LUND UNIVERSITY

In Search of the Climate Change Filter Bubble

A Content-based Method for Studying Ideological Segregation in Google

Genot, Emmanuel; Jiborn, Magnus; Hahn, Ulrike; Volzhanin, Igor; Olsson, Erik J; von Gerber, Ylva

2020

Document Version:
Early version, also known as pre-print

Link to publication

Citation for published version (APA):
Genot, E., Jiborn, M., Hahn, U., Volzhanin, I., Olsson, E. J., & von Gerber, Y. (2020). In Search of the Climate Change Filter Bubble: A Content-based Method for Studying Ideological Segregation in Google. Manuscript submitted for publication.

Total number of authors:
6

Creative Commons License:
Other

General rights
Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.
• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: https://creativecommons.org/licenses/

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

In Search of the Climate Change Filter Bubble: A Content-based Method for Studying Ideological Segregation in Google

Emmanuel Genot[1], Magnus Jiborn[2], Ulrike Hahn[3], Igor Volzhanin[4], Erik J. Olsson[5] and Ylva von Gerber[6]

Abstract: A popular belief is that the process whereby search engines tailor their search results to individual users, so-called personalization, leads to filter bubbles in the sense of ideologically segregated search results that would tend to reinforce the user's prior view (filter bubble hypothesis). Since filter bubbles are thought to be detrimental to society, there have been calls for further legal regulation of search engines beyond the so-called Right to be Forgotten Act (EU, C-131/12, 2014). However, the scientific evidence for the filter bubble hypothesis is surprisingly limited. Previous studies of personalization have focused on the extent to which different users get different results lists without taking the content on the webpages into account. Such methods are unsuitable for detecting filter bubbles as such. In this paper, we propose a methodology that takes content differences between webpages into account. In particular, the method involves studying the extent to which users with strong opposing views on an issue receive search results that are correlated content-wise with their personal view. Will users of with strong prior opinion that X is true on average have a larger share of (top) search results that are in favor of X than users with a strong prior opinion that X is false? We illustrate our methodology at work, but also the non-trivial challenges it faces, by a small-scale study of the extent to which Google Search leads to ideological segregation on the issue of man-made climate change.

Keywords: search engine; computer law; Google; filter bubble; ideological segregation; personalization

1. Introduction

We search for information every day using Google and other search engines. The results we get from search influence what we believe about the world, including what we believe on sensitive topics such as climate change and immigration. A popular view is that search engines and other social media can lead to so-called filter bubbles (e.g. Pariser, 2011). The idea is that since search engines base their search results not only on the search term but also on what the user has been searching on before, she eventually gets only search results that confirm her prior beliefs. Filter bubbles are, moreover, thought to be detrimental to the individual and to society at large.

The mechanism whereby search results are tailored to individual users drawing on their search history is called "personalization". Thus, the fear is that personalization leads to a personalized informational universe that becomes increasingly detached from reality itself. Some authors have argued that this threatens democracy to the extent that there is need for state regulation (Simpson, 2012). Google has already been required by EU law to remove certain search result that compromise an individual's reputation upon request, e.g. links to crimes that were committed a long time ago, under the so-called Right to be Forgotten ruling (C-131/12, 2014). The policy issue at stake in connection with personalization is whether there is need for further regulation of search engines due to filter bubbles.

Despite the obvious societal and theoretical relevance of these issues the scientific literature is surprisingly limited. There are only a handful of empirical studies that attempt to study or identify filter bubbles scientifically. In fact, there seems to be no unproblematic and generally accepted

[1] Department of Philosophy, Lund University, Sweden.
[2] Department of Philosophy, Lund University, Sweden.
[3] Department of Psychological Science, Birkbeck University of London, UK.
[4] Department of Psychological Science, Birkbeck University of London, UK.
[5] Department of Philosophy, Lund University, Sweden.
[6] Department of Philosophy, Lund University, Sweden.

methodology for this purpose currently on offer. While it is uncontroversial that search engines use personalization – searching on "proof of climate change" may turn up different results for an environmental activist and an oil company executive (Pariser, 2011, p. 3) – it has, to our knowledge, not been established whether this leads to filter bubbles.

In this paper, we propose a methodology for studying filter bubbles in search engines such as Google Search. The basic idea is to test the filter bubble hypothesis – the hypothesis that search leads to ideologically segregated results – by comparing the contents of search results for identical searcher for users with strong and opposing view on a particular subject. Do they get results that are systematically correlated with their prior view so that a user with a strong pro-attitude gets more pro-links, and so on? More precisely, will users of with strong prior opinion that X is true on average have a larger share of (top) search results that are in favor of X than users with a strong prior opinion that X is false? We illustrate our methodology at work, but also the non-trivial challenges it involves, by means of a small pilot study on the issue of man-made climate change. In section 2, we give an overview of how personalization has been thought to raise filter bubble worries. We also survey relevant empirical literature. In section 3, we present our own methodological suggestion which we then, in section 4, illustrate in relation to the issue of climate change. Some challenges for our approach and suggestions for improvements are discussed in section 5.

