns that dictate this type of interaction. Here, as in [10], we face the dichotomy between *selective exposure* and *selective avoidance*, as the attractive - repulsive impulses that govern the biased acquisition of information. The heuristic interpretation of confirmation bias under this paradigm is that people favor opinion-reinforcing information or avoid opinion-discouraging information, in an effort to reduce *cognitive dissonance* [5]. Web search and interactive information retrieval systems can be subject to many types of user and cognitive biases (see [2, 12, 13]), but these fall beyond the scope of this study.

2.2 Search Behaviors and Outcomes

Past studies have investigated how selective exposure influences an individual's search behavior as well as its outcomes [15, 22]. Methods have also been developed to mitigate the echo chamber effect. Most studies (e.g. health information search [16, 22], controversial topics in online forum [15]) confirmed the existence of such an effect with the participants favouring opinion-supporting information, and reinforcing their initial opinions.

The methods can be grouped into two types: algorithm based and interface based. Algorithm based methods, such as recommendation systems [8], aim to decrease polarization. Interface based methods focus on encouraging users to access an unbiased spread of opinions or increasing the exposure of opposing opinions in relation to opinionated content. [15] introduces a position marker for information sources and observed that, for individuals who are highly motivated to accurately learn about the topic, the exposure to such a marker increases their tendency to investigate opposing opinions more closely. Similar indicators have been applied in medical information search [16]. [9] mixed these two methods by visualizing the recommendation of politically diverse profiles on social media platforms - Twitter. These studies confirmed that users who are relatively more interested in the topic are more likely to exhibit an unbiased exploration of all the opinions.

3 STUDY DESIGN

We recruited 200 distinct participants from Figure8[2](a premier crowdsourcing platform), to systematically study whether and how users' original opinions evolve through interacting with the search results in a post-query web search session. Participation was restricted to workers from USA. We explored the impact of users' perceptions of opinions on their interactions with the search results in post-query browsing sessions. In the following sub-sections, we introduce the framework of our experiment.

3.1 Search Topics

We chose two controversial and widely discussed topics from US politics (Abortion, and Gun Control) from Wikipedia's List of controversial issues[3]. We chose these popular and controversial topics to ensure that our participants hailing from the US would arguably have some basic understanding of the topic, and potentially an existing opinion on these topics. Note that these topics were also selected to study bias mitigation in an NLP task by prior work [11]. For each of the chosen topics we selected a main statement that reflects the central pro/contra aspect of the controversy, e.g. 'Abortion should be legal', and 'Abortion should not be legal'. We then manually extracted arguments from an online debate forum using two criteria.[4] First, the length of the argument is less than 180 words. Second, a single clear argument is made, rather than mixing multiple arguments together (all statements are publicly available).[5]

Four expert annotators were asked to validate all statements in the final set to ensure that they contained explicit bias. A statement is considered as biased if it is disproportionate in favor of or against the neutral statement without presenting facts. Where necessary, we modified the statements briefly to make them comprehensible out of context and grammatically correct. We split the resulting set of biased statements into *pro* statements that support the main statement for this topic and *contra* statements that oppose the main statement. Additionally, we extracted *neutral* statements from the associated Wikipedia articles that contain facts and statistics pertaining to the topics. We followed the process of open coding to ensure that the statements were reliably identified as pro, contra and neutral [19]. The expert annotators iteratively coded the resulting statements as either 'pro', 'contra', or 'neutral' until unanimous agreement was reached on each statement, thereby forming the ground truth for our experimental tasks. Examples of 'neutral', 'pro', and 'contra' statements are presented in Table 1.

