

Can You Hear Me Now? Mitigating the Echo Chamber Effect by Source Position Indicators

Q. Vera Liao & Wai-Tat Fu

Department of Computer Science

University of Illinois at Urbana-Champaign

201 N. Goodwin Avenue, Urbana IL, 61801

liao28, wfu@illinois.edu

ABSTRACT

We examined how a source position indicator showing both valences (pro/con) and magnitudes (moderate/extreme) of positions on controversial topics influenced users' selection and reception of diverse opinions in online discussions. Results showed that the indicator had differential impact on participants who had varied levels of accuracy motives – i.e., motivation to accurately learn about the topic, by leading to greater exposure to attitude-challenging information for participants with higher accuracy motives. Further analysis revealed that it was mainly caused by the fact that the presence of position indicator increased the selection of moderately inconsistent sources for participants with high accuracy motives but decreased the selection of them for participants with low accuracy motives. The indicator also helped participants differentiate between sources with moderate and extreme positions, and increased their tendency to agree with attitude-challenging information from sources with moderately inconsistent positions. Participants with high accuracy motives were also found to learn significantly more about the arguments put forward by the opposite side with the help of the position indicator. We discussed the implications of the results for the nature of the echo chamber effect, as well as for designing information systems that encourage seeking of diverse information and common ground seeking.

Author Keywords

Selective Exposure; Motivation; Information Diversity.

ACM Classification Keywords

H.5.4 [Information Interfaces and Presentation]: Hypertext/Hypermedia-User issues.

General Terms: Human Factors; Design.

INTRODUCTION

Exposing people to challenging or alternative views is essential for promoting critical thinking and informed decision-making [19, 24], preventing or correcting

inaccurate beliefs or even dangerous radicalization [32]. To encourage exposure to diverse opinions, researchers have stressed the necessity of providing “safe spaces” where people feel comfortable sharing their views and interacting with people of divergent perspectives [16, 32]. This “marketplace of ideas” is central to deliberative democracy [20], and also critical in many other information-seeking or decision-making contexts where people face controversies or choices [14].

Internet can be an ideal venue for such “marketplaces”, given its provision of unprecedented access to information and people of widely varied backgrounds. Latest technology can further promote the access in a more centralized manner. For example, information aggregators allow users to view information from different sources at the same time. Social technologies, such as online discussion forum and social sharing service could assemble people of diverse opinions and open them up to opportunities for communication and education. On the other hand, critics have warned that, by facilitating easy control of one's own information diet, the Internet may exacerbate the “echo chamber effect” [1, 27] – i.e., that people may end up only interacting with others who share similar views. It is often attributed to a phenomenon called *selective exposure* to information, which can be defined as individuals' tendency to favor consonant information and avoid dissonant information. In other words, the increasing availability of diverse information on the Internet does not guarantee an equally diverse exposure to different perspectives if people lack the initiatives to attend to information that challenges their existing attitudes.

From this perspective, it is important for researchers to ask what can motivate Web users to attend to attitude-challenging information. Research in this area suggests that there are two critical aspects that impact selective exposure to information. First, researchers have studied how various design features [15, 17, 22, 28, 35] can promote exposure to diverse perspectives, such as better ways of organizing and presenting information of opposite stances. Second, recent research on Internet selective exposure suggests that there is a divergence on the motives of users to seek diverse information: while some users prefer avoiding dissonant information, some may not favor an exclusive diet of agreeable information [23]. In fact, psychological studies

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
 CSCW '14, February 15–19, 2014, Baltimore, Maryland, USA.
 Copyright © 2014 ACM 978-1-4503-2540-0/14/02...\$15.00.
<http://dx.doi.org/10.1145/2531602.2531711>

have found two competing tendencies when people seek information [12] – a tendency to seek attitude-reinforcing information to avoid cognitive dissonance, and a tendency to achieve an accurate understanding of the reality by seeking relevant information regardless of its position (often called *accuracy motives*). Given that these tendencies often vary among individuals and topics, effects of certain interface designs on users' behavior may also vary.

In this paper, we studied how an interface feature indicating the position of the information source impacted users' selection and reception of attitude-challenging information through a laboratory experiment. More importantly, we studied its differential effects on users who had varied levels of accuracy motives. To the best of our knowledge, few studies have systematically examined the interaction between interface features and users' accuracy motives on the consumption of diverse information. However, given their key roles in influencing selective exposure tendency, the interaction may shed light on what designers can do to encourage attention and consumption of diverse information for users with different backgrounds and information needs.

The indicator feature we introduce here distinguishes itself from those in previous research in two aspects. First, we emphasize its reflection on the position of the information *source* rather than the *content* of a single piece of information. It is particularly pertinent to the design of social technologies. For example, it is common that on political discussion websites people indicate their political affiliation in their profile. Often we can also infer others' positions on particular issues based on more implicit information such as their previous posts, votes, groups they join or users they follow. This kind of data also makes it possible for, e.g., some forms of reputation or recommender system to infer the general position of the information source. We are interested in how explicit indication of such source position information could impact users' interaction with diverse information. Second, the indicator not only shows the valence, i.e., positive or negative, but also the magnitude, i.e., extreme or moderate, of the information source's position. We believe this is more ecologically valid as it is well known that people's attitudes are more likely distributed along a continuum rather than dichotomous. This would also allow us to study how the magnitude of discrepancy between the source position and user's pre-existing attitude impacts the consumption of information from diverse sources.

We expect the source indicator to have impact on different stages of users' information consumption, specifically, their selection, reception and learning outcomes of diverse opinions. For information selection, we are interested in whether the source indicator can motivate users to approach attitude-challenging information, i.e., reduce their selective exposure tendency. For information reception, we examine how the indicator can induce changes in users' perception

of source positions and their agreement with different opinions. Last but not the least, we are interested in whether users can acquire new knowledge from diverse positions.

Biased views and inferior decisions are often caused by false perception of consensus or uniqueness (either side may possibly lead to bias [9, 28]), resistance to appreciating divergent opinions [16], and most importantly, ignorance of alternative views and facts that may challenge the existing views [8, 27, 31, 35]. Therefore, we believe the measurements of selection, position judgment, agreement and knowledge gain are important and especially pertinent to the goal of designing technological interventions that correct biased views and support informed decision-making.

In summary, we asked the following research questions:

RQ1: How does the indicator of positions of information sources impact users' selection of attitude consistent and inconsistent information (**RQ1a**)? Does it have differential impact on people who have high or low accuracy motives (**RQ1b**)?

RQ2: Do both the valence and magnitude of information source position reflected in the indicator affect users' selection of diverse information?

RQ3: Does the indicator have influence on users' perception of the information source's position *after* they process the information?