2. Background

It has been argued that personalized search may potentially result in users being trapped in informational "filter bubbles" (Pariser 2011). In a filter bubble search results tend to confirm and reinforce peoples' prior views, and hide results that challenge them. As Pariser (2011) puts it, "[m]ore and more, your computer mirror is a kind of one-way mirror, reflecting your own personal interests while algorithmic observers watch what you click" (p. 3). A similar worry is raised in Hannak et al (2013): "The increasing personalization is leading to concerns about Filter Bubble effects, where certain users are simply unable to access information that the search engine's algorithm decides is irrelevant" (p. 527).

The underlying worry is that filter bubbles will have detrimental epistemological and social effects. For instance, Pariser (2011), believes that filter bubbles will give rise to fewer opportunities to experience genuine new insight: "In the filter bubble, there's less room for the chance encounters that bring insight and learning" (p. 15). Harvard law professor Cass Sunstein connects filter bubbles to democracy suggesting that "people should be exposed to materials that they would not have chosen in advance" because "[u]nplanned, unanticipated encounters are central to democracy itself" (Sunstein 2001, p. 8). In the opinion of Liljeblad (2012), filter bubbles "favor homophilic networks" (p. 110) and their presence "reduces the potential of the internet to support some interpretations of global civil society, particularly interpretations envisioning consensus, or coalition-building" (ibid.).

However, many of these worries about filter bubbles have not yet been rigorously evaluated, partly because they are precisely worries rather than full-fledged arguments. An admirably detailed objection to personalization in search engines can be found in Simpson (2012). Simpson argues that personalized search threatens the objectivity of search results. Given personalization, search results will be ordered according to the user's preferences rather than "objectively", i.e. according to a defensible judgment of relative relevance. Objectivity, moreover, is a "public good" because it is necessary for understanding. Finally, voters in a democracy need to understand the issues they vote on, so that their votes are not completely uncorrelated with the facts of the matter. Simpson concludes that "there is a prima facie case for justified intervention by government in how search engines work" (p. 441). His concrete proposal in terms of state regulation is that non-personalized search should be the default option, rather than – as is now the case in Google (since 2009) – the other way round.

Turning now to the empirical literature on filter bubbles and related phenomena, selective exposure to news in general has been studied since the 1940s. For instance, in an early study it was found that "people select their exposure along the line of their political disposition" (Lazarsfeld et al 1944). More recently, Garrett and Resnick (2011) found that people are more likely to view news that support their

prior opinions. Gentzkow and Shapiro (2011) studied ideological segregation in news in an American context as measured by "isolation index", an index originally developed to measure racial segregation. Some results that came out of this study were that the online news isolation index is 7,5% (66,5% - 51%) that online news is more segregated that radio, TV and local news, but less segregated than national newspapers (10%) and much less ideologically segregated than neighborhoods, workplaces and voluntary organizations.

Filter bubbles can be approached under the umbrella of ideological segregation, as in the observational study reported in (Flaxman, Goel, and Rao, 2016). The study examined the browsing history of 1.3 million of U.S. users, over a period of 3 months, in order to determine whether or not users were relatively more exposed to news that agreed with their own political views when accessing news items recommended by search and social medial. They found that news access was in fact highly segregated. However, and somewhat counterintuitively, online search via search engines, and social media tools, increased the chances of exposure to opposed views.

Bakshy, Messing and Adamic (2015) studied selective exposure in Facebook. The authors are Facebook employees with access to 10 million active US FB-users with self-reported political affiliation, 900 million news exposures and 59 million clicks. One result that came of this study was that all three factors diminished share of cross-cutting content. Another upshot was that friends had the greatest impact, followed by individual choice. The Facebook algorithm itself was suggested to have a relatively minor influence. The fact that the study was made by researchers with ties to Facebook itself should be noted in this connection.

There are only a few studies of personalization effects in Google and there is no consensus on the most appropriate methodology. As Hannak et al (2013) note, "to date, there has been little scientific quantification of the basis and extent of search personalization in practice" (p.527). Their intriguing study measured differences for 200 actual Google users compared to control computers without search history. The study found that there was an 11.7% percent difference which the authors attribute to personalization in Google. They went on to investigate factors that influence personalization using artificial Google accounts concluding that there was no measurable search history-driven personalization effect. The only factors that were found to have significant effects were being logged in to the Google system (YouTube, Google+ and so on) and geographical location.

An advantage of the approach of Hannak et al is its generality. The methodology is not specific to Google Search but can also be used for other search engines (Bing etc.) and for other web-based services (Twitter search etc.). However, Hannak et al note a number of limitations with their study (p. 536). Crucially, from our perspective, they focus on quantifying literal differences in search results, i.e. whether a.com is different from b.com. They do not address the issue of semantic differences, i.e. whether or not a.com and b.com contain different information content. If so, what is the impact of these differences? They note that "semantic differences and impact are challenging to quantify" (p. 536). In other words, the method used by Hannak et al is in the end unable to detect filter bubbles as such. The method we will propose takes semantic differences of webpages into account. Thus our methodology is a candidate method for detecting filter bubbles as such.

3. A method for studying filter bubbles in search engines

In this section we describe our methodology in the abstract and in the next section we illustrate the methodology at work vis-à-vis the debate about man-made climate change. Consider the following general Filter Bubble Hypothesis for a given search engine E which is assumed to use personalized search:

(FBH) The search engine E gives rise to ideologically segregated search results depending on user's prior views in a way that is likely to reinforce those views.

We can test this hypothesis by focusing on a particular issue X. In other words, a testable consequence of FBH is:

(FBH1) Users of E with strong prior opinion that X is true will, on average, have a larger share of (top) search results that are in favor of X than users with a strong prior opinion that X is false.