3.2 Measuring Users' Attitudes and Interactions

Before starting the search task, users were asked to complete a questionnaire which included questions pertaining to their demographics and their attitude towards the given topic. User self-reported attitudes regarding the topic at hand were gathered before and after the search task using counterbalanced statements at random (e.g., 'Abortion should be legal', 'Abortion should not be legal') with a 5-point Likert scale from *1:strongly disagree* to *5:strongly agree*. A post-task questionnaire was used to collect their attitudes pertaining to the topic immediately after completing each task, and also to gauge whether the users believed the search results were

[2]Figure8: https://appen.com/figure-eight-is-now-appen/

[3]https://en.wikipedia.org/wiki/Wikipedia:List_of_controversial_issues

[4]https://www.debate.org

[5]https://osf.io/5k3wx/?view_only=c2e5cae8ea9a43ecbd09eae811957dcd–'statements.xlsx'

Table 1: Examples of 'neutral', 'pro', and 'contra' statements.

Topic	Neutral	Pro	Contra
Abortion	'26 countries ban abortion altogether'	'Abortion prevents the teen parent issue'	'Abortion is murder'
Gun Control	'Gun laws per US state'	'Gun control helps to identify criminals'	'Guns don't kill people, people kill people'

either 'neutral' or inclined towards the 'pro' or 'contra' side of the statement.

User behavior data (spanning all interactions during the study) was logged using Javascript and the JQuery library, including activity data ranging from mouse movements to keypresses. We took additional measures (e.g. browser fingerprinting) to prevent workers from participating in the study multiple times[6]. We stored only the final hashed fingerprints, to avoid privacy intrusion of workers. Workers were also given an opportunity to opt-out of the Javascript tracking. In this way, we gathered worker activity data and computed two sets of features: number of *clicks* on 'pro' / 'neutral'/'contra' articles during one task, and the total *browsing time* of 'pro' / 'neutral'/'contra' articles during a task.

3.3 Search System

Making use of the statements introduced in Section 3.1, we developed a pseudo searching platform called *Opinions* (as shown in Figure 1 - (a)). To be more specific, we extracted the first two sentences of each statement as the snippets listing in its SERPs (9 snippets per SERP). When users click on a snippet, the corresponding statement would appear as articles in a new tab in *Opinions*. *Opinions* retrieved the articles using Solr[7]. To avoid the effect of the potential bias in users' queries on the ranking of returned search results, the queries were automatically set as either 'Abortion' or 'Gun Control' for the corresponding experiments, and Opinions returned the search results with a balance in stances (i.e. 3 'pro', 3 'contra', 3 'neutral'), with a randomized ranking. We deployed *Opinions* as the search engine that participants could rely on to gather information during the experimental study.

3.4 Study Procedure

We carried out a 2×2 between subjects experiment with 200 participants in total; we divided them into the '*Purposeless*' group wherein the participants were instructed to "browse web search results pertaining to a controversial topic", and the '*Purposeful*' group wherein the participants were informed to "collect supporting information pertaining to a controversial topic". The resulting 4 experimental conditions varied with respect to topics (*Abortion*, *Gun Control*) and the search intent (*Purposeless*, *Purposeful*) in our study, and 50 participants were recurited for each. The overall workflow for participants in our experimental setup is illustrated in Figure 1.

Being directed to our external experimental enviroment, participants were first asked to respond to a few general background questions with regard to their age, gender, education and ethinicity. Next, participants received the pre-study questionnaire (i.e. section 3.2) to gather self-reports of their original opinions corresponding to a controversial topic (i.e. 'Abortion' or 'Gun Control'). Following this, participants were told to use the *Opinions* platform to search and browse as much information as possible about the given topic. To make sure that participants did interact with the search system to a certain extent, we empirically set 5 minutes as the minumum time that users need to spend on using *Opinions* ; participants were also encouraged to proceed to the next stage only once they felt that they had collected enough information to form a clear view about the given topic.

We subsequently logged all the activities of the workers (mouse movements, key presses, clicks, dwelling time, etc.) within the *Opinions* platform. On completing the informational task, we administered the post-study questionnaire as described earlier. Participants who finished the task received a unique code that they could enter on Prolifc to claim their reward. All participants were paid at an hourly rate of 7.5 USD.