RQ4: Does the indicator influence users' agreement with the information from sources of different positions? As people with more extreme positions tend to hold stronger disagreement with opposing views, does the indicator moderate this tendency?

RQ5: Does the addition of the indicator feature lead to differences in users' knowledge gained on either side?

Before discussing the study, we will briefly review the selective exposure theory, previous research on designing interface for diversity seeking, and relevant psychological theories that motivated the design of the study.

RELATED WORK

In psychological research, selective exposure has been mainly explored based on the cognitive dissonance theory, which states that people are motivated to avoid the psychological discomfort incurred by having to reevaluate their prior decisions or existing attitudes [9]. A current focus of this research field is on identifying moderating variables that affect this tendency. In particular, research suggests that in some situations people are less likely to engage in selective exposure: i.e., when they are curious about the topic [9], when the dissonant information has high utility [31], when the decision is related to important outcomes [12, 30], or when the person is generally open-minded and value the norm of diversity [12]. Hart et al. [12] concluded that accuracy motive, defined as the desire to form accurate appraisal and valid representation of reality, is the underlying factor that promotes the tendency to

process information in a more objective, open-minded fashion that fosters uncovering the truth.

Lately, selective exposure theory has attracted increasing attention from Web researchers [10]. The major concern is the fact that Internet users, who have increasing choices and autonomy to filter out disagreeable information, coupled with the wide adoption of personalization technologies that often provide only agreeable information, may isolate themselves into information “filter bubbles” [1, 27]. To cope with this problem, researchers have been developing systems that aim at exposing users to diverse perspectives in various domains such as politics [17, 22, 25], healthcare [15], news [28] and consumer reviews [36].

These designs mainly focused on ways of categorizing and presenting diverse perspectives. In many previous studies, dichotomous valence models (conservative/liberal, pro/con, etc.) were adopted to categorize information. This approach tends to ignore magnitudes of positions, even though attitudes often distribute along a continuum rather than dichotomy, and people with moderate positions may have distinctive perspectives from those with extreme positions [8, 29]. Moreover, there is substantial evidence that lower attitude discrepancy between the communicator and the recipient may decrease cognitive dissonance [0]. Therefore differentiating sources of moderate positions from the extreme ones may reduce users’ resistance to dissonant opinions. A few previous studies provided evidence to support this idea. For example, OpinionSpace developed by Faridani et al. [5], a tool visualizing other voters’ positions, uses spatial distance to represent the magnitude of opinion differences. However the study focused on examining the overall user experience without studying how magnitude information impacted users’ information seeking process. ConsiderIt, developed by Kriplean et al. [17] aggregates public thoughts on controversial political issues, and provides users with top arguments on seven different stances distributed on a 1-support to 7-oppose scale. Qualitative results showed favorable user experience with this categorization method. However, there is still a lack of systematic examination of whether and how adding magnitude information in the categorization of positions will impact users’ information selection and consumption.

Previous studies have explored interfaces that explicitly indicate stances by presenting information in either a two-column view or with inline labels. The former often listed opposite information side by side [17, 28], while the later used highlighting [23], tag [25] and image [11], etc., to indicate positions. The position indicator we introduced here belongs to the later category.

From a cognitive perspective, researchers argue that selective exposure can be attributed to people’s tendency to preserve cognitive resource, as processing and counter-arguing challenging information demands more cognitive resources than consuming agreeable information [7]. Thus, reducing the demand of cognitive resource from other

processes, e.g., interpreting and evaluating the message position [7, 33], could reduce the effect of selective exposure. Providing explicit labels of stances seems to be consistent with this notion of easing the processes involved in attending to and interpreting information. Research also suggests that explicitly indicating the stance of information could serve as a “reminder” that there exists different opinions, which may accentuate the benefit of, as well as the social norm, that one should balance his or her information seeking of diverse opinions [22, 25]. However, previous studies yielded mixed results of the effect of stance labeling on reducing selective exposure [23, 25].

Interestingly, in the Munson et al. [23] study, after distinguishing challenge-averse users (who prefer as much agreeable information as possible) from diverse-seeking users (who prefer at least some amount of challenging information), they examined the effect of highlighting agreeable items on challenge-averse users’ satisfaction with using interfaces presenting information of varied levels of diversity. Surprisingly, overall they observed little effect of highlighting. In the condition in which the interface presented fewer agreeable items, it even slightly decreased challenge-averse users’ satisfaction, possibly because the highlighting made it more explicit that there were fewer agreeable than disagreeable items. These mixed results raised the question of whether explicit indicators of positions could have differential effects on people with high or low accuracy motives. Specifically, we hypothesize that although the indicators can ease the cognitive demands and even “remind” users who have high accuracy motives to seek balanced information, the indicators may have little, or even negative effect on people with low accuracy motives by facilitating them to identify information to avoid. We will test this hypothesis in our study.

METHODOLOGY

In this experiment, we asked participants to browse comments from users who expressed their opinions in discussions of controversial issues in an online forum. We compared the condition where participants saw position bars indicating the users’ positions on the particular issue to the control condition where the bars were hidden. We compared participants’ selection of users and comments, perceived positions of the users, agreement with the comments, and knowledge gained between the two groups.

Participants

By posting recruiting ads in email newsletters, we recruited 32 participants from a Midwest college town in USA for the lab experiment. 20 of them are undergraduate or graduate students. The rest are a mix of faculties, university staffs, engineers, etc. They were randomly assigned to the group with position indicators (group 1) and control group (group 2). No significant difference in gender (group 1 43.8% male, group 2 37.5% male), age (group 1: $M=32.0$, $SD=13.3$, group 2: $M=24.6$, $SD=8.5$; $p=0.28$), education (25% in group 1 and 31.3% in group 2 are graduate students

or have graduate degree, the rest are undergraduate or have bachelor degree) or political leaning (scale 1-conservative to 5-liberal, group 1: M=3.5, SD=0.9, group 2: M=3.5, SD=1.1; $p=0.85$) was observed between the two groups. Participants were compensated for \$8 per hour and spent 1.5-2 hours on pre-tests and the experiment.

Task

In the experiment, we presented participants an interface of “user list” introduced as the result page of “user search” function of an online discussion forum (Figure 1). We will explain in detail how the users were created in the material section. Participants were instructed to imagine that they were visiting an online forum where people discussed controversial issues to learn about topics they were interested. They would eventually write their own posts. For each topic, the user list page suggested 32 users who were active in the discussion of the topic. Participants could click on any user to read his or her comment page (Figure 2), where one comment randomly chosen from the “top 3 comments” of this user would be presented, and participant could click “read more comments from this user” to read the other two, one at a time. Under each comment, there was a question asking participants to rate their agreement with the comment using a 1-disagree to 5-agree scale. They could return to the user list page to check other users at any moment. Once they clicked return (to leave the comment page of current user), they would be prompted to rate the user’s position on the topic using a 1-pro to 5-con scale. They could click on any user and read as many comments as they wanted. When they were ready, they could proceed to the next page to write their own post. After submitting the post, they would proceed to the next task on a different topic. Each user on the experiment interface had a unique user name, which did not appear in more than one topic. The orders of topics, users, and the comments were randomized across participants.