In order to test FBH1 we propose the following general methodology.

Step 1: Select an issue X such that people strongly disagree about facts concerning X.

The issue could in principle be any issue which people find sufficiently interesting so that they use their favorite search engine to search for information on a regular basis. Unless participants search on a regular basis for information, there is little chance that there will be a detectable effect of personalization on their search results because the personalization algorithm will not have a much to go on. The issue we selected for our pilot study (see next section) was whether there is man-made climate change but in principle any reasonably "hot" topic will do (with some qualifications; see section 5 for a discussion of this point.)

Step 2: Devise a questionnaire including questions that would be an indication of subject's prior attitude on X.

The participants should be able to provide answers on some suitable scale, e.g. a scale of 5 going from "strongly agree" to "strongly disagree".

Step 3: Select a set of search terms related to X.

The purpose of the search terms is that the participants should use the search terms for searching information using the search engine E.

Step 4: Recruit test persons that have a strong and opposing opinions on X, and a history of using the search engine E to look for information about it. Also, a control group of participants with no known preferences regarding X needs to be recruited.

The focus on participants with a strong view is useful for identifying an effect that may conceivably not be very pronounced. If there is a filter bubble effect at all, it should be detectable in participants who have strong opposing views on the issue at hand. Again, people with strong views are more likely to be engaged in the issue and use the search engine for searching for information on it, which is crucial if we wish to find an effect at all since personalization that leads to filter bubbles can be expected to depend only or primarily on the user's search history.

Step 5: Instruct the participants to fill in the questionnaire and to search using the key words, and then return the questionnaire and search results list in some way to the experiment leader (e.g. on flash drive in an envelope returned by ordinary mail).

The latter task could, for example, be accomplished by asking participants to download the latest version of Mozilla Firefox and Mozilla Archive Formal-plugin for saving search results with links preserved. It suffices to save the first three search result pages as several studies have found that users rarely go beyond them (Pan et al, 2007; Salmeron et al 2013). Since search engines can be updated and changes, the searches have to be made at the same point in time, minimally on the same day. Also, if search engine E has a personalization option, the participants need to be instructed to have it on.

Once the subjects have enrolled that are asked to follow the instructions uploaded on a USB-memory stick sent out together with the questionnaire by ordinary mail with a pre-stamped return envelope. The tasks include downloading the latest version of Mozilla Firefox + Mozilla Archive Format-plugin for saving search results with links preserved. All searches should be done the same day for all subjects on the computer they normally use, and subjects are to save the first three results pages for each query.

Step 6: Recruit an independent person who lacks knowledge of the purpose of the study to code (categorize) the search results on the participants' lists in terms of content on webpages in question as belonging to one of three categories: pro X, contra X or neutral.

Step 7: FBH1 can now be tested using the questionnaires and coded result lists returned by the participants.

For example, do the participants with a strong prior opinion that X is true (according to the questionnaire) have a larger share of search results that are pro X (according to the coded search result lists) than users with a strong prior opinion that X is false? If so, FBH1 (FBH) has been confirmed. If not, it has been disconfirmed.

Figure 1 describes the tasks to be carried from a participant perspective.

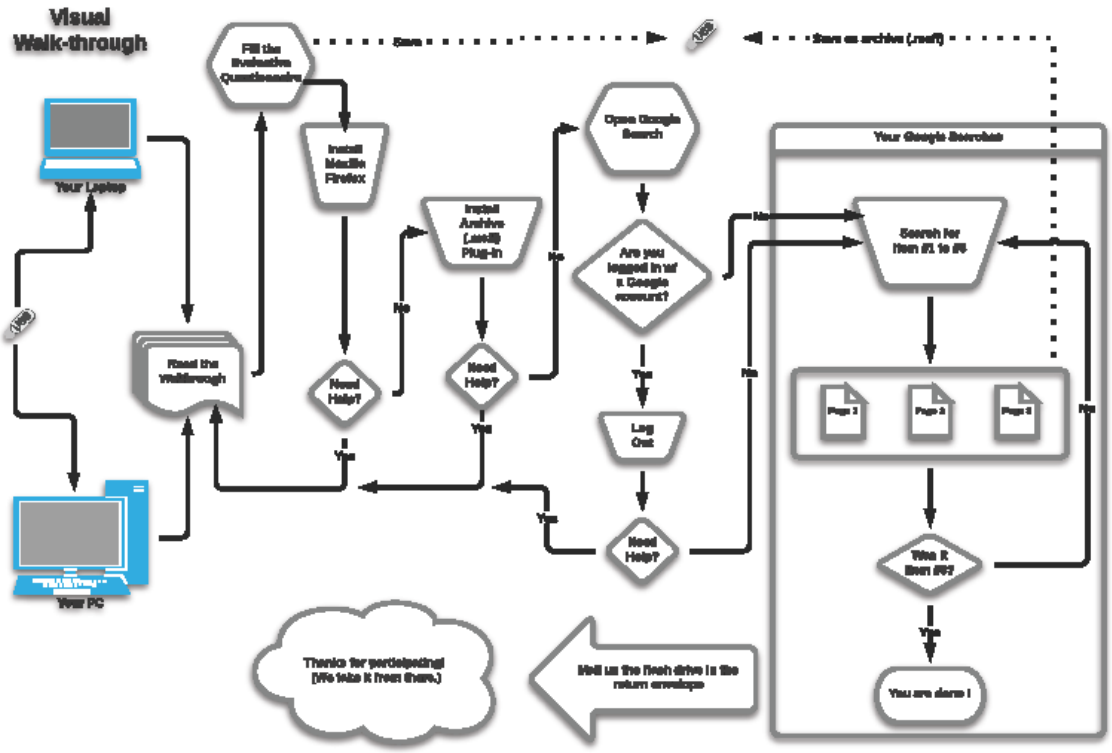

Figure 1: Visual walk-through for participants.