3.5 Data Collection

To ensure reliability of the resulting data, we restricted the participation to Level-3 workers [8] from English-spearking countries on Figure8. 32% of users were female. 17% of users were aged 18 - 25; 41% of users were aged 26 - 35; 23% of users were aged 36 - 45; 15% of users were aged 46 - 55; 4% of users were aged 56 - 65. Of all the users, 61.5% received college or higher education while the rest did not. All users in our experiment reported that they used search engines frequently and several times per day.

Enforcing reputation restrictions is a standard method adopted by requesters to ensure reliability [14]. We followed the guidelines laid down by prior work [7] and used attention check questions to label unreliable participants [6] in this study. We examined the responses of partipants to flag those with an overall accuracy of 0% as being untrustworthy. We further filtered out those participants who had no interactions with the *Opinions* platform. In total, we filtered out 64 participants due to the aforementioned criteria, resulting in 136 participants across the two topics. We henceforth refer to these filtered participants as users in our experimentally orchestrated search sessions.

4 RESULTS

4.1 Bias Dynamics

Previous works has showed that users' opinions can change through searching [20]. In this section, we further discuss whether the change of users' opinions happens suddenly in a single post-query browsing session. We calculate the difference between users' opinion ratings in the post- and pre-study questionnaires to measure the change of users' opinions.

[6]http://github.com/Valve/fingerprintjs

[7]Solr: https://lucene.apache.org/solr/

[8]Level-3 contributors on Figure8 comprise workers who completed over 100 test questions across hundreds of different types of tasks, and have a near perfect overall accuracy. They are workers of the highest quality on Figure8.

Figure 1: Workflow of our study on Figure8.

In the pre-study questionnaires, 35% of users' responses leaned toward *disagree* with the statement "Abortion should be legal" and 54% leaned toward *agree* with it (11% of users reported neutral for this topic); 70% of respondences leaned toward *disagree* with the statement "Guns should be controlled" and 20% leaned towards *agree* with it (11% of users reported neutral).

Table 2 presents the proportion of users who changed their opinions after the post-query browsing session. Generally, most of the users kept their attitudes towards the topics after a short-term of interactions with the search results. For 'Abortion', around 18% of users changed their opinions (i.e. report different opinion ratings) after interacting with the web search results in *Purposeless* browsing situation, and 26% in *Purposeful* browsing situation. In case of 'Gun Control', around 30% of users changed their opinions in either *Purposeless* or *Purposeful* browsing situation. In most cases, we found that the vast majority of users who changed their opinions actually became more positive (i.e. 4% lean disagree *VS* 26% lean agree) after interacting with the web search results. However, there was an exception for topic 'Abortion', in case of *Purposeless* browsing (i.e. 14% lean disagree *VS* 12% lean agree, 75% lean disagree *VS 25%* lean disagree among neutral users), which is also different from the findings in previous work [20], where most users with neutral opinions tend to be more positive after search. Thus, we further analyze the topic effect on users' opinion change in *Purposeless* and *Purposeful* situations. A Mann-Whitney U test indicated that when they purposelessly browse the search results, there was a significant effect of topics on the changes of users' opinions ('Abortion' *VS* 'Gun Control', $p = 0.023$). But when users interact with search results purposefully, for example to collect supporting information for their opinion, their opinions still depict a positive skew in general.

In addition, it seems that users were more likely to change their opinions in *Purposeless* browsing situations than in *Purposeful* situations (e.g. Table 2, 18% opinion changes in case of 'Purposeful, Abortion' *VS* 26% in 'Purposeless, Abortion'; 0 drastic opinion changes in 'Purposeful, Gun control' *VS* 8% in 'Purposeless, Gun control'). One initial explanation is, the users' opinions could be more 'liberal' in a *Purposeless* situation, while confirmation bias was more likely to occur when users browse the search results with purpose. However, this analysis is insufficient to describe the full picture of the user's opinion dynamics.

Table 2: Proportion of users who changed opinions (black); drastically changed opinions(i.e. changed from opposite or from neutral to extreme, *italic,* light gray); proportion of neutral users who changed opinions (blue).