Figure 1 presented the interface for the user list page. All 32 users were listed in this fashion, with a user name and an

opinion summary generalizing the person’s view. In the condition with position indicators, a bar showing the position of this user was added to the user profile, which indicated both the valence and the magnitude of this user’s position, i.e., there were four kinds of position bars: 1) strong pro-a full green bar; 2) moderate pro- a ¾ green and ¼ red bar; 3) moderate con- a ¾ red and ¼ green bar; 4) strong con- a full red bar. As shown in Figure 1, images of thumb-up and thumb-down were added in the bar to indicate the meaning of the colors. Participants were told in the instruction that the website allowed other users to rate the position of users who participated in the discussions, and the position bar was calculated based on these ratings. In the control condition, this position bar was hidden in both the user list page and user comment page.



Figure 1. Interface for user list page



Figure 2. Interface for user comment page

Material

In a previous study [18], we identified a number of social, political and health related topics that have a balanced distribution of supportive and opposing attitudes in the similar population. We selected 6 topics from them (see Appendix A) for the experiment. Examples are “should

Category	Opinion Summary	Comment Examples
Strong pro	Drug advertisements inform patients about medical issues therefore should be encouraged.	An important benefit of direct to consumer advertising is that it fosters an informed conversation about health, disease and treatments between patients and their health care practitioners. Pharmacy members want patients and consumers to talk to their physicians about the medicines that may help them and to fully understand the known risks regarding these medicines.
Moderate pro	I do not have problem with advertisement about prescription drug as long as it provides unbiased information.	I do not have problem with advertisement about prescription drug as long as it provides unbiased information, meaning including both its benefits and side effects in a honest manner. A prescription drug is something that consumers should be making a rational decision about. And the more information consumers have, the better decisions they can make.
Moderate con	We should ban, or at least limit the advertisement of prescription drugs to avoid people making biased medical decisions.	I think that U.S. should limit the television commercials for prescription medications for this may influence the doctors and make a medication "more popular" without proper reason. Excessive promoting of medications using commercials may influence the "popularity" of a certain medication. It is a bad thing because people would more frequently ask for it.
Strong con	Drug advertisement is dangerous and should be banned as it may mislead patients.	Direct to consumer prescription drug ads, like most advertisements, are intended to sell the product being advertised. Such ads use marketing tactics that manipulate, create false impressions, and therefore mislead consumers instead of educating them about the drugs. It can be dangerous that patients start self-diagnosis by the information they get from advertisement.

Table 1. Example of Comments and Opinion Summary

social security be privatized?” and “should prescription drugs be allowed to directly advertise to consumers?”.

In the experiment, we created fictional users to control for the consistency of information and distribution of positions. To be ecologically valid, we obtained the material from real user posts on the Internet. We started by collecting 150-200 comments, with 60-100 words each, expressing varied opinions for each of the topics, from the online forums of debate.org, procon.org and Yahoo! Answers¹. One graduate student and two undergraduate students worked on categorizing these comments into four categories: strong pro, moderate pro, moderate con and strong con (see Table 1 for examples). The inter-rater reliability (Fleiss Kappa) reached 0.85, which is often considered good agreement [6]. We categorized the comments with the following features to be moderate: 1) expressing indecisiveness by mentioning merits of both sides, but leaning towards one side; 2) supporting one side but with certain condition (e.g., “it is good if certain restriction applied”); 3) modest and/or uncertain tone (e.g., “I guess”, “maybe”). In contrast, strong comments expressed one-sided and assured opinions and often with confident and strong tone (e.g., “I believe”, “definitely”).

For each topic, we created 32 users, with 8 under each of the four categories: strong pro, moderate pro, moderate con and strong con. For each user, we selected 3 comments from the corresponding comments pool to be his or her “top comments”. As a priority, we grouped comments collected from the same real users together. Otherwise, we carefully selected and modified the comments minimally, if necessary, to make them reasonably consistent for a particular user. For each user, we created an “opinion summary” (see Table 1) by either choosing the first sentence of one of the comments or writing by ourselves.

Measurement

In this section, we will only introduce variables that were measured directly by questionnaires or the experiment platform. More indexes we defined for analyzing specific RQs will be introduced in the result section. A summary of all the variables used in the study can be found in Table 2.

Participants completed a demographic information questionnaire and a topic related questionnaire before the experiment. In the later, three attributes were measured for each topic: prior attitude, accuracy motive and knowledge.

Prior Attitude Index

Following [13], we used a 5-item semantic differential scale to measure participants’ attitude on each topic. For example, when measuring participants’ attitudes on vegetarianism, instead of directly asking whether they held positive or negative attitudes, we asked them to choose their opinions on vegetarianism based a 7-point Likert scale for

five pairs of bipolar adjectives: unfavorable-favorable, unhealthy-healthy, bad-good, unnecessary-necessary, harmful-beneficial. We calculated the average ratings of the five items to be the participant’s *prior attitude index* for the topic (Cronbach’s $\alpha=0.96$, good internal consistency [4]).

Accuracy Motive Index

Based on previous research [18], accuracy motive was measured by two items: 1) *how much are you interested in learning more about the topic*, and 2) *how much do you desire to know the truth of the topic regardless of your own position*. The answers were based on a 1-none to 7-a lot scale and we averaged the ratings to be the *accuracy motive index* (Cronbach’s $\alpha = 0.88$, good internal consistency).

Knowledge

For each topic, we asked participant to write down “what reasons or arguments immediately come to your mind for people who support or oppose the issue” using simple sentences such as “religious reason”. They were asked to write pro-arguments and con-arguments in two separate bulleted lists. The same questions were asked again after the experiment. We counted the number of points that appeared in the post-experiment but not in the pre-experiment questionnaire as a proxy of participants’ *knowledge gain* after browsing the website.

User and Comment Selection

One main focus of this study was participants’ selection of users and comments of varied positions. The experiment platform, i.e., the website we used, was able to record which users and comments were clicked on.

RESULT

Participants in the condition with position bars selected a mean of 5.1(SD=2.7) users and read 1.8 comments (SD=0.8) from each user for each topic. Participants in the control condition selected a mean of 5.3(SD=2.7) users and read 1.7 (SD=0.8) comments from each user. No significant difference in the number of users ($p=0.69$) or comments per user ($p=0.60$) was observed between the two groups.