There are various other techniques available for testing the main hypothesis. However, rather than describing these techniques in general terms we will show how they work in the context of a concrete example.

4. In search of the climate filter bubble: a pilot study

Pariser (2011), p. 3, mentions the climate change debate as an area that might be affected by filter bubbles: "With Google personalized for everyone, the query 'stem cells' might produce diametrically opposed results for scientists who support stem cell research and activists who oppose it. 'Proof of climate change' might turn up differently for an environmental activist and an oil company executive."

Following Pariser, we decided to use our method to test the Filter Bubble Hypothesis for Google Search (FBHGS) with regard to climate change. Thus, we focused on the following hypothesis:

(FBHGS) Google Search gives rise to ideologically segregated search results on climate change depending on user's prior views in a way that is likely to reinforce those views.

More specifically, we focused on the claim that climate change occurs and is man-made. We will refer to this hypothesis as the Man-Made Climate Change Hypothesis (MMCCH). We considered the following testable consequence:

(FBHGS1) Users of E with strong prior opinion that MMCCH is true will, on average, have a larger share of (top) search results that are in favor of MMCCH than users with a strong prior opinion that MMCCH is false.

In order to test this hypothesis we followed the general methodology as described in the previous section with participants from Sweden. Since our aim was to study the effect of Google Search per se we asked participants not to log into the Google system. The next step was to devise a questionnaire that would reveal the participants' prior attitudes to man-made climate change. We have included the questions that we settled upon in Appendix 1.

Concerning search terms we decided for "arctic ice", Antarctic ice", "climate change", "climate change East Anglia", "global warming", "greenhouse effect", "man made climate change" and "severe climate". Some search terms were chosen for their potential of higher occurrence on one side of the debate, such as "arctic ice" and "Antarctic ice". Others were chosen for their potential of returning different results depending of the opinion of the participant, such as "climate change East Anglia", the university of East Anglia being not only an international research center on climate change but also the origin of the so-called "climategate emails" in which researchers seem to express doubt about climate change being man-made.

In order to recruit subjects we used known networks of climate activists and sceptics in Sweden, including bloggers, environmental organizations and climate researchers. A control group with no known sympathies in either direction was recruited at the university café and library. The method used to recruit participants with known opinion was that of snowball sampling, although on a rather small scale. Specifically, one of the authors had had prior contacts with people involved in both sides of the climate change debate in Sweden, and contacted them. The inclusion of bloggers, members of environmental organizations and climate researchers, rather than merely strongly opinionated persons, was meant to guarantee that our participants would have actually, in the past, searched for climate-change-related topics, and generated a sizeable amount of data that personalization algorithms could operate upon. A second questionnaire was given to our participants, including questions about how many other people used the participants' terminal, preferences for sources of information, and overall perception of the search result. This questionnaire can be found in Appendix 2.

Once the subjects had enrolled they were asked to follow the instructions that were uploaded on a USB-memory stick that was sent out by ordinary mail, together with questionnaire (views on climate change and internet use) and pre-stamped return envelope. The text introducing the participants to the experiment can be found in Appendix 3. The tasks included downloading the latest version of Mozilla Firefox + Mozilla Archive Format-plugin for saving search results with links preserved. All searches should be done in Google, without being logged into the Google system, the same day for all subjects on the computer they normally use, and subjects were asked to save the first three results pages for each query. The final instruction was to save the search results and questionnaires to the USB and return it without a name.

28 persons initially agreed to participate. 13 subjects were found to be pro and 8 contra vis-à-vis MMCCH. The control group contained 7 subjects. Among those who agreed to participate, 18 subjects completed the task. Two of the submitted responses could not be used due to procedural issues. In the end we had 8 pro, 3 contra and 5 control group responses. This was fewer than expected, although every user generated a respectable amount of data.

Once the responses were in we recruited an external person without knowledge about the purposes or setup of the experiment to code (categorize) links in search results on the basis of the content of the corresponding webpage as belonging to one of the following three categories: pro MMCCH, contra

MMCCH or neutral. We then went on to compare the share of links of each category in the different groups.

The total number of links in the search lists was 3840. Of these links 3168 could be analyzed; the others were broken. There was a massive overweight of pro or neutral links. The most common sceptic search results resulted from searching on "Manmade climate change". Figure 2 shows the number of working links in each category and Table 1 the number of unique links per search term.

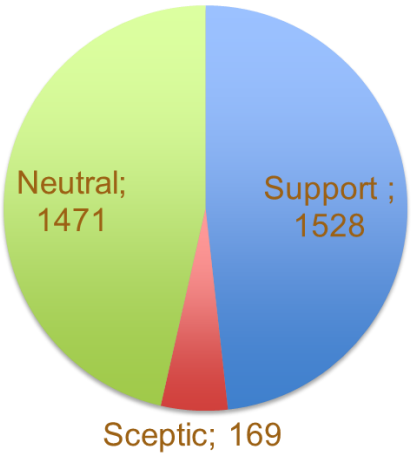

Figure 2: Number of working links in each category.