	Abortion		Gun Control	
	lean disagree	lean agree	lean disagree	lean agree
	2%	16%	4%	26%
Purposeful	*0*	*2%*	*0*	*0*
	17%	33%	25%	25%
	14%	12%	4%	27%
Purposeless	*5%*	*2%*	*2%*	*6%*
	75%	25%	0	50%

4.2 Roles of Users' Opinions

Next, we analyze whether and how the users' opinions affect their interactions with the search results and their perception of the search results.

Clicks. We found that users on average issued 11 clicks (4 / 3 / 4 clicks on 'pro' / 'neutral' / 'contra' search results) during their post-query browsing session. To analyze the effect of users' opinions on their clicking behavior, we conducted a one-way between users ANOVA. We found a significant difference in the number of clicks issued by users with varied pre- and post-study opinions at $p <$ 0.05 level; $F(19, 136) = 4.264$. Post-hoc comparisons using the Tukey-HSD test revealed a significantly larger number of clicks issued by users corresponding to *'pre-pro, post-pro'* in comparison to each of the other groups at the $p < .05$ level. However, we did not find a significant linear relationship between the number of clicks and the opinions of users in the post-query browsing session, using Spearmann Correlation Coefficient. We also found no significant difference on the ranking of clicked search results and users' opinions.

Browsing time. We found that the average browsing time that users spent on search results was around 394 seconds long (201 / 111 / 144 seconds on 'pro' / 'neutral' / 'contra' search results).

To investigate the effect of users' opinions on their browsing time on the search results, we conducted a one-way ANOVA. Results confirmed a significant difference in the amount of time that users with different opinions spent on browsing search results at $p < .05$ level; $F(4, 136) = 2.576$. Post-hoc comparisons using the Tukey-HSD test revealed that the time users with *'pre-pro'* opinions spent on tasks were significantly different in comparison to the time users with *'lean contra'* opinions at $p < .05$ level. In addition, users with *'pre-pro'* opinions spent significantly more time on the 'contra' search results than each of the other groups except the users with *'pre-contra'* opinions (one-way ANOVA, $F(4, 136) = 2.845, p < .05$)).

Then, we investigated the effect of users' opinions on their interactions with search results, in case of *'Purposeless'* and *'Purposeful'* situation. We found that on average users issued 12 clicks on search results in *'Purposeless'* situation which is more than the 9 clicks in *'Purposeful'* situation. We compared the effect of opinions on the number of users' clicks across the two situations using a one-way between users ANOVA. We found no significant effect across the two conditions. On average, users spent 173 seconds longer in *'Purposeless'* situation than in *'Purposeful'* situation. Through a Mann-Whitney U test, we found that the amount of time users spent on 'neutral' search results in *'Purposeless'* situation (92s) was significantly more than that in *'Purposeful'* situation (46s). From this, we reason that a purposeless browsing situation might provide users with a relatively more free atmosphere to learn about different viewpoints in detail, while users with a clear intent may potentially prefer to understand the reasons to support or refute an idea and form their own opinions.

Finally, we collected users' feedback from the standpoint of search results. Among all the users, only 28% thought the search results presented in SERPs were neutral in general. 44% of participants believed the search results were inclined to 'pro' the statement while 28% believed the opposite. We didn't find any significant topic effect on users' bias towards the attitudes of returned search results. For both topics, we found that the users with different pre-study opinions had significantly different feelings towards the search results (one-way ANOVA, $p < .05$); users with 'pro' ('contra') pre-study attitudes tend to perceive the returned results as 'inclined to pro (contra)'.

5 CONCLUSION

In this paper, we presented a study to investigate the effect of users' opinions on their interactions with web search results. To this end, we developed *Opinions*, a pseudo search engine and prepared statements with attitudes of 'pro', 'neutral' and 'contra' as the search results returned by it. Making use of *Opinions* and two popular controversial topics ('Abortion' and 'Gun Control'), we constructed a user study with 200 participants. We found that: (i) users' opinions seem to be more liberal and easier to be affected by search results when they browse the search results purposelessly; (ii) users who strongly support a topic issued significantly more clicks and spent more time on the search results; (iii) users' perceptions of the standpoint of search results could also be affected by their opinions, in that users would like to believe the search results were inclined to support their own opinions. We expect that our findings will help in better understanding the role of opinions on users' interaction behaviors with the web search results in the evaluation, analysis and study of users' opinions during web search.