We were interested in the effect of the source position indicator on both the selection and processing of diverse information. To that end, we first investigated participants’ selections of users of varied positions (**RQ1**), specifically how the valence and magnitude information showed by the indicator influenced the selection (**RQ2**). We then analyzed whether the position indicator had impact on participants’ perception of the users’ position *after* reading the comments (**RQ3**), and their agreement with the comments (**RQ4**). Lastly, we analyzed whether the selection of different users led to differences in participants’ knowledge gained on either side (**RQ5**).

Prior Attitude Side and Relative User Position

¹ www.debate.org, www.procon.org, answers.yahoo.com

Index	Definition	Calculation
<i>Prior Attitude Index</i>	Reflects the participant’s prior attitude on the topic based on 1-con to 7-pro scale.	Mean of ratings of the 5-item semantic differential scale on the particular topic measured in the pretest.
<i>Accuracy Motive Index</i>	Reflects the motivation the participant has for accurately learning about the topic.	Mean of the ratings of the 2-item accuracy motive scale on the particular topic measured in the pretest.
<i>Prior Attitude Side</i>	Reflects whether the participant has pro or con prior attitude on the topic.	Pro, if <i>Prior Attitude Index</i> >4 Con, if <i>Prior Attitude Index</i> <4
<i>Attitude Extremity Index</i>	Reflects extremity of prior attitude	<i>Prior Attitude Index</i> -4
<i>Relative User Position</i>	Reflects user’s position as relative to participant’s side.	If <i>Prior Attitude Side</i> =Pro, extremely/moderately pro (con) users coded as <i>extremely/moderately consistent, (inconsistent)</i> If <i>Prior Attitude Side</i> =Con, extremely/moderately pro (con) users coded as <i>extremely/moderately inconsistent (consistent)</i>
<i>Selective Exposure Index</i>	Measures user’s preferential selection between attitude consistent and inconsistent users (or comments).	N (extremely consistent users) + N (moderately consistent users) –N (extremely inconsistent users) – N (moderately inconsistent users)
<i>Position Judgment Index</i>	Measures participant’s position judgment of the particular type of users.	Mean of position ratings given to all users of the particular type
<i>Relative Position Index</i>	Measures participant’s perception of the position of the type of users (extremely/moderately consistent or inconsistent) as relative to their own side	If <i>Prior Attitude Side</i> =Pro, <i>Relative Position Index</i> = <i>Position Index</i> If <i>Prior Attitude Side</i> =Con, <i>Relative Position Index</i> = 6- <i>Position Index</i>
<i>Agreement Index</i>	Measures participant’s agreement with comments from the type of users	Mean of agreement ratings given to all comments from the type of users
<i>Knowledge Gain Index</i>	Reflects the new knowledge of arguments on either side gained by using the system	Number of points appeared in post-experiment recall test but not pre-experiment recall test listed on either side

Table 2. Summary of the measurements used in the current study.

First we coded each participant’s prior attitude on each topic to be *pro* or *con* based on whether his or her *prior attitude index* is below or above 4, which is the neutral point of the scale and also the median among all participants. 20 out of the 192 cases where participants scored exactly 4 were removed from the analysis. Then we coded each user case to be *extremely consistent, moderately consistent, moderately inconsistent* or *extremely inconsistent* based on whether the user is on the same or opposite side of the participant’s prior attitude, and their attitude magnitude (see Table 2 *Prior Attitude Side* and *Relative User Position*).

User Selection

Selective Exposure Index

To answer RQ1, how the position indicator impacted participants’ selective exposure tendency, we created *selective exposure index*, calculated by the difference in the number of attitude consistent and attitude inconsistent users selected (see Table 2 *Selective Exposure Index*). A positive index indicated the exhibition of selective exposure, while negative indicated selection of more attitude inconsistent users, and its magnitude reflected the size of the difference.

Overall Selectivity of All Participants (RQ1a)

We started by examining the indicator’s overall impact on participants’ selective exposure tendency. We ran a mixed-

effect regression model on the *selective exposure index* by having the presence position bars (present=1, absent=0) as the fixed-effect independent variable (N=172). In all the mixed-effect regression models described in this paper, participants were included as random effects. We will only mention variables included as fixed effects in the remaining of the paper. We found a main effect of the presence of the position bars ($\beta=-0.51, t(30)=-2.20, p=0.04, f^2=0.03$), suggesting that, overall, position bar decreased participants’ selective exposure.

Selectivity of Participants with High Versus Low Accuracy Motives (RQ1b)

In RQ1b, we asked whether the position indicator would have differential impacts on users with varied levels of motivation to accurately learn about the topic (*accuracy motive*). To answer this question, we further examined the result by including the *accuracy motive index* and presence of position bars in the mixed-effect regression model on *selective exposure index* (detailed definition can be found in Table 2). Indeed, we found a significant interaction between *accuracy motive* and presence of position bars on the *selective exposure index* ($\beta=-0.37, t(28)=-2.22, p=0.03, f^2=0.09$). The significant interaction answered our research question in the affirmative, i.e., the effect of the position indicator on decreasing selective exposure depended on

whether the participant had a high or low motivation to learn about the topic.

To more clearly visualize this two-way interaction, we chose to use median splits to show how the effect of the position bar depended on the level of accuracy motive. We should point out that the median splits were performed *only* for the visualization. The *accuracy motive index* was treated as a continuous variable in *all* regression models. In figure 3, we first preformed median splits to create the high or low accuracy motive groups, and we presented the mean *selective exposure index* for each group. This visualization highlighted the two-way interaction between accuracy motive and presence of position bars: the presence of position bars decreased the selective exposure for topics in which participants had higher than average accuracy motive, as they selected more attitude inconsistent users. However, the bars had little impact on participants who had lower than average accuracy motive. The results suggested that the position bar was effective in mitigating the selective exposure tendency, but *only* for participants who had stronger motivation to learn about the topic.




Figure 3. Average *selective exposure index*

To confirm this effect, we further investigated whether the same pattern of result was observed in the selection of the individual comments from the users. Since participants had the option to read from 1 to 3 comments from each user, we calculated the *selective exposure index* based on the number of comments they read as well. We observed the same trend that there was a significant interaction between *accuracy motive* and presence of position bars ($\beta=-0.68$, $t(28)=-2.20$, $p=0.03$, $f^2=0.10$), indicating that with the presence of position bars participants were more likely to read comments of the opposite side for topics they had relatively high accuracy motives than those without the position bars.

To summarize the answer for RQ1, we found that the position bar had different impact on participants with varied levels of accuracy motives. Specifically, it increased their exposure to attitude-challenging information when participants *had high motivation to accurately learn about the topic*. However, for participants who had little interest in learning about the topic, the bar had little impact.