Search term	# unique links
Antarctic Ice	60
Arctic Ice	48
Climate Change	48
Climate Change East Anglia	65
Global Warming	52
Greenhouse Effect	42
Manmade Climate Change	80
Severe Climate	75

Table 1: number of unique links per search term

An obvious way of analyzing the results is to compare the share of links pro (support) or con (sceptic) vis-à-vis MMCCH in the result lists of different subjects. If FBHGS1 is correct, we should expect subjects from the sceptic group to have a larger share of sceptic links and a smaller share of supporting links in their results than subjects from the pro-group. Figure 3 shows the share of links in the various groups of subjects. The shares were essentially the same in the pro group and the control group.

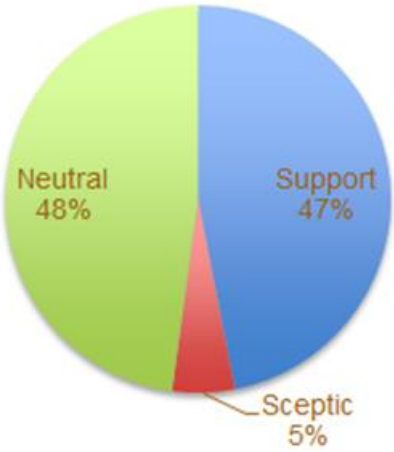

Figure 3: Share of links in the sceptic group.

There could still be a difference in the ordering of the links however, so that sceptic links where higher ranked for sceptic subjects and vice versa. To investigate this possibility, we calculated the share of pro/con links on the first result page for subjects from the different groups, and also the average position of the first sceptic link on the first page. Figure 4 shows the shares of first page results in the various groups of subjects.

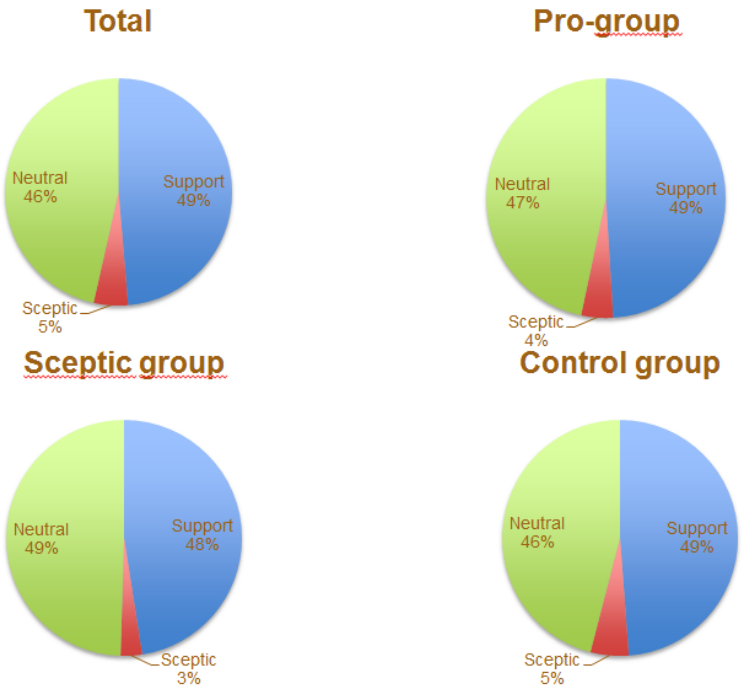

Figure 4: Share of first page results in the various groups of subjects.

Again, the differences were very small, with sceptics having slightly *less* sceptic links on their first result pages than subjects from the pro-group. The average location of first sceptic link on first page for the groups was 4.18 for the pro group, 4.33 for the sceptic group and 4.2 for the control group.

Another way of testing the hypothesis is by studying the heat map of search results (figure 5).

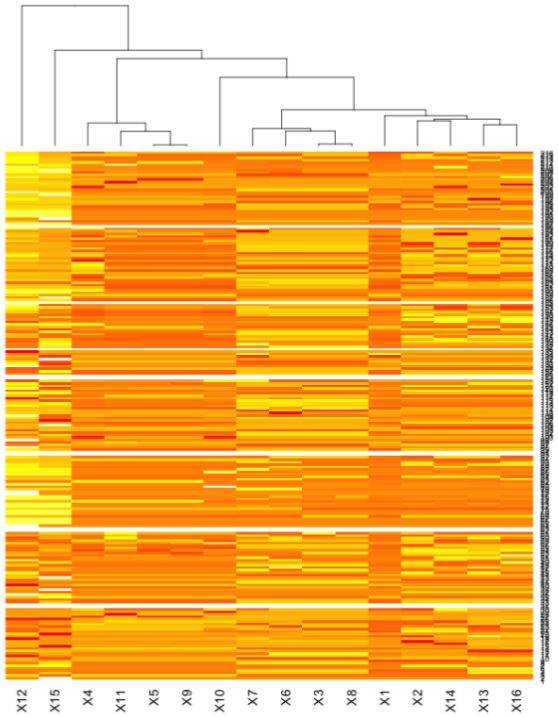

Figure 5: Heatmap for the search term “global warming”.