REFERENCES

[1] Josef Al-Sader and Neda Ashrafi Amiri. 2016. Effects of Confirmation Bias on Web Search Engine Results and a differentiation between Non assumptive versus Assumptive Search Queries. (2016).

[2] Leif Azzopardi. 2021. Cognitive biases in search: a review and reflection of cognitive biases in Information Retrieval. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. 27–37.

[3] Ricardo Baeza-Yates. 2018. Bias on the web. *Commun. ACM* 61, 6 (2018), 54–61.

[4] Michael D Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*.

[5] Leon Festinger. 1957. *A theory of cognitive dissonance*. Vol. 2. Stanford university press.

[6] Ujwal Gadiraju, Ricardo Kawase, Stefan Dietze, and Gianluca Demartini. 2015. Understanding malicious behavior in crowdsourcing platforms: The case of online surveys. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1631–1640.

[7] Ujwal Gadiraju, Jie Yang, and Alessandro Bozzon. 2017. Clarity is a worthwhile quality: On the role of task clarity in microtask crowdsourcing. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media*. 5–14.

[8] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing controversy by connecting opposing views. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 81–90.

[9] Eduardo Graells-Garrido, Mounia Lalmas, and Ricardo Baeza-Yates. 2016. Data portraits and intermediary topics: Encouraging exploration of politically diverse profiles. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 228–240.

[10] William Hart, Dolores Albarracín, Alice H Eagly, Inge Brechan, Matthew J Lindberg, and Lisa Merrill. 2009. Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin* 135, 4 (2009), 555.

[11] Christoph Hube, Besnik Fetahu, and Ujwal Gadiraju. 2019. Understanding and mitigating worker biases in the crowdsourced collection of subjective judgments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[12] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, Filip Radlinski, and Geri Gay. 2007. Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *ACM Transactions on Information Systems (TOIS)* 25, 2 (2007), 7.

[13] Rishita Kalyani and Ujwal Gadiraju. 2019. Understanding user search behavior across varying cognitive levels. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media*. 123–132.

[14] Aniket Kittur, Jeffrey V Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 1301–1318.

[15] Q Vera Liao and Wai-Tat Fu. 2014. Can you hear me now?: mitigating the echo chamber effect by source position indicators. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 184–196.

[16] Q Vera Liao, Wai-Tat Fu, and Sri Shilpa Mamidi. 2015. It is all about perspective: An exploration of mitigating selective exposure with aspect indicators. In *Proceedings of the 33rd annual ACM conference on Human factors in computing systems*. ACM, 1439–1448.

[17] Raymond S Nickerson. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2, 2 (1998), 175–220.

[18] Scott Plous. 1993. *The psychology of judgment and decision making*. Mcgraw-Hill Book Company.

[19] Anselm L Strauss. 1987. *Qualitative analysis for social scientists*. Cambridge university press.

[20] Ryen White. 2013. Beliefs and biases in web search. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. 3–12.

[21] Ryen W White and Eric Horvitz. 2015. Belief dynamics and biases in web search. *ACM Transactions on Information Systems (TOIS)* 33, 4 (2015), 1–46.

[22] Ryen W White and Eric Horvitz. 2015. Belief dynamics and biases in web search. *ACM Transactions on Information Systems (TOIS)* 33, 4 (2015), 18.

[23] Yusuke Yamamoto and Takehiro Yamamoto. 2018. Query Priming for Promoting Critical Thinking in Web Search. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (New Brunswick, NJ, USA) *(CHIIR '18)*. ACM, New York, NY, USA, 12–21. https://doi.org/10.1145/3176349.3176377