Valence and Magnitude of Source Position on Participants' Selectivity (RQ2)

Since the position bar not only provided the valence (pro/con) but also the magnitude (moderate/extreme) of users' stances, RQ2 asked whether participants differentially selected users indicated as one of the four *relative user positions*: *extremely consistent*, *moderately consistent*, *moderately inconsistent* and *extremely inconsistent*, coded as *Relative User Position* described in Table 2. We included *accuracy motive index*, presence of position bars (present=1, absent=0), the selected users' position valences (consistent=1, inconsistent=0) and magnitudes (moderate=0, extreme=1) in the mixed-effect linear regression model as independent variables to test how they impacted the percentage of the type of users selected (N=668, 4 categories for each of the 172 tasks). We found a significant three-way interaction among the presence of position bars, *accuracy motive*, and the selected users' valences of positions, ($\beta=-0.08$, $t(16)=-2.95$, $p<0.01$, $f^2=0.04$), and a significant three-way interaction among the position bars, *accuracy motive*, and the selected users' magnitudes of positions ($\beta=-0.06$, $t(16)=-2.07$, $p=0.04$, $f^2=0.07$). These significant interactions provided answers to RQ2 in the affirmative: both the valence and magnitude reflected in the position indicator affected participants' selection of users, and the selection differed between participants with varied levels of accuracy motives.

Given that we were interested in how the position bars would encourage exposure to attitude-challenging positions, we separated the selected users into two groups: *attitude consistent* users and *inconsistent* users. We then ran the regression analysis in each group to study how the interactive effects of the presence of position bars and *accuracy motive* would lead to different patterns of selection of users with extreme or moderate positions (i.e., position magnitudes).

For attitude inconsistent users, we ran a mixed-effect regression model on the selection percentages by including the presence of position bars, *accuracy motive index*, and the user's position magnitudes as independent variables. We found a significant three-way interaction among all the independent variables ($\beta=-0.06$, $t(24)=-2.18$, $p=0.03$, $f^2=0.08$). It suggested that the position bar had differential effects on participants who had high and low accuracy motives on selecting inconsistent users with different attitude magnitudes. Figure 4 provided a clearer picture- *the position bars encouraged participants who had higher accuracy motives to select more moderately inconsistent users than those with lower accuracy motives*. No such difference was observed when the position bars were absent. To verify the conclusion, we ran a mixed-effect regression model on the selection percentages for *moderately inconsistent* users only by including *accuracy motive index* and presence of position bar as independent variables. Results showed a significant two-way interaction between the independent variables on the selection percentages ($\beta=0.05$, $t(28)=2.72$, $p<0.01$, $f^2=0.10$). For extremely inconsistent users, we did not

observe any statistically significant effect of the presence of position bars or *accuracy motives*. These results therefore provided further support for this conclusion.




Figure 4. Selection percentage of attitude inconsistent users

For attitude consistent users, we ran the same mixed-effect regression model but did not find the three-way interaction to be significant. However, we found a significant two-way interaction between presence of position bars and *accuracy motives* ($\beta=-0.03$, $t(27)=-1.93$, $p=0.05$, $f^2=0.03$), and a main effect of users' position magnitudes ($\beta=0.09$, $t(27)=4.52$, $p<0.01$, $f^2=0.06$). The two-way interaction suggested that the position bars led to lower selection percentages only for participants with high accuracy motives, but not for those with low accuracy motives. This was consistent with the conclusion from the analysis of the *selective exposure index*: the position bar led participants with high accuracy motives, but not those with low accuracy motives, to select lower percentages of users with consistent positions. The main effect of position magnitude was caused by the fact that, regardless of the accuracy motive or presence of position bars, participants consistently selected more users with extremely consistent than moderately consistent positions. We speculated that it was caused by the fact that attitude consistent information would not lead to cognitive dissonance, thus neither the accuracy motives nor the presence of position bar impacted how likely participants selected users with moderately consistent or extremely consistent positions.

Interestingly, in the regression model for all users, the three-way interaction among users' attitude magnitudes, presence of position bars, and accuracy motives on the selection percentage implied that the position bar made participants who had low accuracy motives to be more selective towards users with extreme positions rather than moderate positions. By analyzing the total number of users selected, we found the accuracy motive had a marginally significant positive effect ($\beta=0.22$, $t(30)=1.69$, $p=0.08$, $f^2=0.03$), showing that participants overall selected fewer users when they had relatively low accuracy motives. Given the low accuracy motives, it is likely that they were only willing to spend limited effort to understand the topic and may not be interested in knowing and weighing the detailed arguments, and therefore strategically chose to focus mostly on those standing at the extreme ends. In contrast, the greater motivation to form an accurate understanding of the topic might intrigue participants to seek for more detailed

arguments from users of varied positions, including those with moderately inconsistent attitudes.

At the comment level, we performed the same analysis with the percentages of the selection of comments as the dependent variable in the regression model. We found a significant four-way interaction among presence of position bars, *accuracy motive*, users' valences and magnitudes of positions ($\beta=0.08$, $t(16)=1.93$, $p=0.05$, $f^2=0.08$). These results again suggested that the position bar had differential effects on participants with high and low accuracy motives on their tendency to read comments from different types of users. Closer examination revealed similar patterns as in the selection of users: with the position bars, participants with high accuracy motives read lower percentages of attitude consistent comments, but higher percentages of moderately inconsistent comments than participants with low accuracy motives, which resulted in overall lower selective exposure tendency of the former than the later.

In summary, our results provided affirmative answers to RQ2 by showing that both valence and magnitude information showed in the source position indicators had differential effects on users who had varied levels of accuracy motives. The most important result was that *with the position bars, participants with high accuracy motives selected more moderately inconsistent users than those with low accuracy motives*. As a result, the position indicators increased the exposure to attitude-challenging information for participants with high accuracy motives, but had no such impact for participants who had low accuracy motives.

Position Judgment

After examining the effect of the position bars on participants' selective exposure, the next question is to what extent the position bars could lead to changes in their *reception* of attitude inconsistent information. This leads to our RQ3: whether the source position indicator had influence on participants' perception of users' positions *after* they read their comments. To answer this question, we examined participants' position ratings (on a 1-pro to 5-con scale) given to each user after they read the comments. First, we looked at the correlation between participant's judgments and the "ground truth" (i.e., which we manipulated and checked in a pilot study, and coded as extremely pro user=1, moderately pro user=2, moderately con user=3, extremely con user=4). We found that the position bars significantly increased this correlation from 0.64 to 0.78 ($Z=4.77$, $p<0.01$). It indicated that after seeing the position bars, participants judgments of users' stances were closer to those indicated by the position bars, suggesting that the bar did influence their judgment of the positions of the users.