The heatmap in figure 5 results from the search term “global warming”. The numbers are mapped onto colors. The heatmap gives a gross visualization of variability of links returned across participants. On the y-axis are the link position, numbered consecutively, i.e. 1 is the first items returned to first search term. On the x-axis are the participants. If all search results were the same, we would see uniform stripes across the entire row of the image, for all rows. This is not what we see in figure 5, meaning that the variability is considerable which is what we would expect from personalization.

More interesting is the question of how the variability seen in figure 5 relates to groups. The participants have been ordered according to a cluster analysis on the link rank order vector (i.e. the link number that goes in each slot from 1 to 340). If results were strongly influenced by filtering according to prior views, the clusters that can be seen in the heat map would largely coincide with subject groups, at least for the pro-group (participants X1- X8) and the sceptic group (participants X14-X16). This, however, is not the case.

5. Discussion

Our pilot study was based on a small sample and gave no conclusive support either for or against the filter bubble hypothesis in the context of Google Search. In particular, we found no evidence that prior opinions affected either the share of supporting/sceptic results that the subjects received or the ordering of the results. This said, one might speculate that, had there been a very strong filter bubble effect, even our limited study would have identified it. The effect, if there is one, is not very strong, it would seem.

In a second experiment we used online recruitment in collaboration with an NGO that agreed to announce the study on its homepage and Facebook page. Participants were asked to sign up on Google Form. There was no explicit control group and the opinions were mapped from a questionnaire. To reduce complexity, instructions – including a detailed video description of the procedure – were made available online 1-2 weeks before the actual search experiment took place. Virtual flash drives were sent out before the experiment. Participants who finished the experiment were sent physical flash drives with pre-stamped return envelope to save and return results. 31 participants registered and 16 confirmed participation. However, only 6 flash drives were actually returned. The cause of this

unexpected dropout is unknown. Since many registered participants were university employees and Swedish universities generally disallow staff members to install their own software, one problem may have been the instruction to install Mozilla Firefox for saving search results with links included.

Many issues arise when actually implementing our methodology in a concrete case. Recruiting participants is often one of the main obstacles in empirical studies, and ours is no exception, as illustrated by our pilot study on climate change. Participants need to be found that not only have strong opposing views on the underlying issue (climate change etc.) but who have also used the search engine for searching about information on that issue. Furthermore, our method presupposes that we have access to information about participants views on the issue at hand. Depending on what the underlying issue is, there may be ethical or legal constraints regarding how the information about participants' prior views can be collected or stored, and it may be difficult to secure the trust of the participants that the data will be handled correctly and not misused. For instance, the Swedish constitution disallows storage of public records on citizens' political views. In many cases, it will be desirable or even obligatory for participants to return questionnaires and search results in a way that secured anonymity. This normally means that the study has to be conducted in such a way that the search lists cannot be traced back to individual subjects while still making it possible to connect search results to prior views. In our pilot study, we solved these problems by having participants returning their data on a USP flash drive enclosed in an anonymous envelope.

Snowball sampling is particularly well-suited to recruit participants within a particular group. As it turned out, the particular challenge we faced in our pilot study was that there are rather few climate skeptics in Sweden, and probably fewer still expressing their views online. We therefore had a limited pool of suitable participants to enroll on the sceptic side. A solution to this problem would be to choose an issue on which people are divided more evenly, such as the immigration policy issue. Future studies will explore this possibility.

A further difficulty with our approach has to do with the fact that search engines and Google Search in particular is in a certain sense a "moving target" and may be subject to change from one day to another. Therefore, the searches would ideally have to be made at the same point in time. We needed to solve this problem in a way that allowed for easy collection and processing of search results and yet avoid drop-outs due to complex procedure, which was no easy task.

In future work, we will consider relaxing the requirement that the searches will have to be made on roughly the same point in time. If the number of participants from the various groups is sufficiently large, any effect in search results that is due to a change in the search engine over time could be expected to cancel out. We would also like to extend our study in a number of directions. First, we would like to explore the option of having subjects log into the Google system. There are reasons to believe that the degree of personalization is greater if users are logged into the system than if they use Google Search as a stand-alone resource, as suggested in Hannak et al (2013). We would also like to explore different ways of generalizing our methodology beyond search engines, such as social networks like Facebook and other online resources.

6. Conclusions

Previous studies of personalization in search engines have focused on quantifying literal differences between web pages, i.e. whether one webpage is different from another, without taking the content on the webpages into account. In this paper, we proposed a methodology that takes content differences between webpages into account. The main advantage of our methodology is that it makes it possible to study the extent to which personalization gives rise to filter bubbles, in the sense of ideologically segregated search results that would tend to reinforce the users' prior views. If personalization in search engines gives rise to filter bubbles, a case can be made for legal regulation of search engines. We illustrated our methodology and the non-trivial challenges it involves by a small-scale pilot study of the extent to which Google Search leads to ideological segregation on the issue of man-made climate change. Our pilot study, which was based on a small sample, gave no conclusive support either for or against the filter bubble hypothesis in the context of Google Search. In particular, we

found no evidence that prior opinions affected either the share of supporting/sceptic results that the subjects received or the ordering of the results. While our pilot study is inconclusive, we conjecture that had there been a very strong filter bubble effect, even a small study like ours would have identified it. Finally, we suggested some ways in which our method could be improved.