Relative Position Index

Another possible influence of the position indicator on participants' information reception is the extent to which a user was perceived to be different from one's pre-existing attitude. For example, if an attitude-inconsistent opinion

had low reception, it might be judged to be more extreme than information the user had high reception with. To examine such judgment, we created an index called *relative position index* to measure participant's judgment on user's positions relative to the side of his or her pre-existing attitude. To do so, we first calculated *position judgment index* by the average ratings given to the four categories of users (extremely/moderately consistent or inconsistent). Since the position judgment used a 1-pro to 5-con scale, if the participant had positive prior attitude, the *relative position index* of each type of users would be equal the *position judgment index*. If the person had negative prior attitude, the *relative position index* would be calculated by 6 minus *position judgment index*. *Relative position index* reflected how far the participant perceived the type of users to be different from his or her prior attitude, with a higher number indicating that the users were perceived to be more different (see Table 2 *Position Judgment Index* and *Relative Position Judgment Index*).

Participants' Judgments of Users' Positions (RQ3)

We ran a mixed-effect regression model on the *relative position index* (N=535, 153 cases where the corresponding type of users were not selected were removed) by including the presence of position bars, user's position valences and magnitudes as independent variables. We found a significant three-way interaction among all the independent variables ($\beta = -0.78$, $t(24)=-2.46$, $p=0.01$, $f^2=1.50$). All the other two-way and main effects, except for that between presence of position bars and users' position valences, were also significant. It suggested that the position bar had differential impact on participants' perceived relative positions of the users (see Figure 5). The Cohen's f^2 indicated that the interaction had a large effect size [3].

Given that we were interested in the difference in reception between attitude-consistent and inconsistent information, we unpacked the three-way interaction by running mixed-effect linear regression models on the *relative position index* among attitude consistent users and inconsistent users separately. We found a significant two-way interaction between users' opinion magnitudes and presence of position bars in the attitude inconsistent group ($\beta = 0.59$, $t(28)=3.10$, $p<0.01$, $f^2=0.25$), but not in the attitude consistent group ($\beta = -0.21$, $t(28)=-0.98$, $p=0.33$). The results confirmed that the three-way interaction was caused by the difference in the reception of attitude-consistent and inconsistent information. In other words, *the position bars only led to significant differences in the position judgments of users with inconsistent positions*. Figure 5 illustrated this result: with the help of position bars, participants were better at differentiating between extremely and moderately inconsistent users. In fact, by examining moderately inconsistent users only, we found a significant main effect of the presence of position bars on the *relative position index* ($\beta = -0.33$, $t(28)=-2.00$, $p=0.05$, $f^2=0.05$), suggesting

that participants tended to perceive this group of users to be closer to themselves when the position bars were present than absent. This is consistent with the notion that people are inclined to perceive others of different opinions to be more extreme than they actually are. The position bar seemed to help to moderate their perception of moderately inconsistent sources to be less extreme.




Figure 5. Average *relative position index*

Agreement

To further study how the position indicator influenced reception of diverse opinions, RQ4 asked whether the position bars also influenced participants' agreement with the comments from different types of users. To answer this question, we investigated participants' agreement ratings given to the comments (by a question under each user comment on the experiment interface). We expected that providing agreement rating would encourage more reflection of personal value than position judgment, for the later of which one could adopt a more third-person position. We therefore expected that agreement ratings would be impacted by the extremity of prior attitudes, as participants with more extreme prior attitudes would more likely have stronger disagreement with users who take opposite positions. Moreover, intuitively, if the participant was only moderately leaning towards one side, he or she might agree with users who were also in moderate positions more than those with extreme positions. Therefore, we examined participants' agreement by taking their prior attitude extremity into consideration.

Agreement Index and Prior Attitude Extremity Index

Agreement index was calculated by the average agreement ratings participants gave to all the comments they read and rated for the corresponding type of users for each topic (Table 2 *Agreement Index*).

Prior attitude extremity index was defined as the distance of one's *prior attitude index* from the neutral point, and was calculated by the absolute value of the *prior attitude index* minus neutral value 4. For example, if a participant scored 5 in the *prior attitude index*, he or she would be coded as *con* prior attitude with *prior attitude extremity index* of 1 (Table 2 *Prior Attitude Extremity Index*).

Participants' Agreement with Comments from Users of Varied Positions (RQ4)

We ran a mixed-effect regression model on the *agreement index* (N=535), with the presence of position bars, *prior attitude extremity*, users' position valences and magnitudes as independent variables. We found a significant three-way interaction among presence of position bars, *prior attitude extremity index*, and users' position valences ($\beta = -0.86$, $t(16)=-3.65$, $p<0.01$, $f^2=0.36$), and a significant three-way interaction among presence of position bars, *prior attitude extremity*, and users' position magnitudes ($\beta = -0.47$, $t(16)=-2.12$, $p=0.04$, $f^2=0.10$). The results suggested that the position bars differentially influenced users with moderate and extreme prior attitudes on their agreement with the comments from users of different positions.

To understand whether the position bars differentially impacted agreement with users of different positions, we separated the consistent and inconsistent users into two groups. For attitude inconsistent users, we ran a mixed-effect regression model on *agreement index* by including presence of position bars, *prior attitude extremity index* and users' position magnitudes as independent variables, and found a significant three-way interaction among them ($\beta=-0.41$, $t(24)=-1.91$, $p=0.05$, $f^2=0.35$). The two-way interaction between *prior attitude extremity* and presence of position bars, and the main effect of *prior attitude extremity* were also significant.

Figure 6² explained the three-way interaction we found above: while the position bars did not have a significant impact on participants with moderate prior attitudes, they increased the agreement with moderately inconsistent users for participants with extreme prior attitudes, who tended to have lower agreement with attitude inconsistent comments than participants with moderate prior attitudes. To further verify this conclusion, we looked at the *agreement index* for moderately inconsistent users only. Indeed, we found the two-way interaction between *prior attitude extremity* and presence of position bars to be significant ($\beta=0.65$, $t(28)=3.73$, $p<0.01$, $f^2=0.19$). It implied that, by helping participants to differentiate moderately inconsistent users from extremely inconsistent ones, participants who had extreme prior attitudes more likely agree with users with moderately inconsistent positions when the position bars were present than absent. For *agreement index* for extremely inconsistent users, we only observed significant main effect of *prior attitude extremity* ($\beta=-0.46$, $t(31)=-5.89$, $p<0.01$, $f^2=0.25$), suggesting that the position bars did not have the same effect on agreement with users with extremely inconsistent positions. As expected, participants with extreme prior attitudes tended to disagree more with extremely inconsistent users than those with moderate prior attitudes regardless of the presence of position bars.

² As in previous figures, median splits on *prior attitude extremity* were used to illustrate the interaction in the figure, but not in the regression models.

For attitude consistent users, we ran the same mixed-effect regression model and did not observe the three-way interaction, but found a significant two-way interaction between participants' *prior attitude extremity* and users' position magnitudes ($\beta = 0.31$, $t(28)=2.89$, $p<0.01$, $f^2=0.07$), indicating that regardless of the presence of position bars, participants with more extreme prior attitudes tended to agree more with users on the same side with extreme positions, while participants who had moderate prior attitudes tended to agree with users on the same side with moderate attitudes. This is consistent with the intuition that people tend to agree with others of similar positions.




Figure 6. Average *agreement index* attitude inconsistent users

The analysis performed on position judgment and agreement revealed a consistent theme: the indicator showing both source position valence and magnitude was especially helpful for participants to differentiate between those who had moderately opposite stances from those with more extreme ones. Ultimately, the indicator seemed to help moderate participants' tendency to over-estimate the discrepancy of positions between themselves and the attitude-inconsistent information source, and induce more agreement with sources with moderately opposite positions. In general, the results showed that not only did the position bars reduce selective exposure, but also increased reception of attitude-challenging information.

Knowledge Gain (RQ5)

One important goal of exposing people to diverse opinions is to promote awareness of arguments on both sides in order to help them make informed decisions or form unbiased views. Last but not the least, we analyzed participants' knowledge gained on both sides as the outcome measurement. Specifically, we were interested in whether the differentiation in selection of information led to the difference in knowledge gained.

In the survey before experiment, we asked participants to list points that could "immediately come to mind" to support or oppose the topic. After the experiment, we asked them to do the same. We compared the two answers and counted how many points of attitude-consistent argument and inconsistent argument, respectively, appeared in the post-experiment questionnaire but not in the pre-experiment questionnaire. We used this number as a proxy measure of

knowledge gain (see Table 2 *Knowledge Gain Index*) from using the system.

We ran a mixed-effect regression model on the *knowledge gain index* for the side opposite to participants' prior attitude (*attitude inconsistent knowledge gain*) by including presence of position bars and *accuracy motive index* as independent variables, and found a significant interaction between the two ($\beta=0.33$, $t(28)=3.10$, $p<0.01$, $f^2=0.04$). Figure 7 illustrates the two-way interaction after we performed median splits of *accuracy motive index*: with the help of position bars, participants gained more knowledge of opposite side for topics they had high accuracy motives, but not for those who had low accuracy motives. For knowledge gained on arguments that people of similar attitudes make, we didn't observe any significant effect of position bars or accuracy motives. The result was consistent with the finding that participants, with the help of position bars, were exposed to attitude inconsistent users more for topics they had high accuracy motive. The results therefore further supported the conclusion that providing indicators of information source position is beneficial for users who are interested in accurately learning about the topic.




Figure 7. Average attitude inconsistent *knowledge gain*

DISCUSSION

Scholars have warned that the “echo chamber effect” on the Internet can exacerbate attitude polarization, as people may end up only interacting with others who share similar attitudes or opinions. In this experiment, in the hope of mitigating the problem, we introduced and tested the effects of an interface feature that indicated the information source’s (in this study, the user) position, including both its valence and magnitude, on reducing the potential echo chamber effect. We found that it had differential effects on participants who had varied levels of motivation to accurately learn about the topic. Specifically, the indicator significantly decreased the selective exposure of participants who had high accuracy motives but had no such effect on participants who had low accuracy motives. The difference was caused by the fact that the multi-level position indicator encouraged participants who had high accuracy motives to select users with moderately inconsistent attitudes, but discouraged participants who had low accuracy motives from selecting information sources that were more neutral.

At the information reception stage, the indicator had influence on participants' perception of users and evaluation of information. Specifically, it helped participants to differentiate users who took moderately inconsistent stances from those who took extreme ones. As a result, participants eventually had a lower disagreement with comments from these moderately inconsistent users. Moreover, we found that, by encouraging exposure to attitude-challenging information, the indicator led to greater knowledge gained on the opposite side of their original positions for participants who had high accuracy motives.

The study has two main contributions. First, the results suggest that providing not only the valence but also the magnitude of information source position is useful for encouraging exposure to diverse information, which may facilitates the seeking of common grounds among people with different perspectives. Based on this result, we speculate that the group of people who are only moderately leaning towards the opposite side are critical for bringing in different perspectives without invoking strong resistance arising from cognitive dissonance. It is also possible that these people may provide distinctive insights that are useful for others to understand arguments from both sides. We believe that in many situations, identifying and explicitly indicating sources or arguments with moderate positions may allow them to function as “bridges” that increase the cohesion of groups with diverse opinions and contribute to more efficient “marketplaces” for idea exchanges.

Second, our results echoed suggestions from previous research [10, 23] that interfaces should be tailored for users with different levels of accuracy motives to achieve the goal of exposing users to challenging views, which highlighted the importance of personalization when designing for seeking diverse opinions. Our results showed that, when users are motivated to accurately learn about a topic thus interested in getting insights from people of varied positions, a non-dichotomous, more elaborate categorization of information combined with explicit indicators of such categories is useful.

Previous research suggests that high accuracy motive is often related to prioritizing the utility of information [34], and/or the social norm of valuing diversity [22]. By explicitly labeling the stances, it may better support the utility goal of seeking new, useful information. In addition, it may serve as a “reminder” that there exist diverse opinions, so as to further reinforce the normative goal of diversity seeking. However, when users have little interest in accurately learning about the topic, providing explicit labels of positions may sometimes exacerbate their biased selectivity between different stances.

From a design perspective, two issues emerged from this set of conclusions: the importance of inferring users' accuracy motivation levels and inferring the information sources' stances. Learning from research on personalization technologies, users' accuracy motives can be identified

from either explicit inquiry or implicit inference based on their user profiles or previous behavior [21]. There are also conditions in which people are in general more motivated to accurately learn about a topic. For example, decision-making support tools will more likely target users who have high accuracy motives as the decisions are often related to important outcomes. Moreover, we argue that it is possible for designers to actively increase users' accuracy motive by, for example, providing simple interface cues that emphasize the social norm of valuing diversity [22] or highlight the utility of diverse or challenging information [34]. To infer information sources' stances, the literature has shown various ways of classifying, e.g., political positions or consumer opinions, based on either text based methods (e.g., machine learning, opinion mining, sentiment analysis [26]) or non-text features such as user voting or social network [37]. One future direction would be to develop methods that can accurately identify the "bridges" in an online opinion space, or information sources that can provide insights on both sides.