Appendix 1

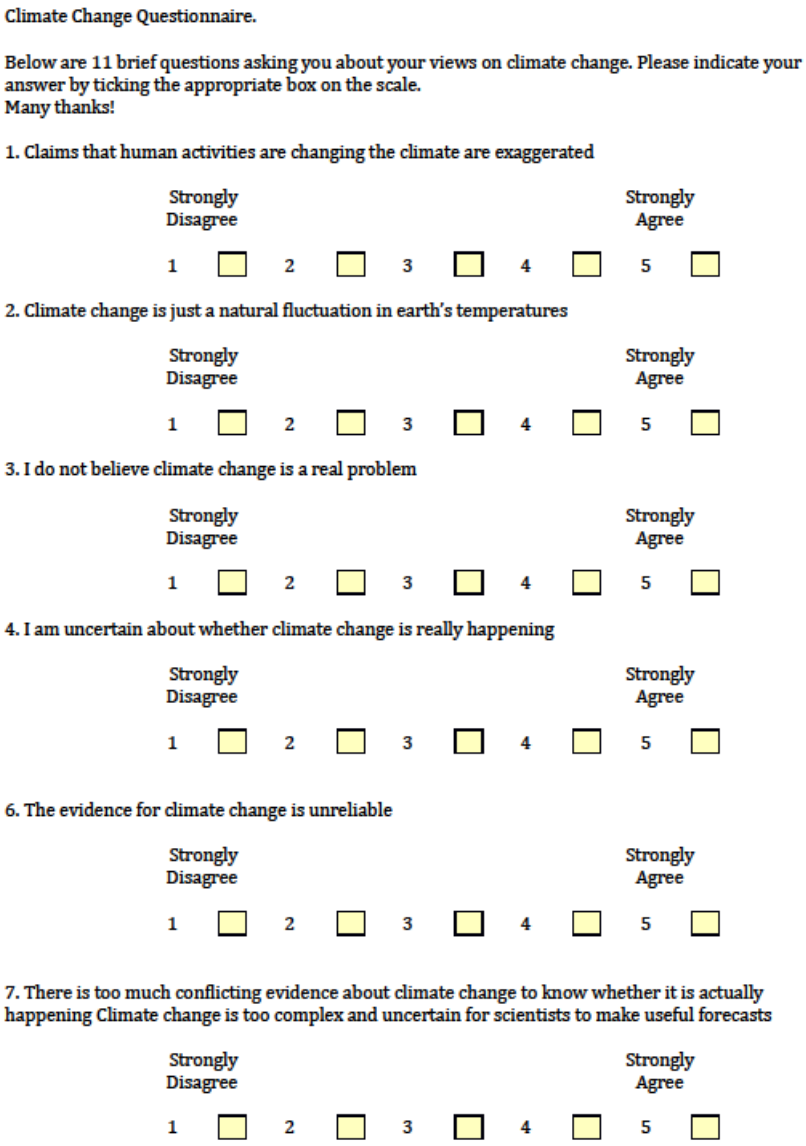

Climate Change Questionnaire.

Below are 11 brief questions asking you about your views on climate change. Please indicate your answer by ticking the appropriate box on the scale.
Many thanks!

1. Claims that human activities are changing the climate are exaggerated

Strongly Disagree — Strongly Agree

1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

2. Climate change is just a natural fluctuation in earth's temperatures

Strongly Disagree — Strongly Agree

1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

3. I do not believe climate change is a real problem

Strongly Disagree — Strongly Agree

1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

4. I am uncertain about whether climate change is really happening

Strongly Disagree — Strongly Agree

1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

6. The evidence for climate change is unreliable

Strongly Disagree — Strongly Agree

1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

7. There is too much conflicting evidence about climate change to know whether it is actually happening Climate change is too complex and uncertain for scientists to make useful forecasts

Strongly Disagree — Strongly Agree

1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

Figure 6: Climate change questionnaire.