While much previous research alarmed the problem of selective exposure on the Internet, other questioned its existence as a universal phenomenon [23, 31]. Our study generated consistent results with the later, suggesting that there are important moderating variables that influence the exhibition of such behavior. In addition to exploring ways to mitigate selective exposure, future research should also focus on identifying and understanding these moderating variables, and examining their interactive effects with interface features.

In this study, given the limited time of use, we could not conclude whether participants became less polarized after using the system. Future work should focus on longer-term field studies, in which effects of information aggregator showing source positions on users' information seeking behavior as well as attitude change can be studied.

REFERENCES

1. Van Alstyne, M., & Brynjolfsson, E. (1996). Electronic communities: Global village or cyberbalkans?. In Proceedings of the 17th International Conference on Information Systems.
2. Aronson, E., Turner, J. A., & Carlsmith, J. M. (1963). Communicator credibility and communication discrepancy as determinants of opinion change. *The Journal of Abnormal and Social Psychology*, 67(1), 31.
3. Cohen, J. (1992). A power primer. *Psychological bulletin*, 112(1), 155.
4. Cronbach, L.J., Coefficient Alpha and The Internal Structure of Tests. *Psychometrika* 16, 3 (1951), 297-334.
5. Faridani, S., Bitton, E., Ryokai, K., & Goldberg, K. (2010). Opinion space: a scalable tool for browsing online comments. In Proc. CHI2010, 1175-1184.
6. Fleiss, L. (2011). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76 (5), 378.
7. Fischer, P. (2011). Selective Exposure, Decision Uncertainty, and Cognitive Economy: A New Theoretical Perspective on Confirmatory Information Search. *Social and Personality Psychology Compass*, 5(10), 751-762.
8. Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 3(4), 552.
9. Frey, D. Recent Research on Selective Exposure to Information. *Advances in Experimental Social Psychology* 19(1986), 41-80.
10. Garrett, R., Resnick, P. Resisting Political Fragmentation on the Internet. *Daedalus* 140, 4 (2011) 108-120.
11. Garrett, R. K., Stroud, N. (2012). Decoupling selective approach and selective avoidance. Paper to be presented at the National Communication Association, Political Communication Division, Orlando, FL.
12. Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin*, 135(4), 555.
13. Himmelfarb, S. The measurement of attitudes. In A.H. Eagly & S. Chaiken (Eds.), *Psychology of Attitudes* (1993), 23-88. Thomson/Wadsworth
14. Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment* (p. 81). New York: Free Press.
15. Jiang, Y., Liao, Q., Cheng, Q., Berlin, R., Schatz, B. Designing and Evaluating a Clustering System for Organizing and Integrating Patient Drug Outcomes in Personal Health Messages. In Proc. AMIA 2012
16. Klayman, J. (1995). Varieties of confirmation bias. *Psychology of learning and motivation*, 32, 385-418.
17. Kriplean, T., Morgan, J., Freelon, D., Borning, A., Bennett, L. Supporting Reflective Public Thought with ConsiderIt. In Proc. CSCW2012, 265-274
18. Liao, Q. V., & Fu, W. T. (2013). Beyond the filter bubble: interactive effects of perceived threat and topic involvement on selective exposure to information. In Proc. CHI2013, 2359-2368.
19. Maass, A., & Clark, R. D. (1984). Hidden impact of minorities: Fifteen years of minority influence research. *Psychological Bulletin*, 95(3), 428-450.
20. Mendelberg, T. (2002). The deliberative citizen: Theory and evidence. *Political decision making, deliberation and participation*, 6(1), 151-193.
21. Mulvanna, M. D., Anand, S. S., & Büchner, A. G. (2000). Personalization on the Net using Web mining: introduction. *Communications of the ACM*, 43(8), 122-125.

22. Munson, S. A., Lee, S. Y., & Resnick, P. (2013). Encouraging Reading of Diverse Political Viewpoints with a Browser Widget. *In Proc. ICWSM 2013*
23. Munson, S. & Resnick P. Presenting Diverse Political Opinions: How and How Much. *In Proc. CHI2010*, 1457-1466.
24. Nemeth, C. J., & Kwan, J. L. (1987). Minority influence, divergent thinking and detection of correct solutions. *Journal of Applied Social Psychology*, 17(9), 788-799.
25. Oh, A., Lee, H., Kim, Y. User Evaluation of a System for Classifying and Displaying Political Viewpoints of Weblogs. *In Proc. ICWSM 2009*, 282-285
26. Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2), 1-135.
27. Pariser, E. *The Filter Bubble: What the Internet Is Hiding from You*, Penguin Press, 2011
28. Park, S., Kang, S., Chung, S., Song, J. NewsCube: Delivering Multiple Aspects of News to Mitigate Media Bias. *In Proc. CHI2009*, 443-452
29. Petty, R. E., & Krosnick, J. A. (Eds.). (1995). *Attitude strength: Antecedents and consequences* (Vol. 4). Mahwah, NJ: Lawrence Erlbaum Associates
30. Sanders, GS; Mullen, B. 1982. Accuracy in perceptions of consensus: Differential tendencies of people with majority and minority positions. *Eur J Soc Psychol* 13(1): 57-70.
31. Sears, D. O., & Freedman, J. L. (1967). Selective exposure to information: A critical review. *Public Opinion Quarterly*, 31(2), 194-213.
32. Sunstein, C. R. (2009). *Going to Extremes: How Like Minds Unite and Divide: How Like Minds Unite and Divide*. Oxford University Press, USA.
33. Sweeny, K., Melnyk, D., Miller, W., & Shepperd, J. A. (2010). Information avoidance: Who, what, when, and why. *Review of general psychology*, 14(4), 340-353
34. Valentino, N. A., Banks, A. J., Hutchings, V. L., & Davis, A. K. (2009). Selective exposure in the Internet age: The interaction between anxiety and information utility. *Political Psychology*, 30(4), 591-613.
35. Vydiswaran, V., Zhai, C., Roth, D., Pirolli, P. Unbiased Learning of Controversial Topics. *In Proc. ASIST2012*
36. Yatani, K., Novati M., Trusty, A., Truong, K. Review Spotlight: A User Interface for Summarizing User Generated Reviews Using Adjective-Noun Word Pairs. *In Proc. CHI2011*, 1541-1550
37. Zhou, D. X., Resnick, P., & Mei, Q. (2011, May). Classifying the political leaning of news articles and users from user votes. *In Proc. ICWSM2011*, 417-424.

APPENDIX A: TOPICS USED IN THE EXPERIMENT

1. Should death penalty be allowed?
2. Should prescription drugs be advertised directly to consumers?
3. Should euthanasia be legal?
4. Do violent video games contribute to the increase of youth violence?
5. Should people become vegetarian?
6. Should social security be privatized?