Appendix 2

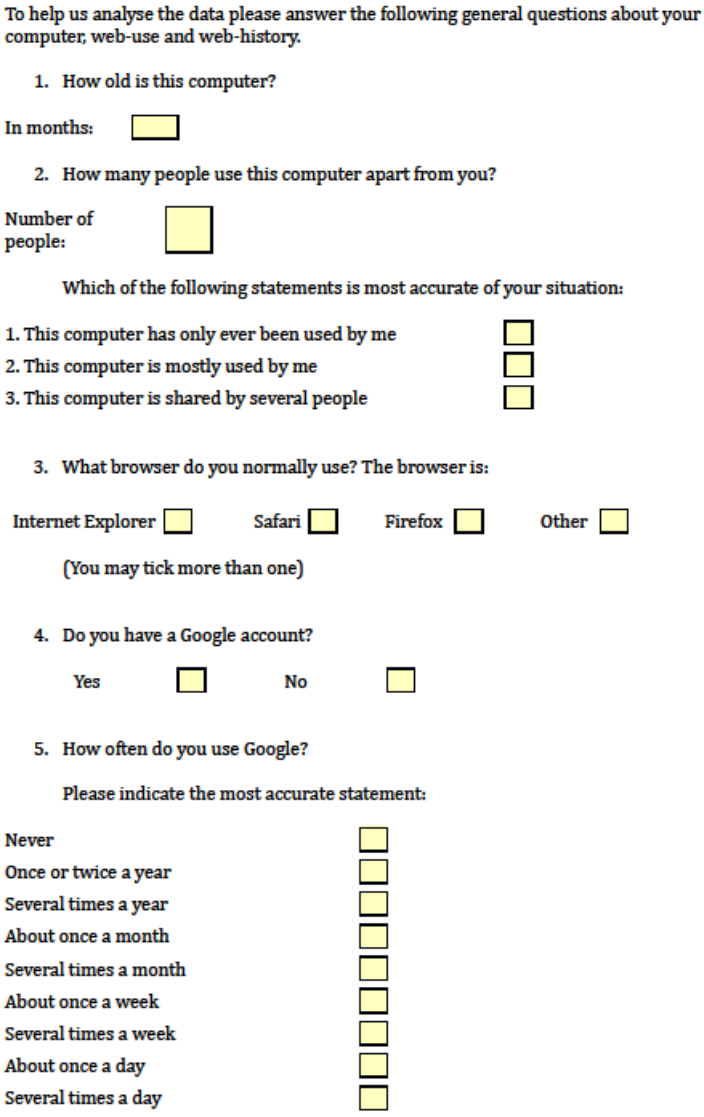

To help us analyse the data please answer the following general questions about your computer, web-use and web-history.

1. How old is this computer?

In months: ☐

2. How many people use this computer apart from you?

Number of people: ☐

Which of the following statements is most accurate of your situation:

1. This computer has only ever been used by me ☐
2. This computer is mostly used by me ☐
3. This computer is shared by several people ☐

3. What browser do you normally use? The browser is:

Internet Explorer ☐ Safari ☐ Firefox ☐ Other ☐

(You may tick more than one)

4. Do you have a Google account?

Yes ☐ No ☐

5. How often do you use Google?

Please indicate the most accurate statement:

Never ☐
Once or twice a year ☐
Several times a year ☐
About once a month ☐
Several times a month ☐
About once a week ☐
Several times a week ☐
About once a day ☐
Several times a day ☐

Figure 7: Web-use questionnaire.

Appendix 3

Instructions to participants (as presented in a file labelled “Read Me First”):

“Welcome, and thank you for your participation to this exciting experiment! You will be asked to conduct a series of pre-specified web searches and save their results. How this is done needs to be standardized across all our participants, so we provide very detailed step by step instructions. It is important that you adhere exactly to the instructions in all detail; otherwise the results will not be interpretable.

For now, here is what you should do:

1. Fill the consent form that you will find in one of the two stamped envelopes, and send it as soon as possible. Keep the other envelope to send flash drive (separate envelops are needed in order to maintain anonymity).

2. Plug the flash drive in your computer, and open the folder 1_Open This Folder First.

3. Open the file 1.2_Visual Walkthrough.pdf and print it. It will remind you of the essential

steps of the experiment.

4. Find the document 2.1_Step-by-Step Walkthrough.pdf in the folder 2_Helping Stuff and read it carefully. It contains all the instructions for carrying out the experiment.

Be aware that the process of installing Mozilla Firefox and the Add-on needed for the experiment may be time-consuming. We thus strongly recommend that you install Firefox as soon as possible as detailed in the document 2.1_Step-by-Step Walkthrough.pdf. Finally, don’t throw away the flash drive blister: you will need it when you send us the flash drive in the return envelope, in order to provide it with some protection. Thanks again for your participation. We are looking forward to receiving your data!”

References

Bakshy, E., Messing, S., and Adamic, L. A. (2015), ”Exposure to ideologically diverse news and opinion on Facebook.” Science 348 (6239): 1130-1132.

Flaxman, S., Goel, S., and Rao, J. M. (2016), “Filter bubbles, echo chambers, and online news consumption”, Public Opinion Quarterly 80: 298-320.

Garrett, R. K., and Resnick, P. (2011), “Resisting political fragmentation on the internet”, Daedalus 140: 108-120.

Gentzkow, M., and Shapiro, J. M. (2011), “Ideological segregation online and offline”, Quarterly Journal of Economics 126, 1799-1839.

Hannak, A., Sapiezynski, P., Kakhki, A. M., Krishnamurthy, B., Lazer, D., Mislove, A., and Wilson, C. (2013), “Measuring personalization of web search”, Proceedings of the 22nd international conference on World Wide Web.

Lazarsfeld, P., Berelson, B., and Gaudet, H. (1944), The people’s choice: how the voter makes up his mind in a presidential campaign, New York: Duell, Sloan and Pearce.

Liljeblad, J. (2012), “The implications of personal internet search for theories of global civil society”, International Journal of Technology, Knowledge and Society 8 (1): 103-113.

Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., and Granka, L. (2007), “In Google we trust: users’ decisions on rank, position, and relevance”, Journal of Computer-Mediated Communication 12 (3): 801-823.

Pariser, E. (2011) The filter bubble: what the internet is hiding from you. Penguin Books.

Salmerón, L., Kammerer, Y., and García-Carrión, P. (2013), ”Searching the web for conflicting topics: page and user factors”, Computers in Human Behavior 29: 2161–2171.

Simpson, T. W. (2012), “Evaluating Google as an epistemic tool”, Metaphilosophy 43: 426-445.

Sunstein, C. R. (2001), Republic.com. Princeton University Press.