

Encouraging Reading of Diverse Political Viewpoints with a Browser Widget

Sean A. Munson¹, Stephanie Y. Lee², Paul Resnick³

¹Human Centered Design & Engineering, ²Sociology
DUB Group, University of Washington
{ smunson, syl3 }@uw.edu

³School of Information
University of Michigan
presnick@umich.edu

Abstract

The Internet gives individuals more choice in political news and information sources and more tools to filter out disagreeable information. Citing the preference described by selective exposure theory — people prefer information that supports their beliefs and avoid counter-attitudinal information — observers warn that people may use these tools to access only agreeable information and thus live in ideological echo chambers. We report on a field deployment of a browser extension that showed users feedback about the political lean of their weekly and all time reading behaviors. Compared to a control group, showing feedback led to a modest move toward balanced exposure, corresponding to 1-2 visits per week to ideologically opposing sites or 5-10 additional visits per week to centrist sites.

Introduction

Diversity in the information people see and attend to can help people and groups make better decisions, can help people learn and correct inaccurate beliefs, and can help people see ideas with which they do not agree as legitimate. Media policy in the United States has long promoted audiences' exposure to diverse information. The rationale for this goal has been that accurate beliefs and perceptions are necessary for good decision-making and for good governance in democratic society, and that development of such beliefs requires some degree of exposure to information that challenges one's existing beliefs and opinions (Frey 1986, Hart 2009).

The Internet has brought more choice for what news and information individuals can access. Observers have warned that existing media policies are ill-suited for guaranteeing exposure to political diversity in the Internet age, as people are able to choose more freely from an ever-increasing variety of sources, many of which cater to and present a very narrow range of viewpoints. Given this range of

choice, they argue, Americans will increasingly live in ideological echo chambers and polarization of different political groups will increase (Prior 2007, Sunstein 2001). Republicans and Democrats read different newspapers and watch different news stations (Iyengar et al. 2009, Morris 2007). If people prefer to avoid hearing challenging views, we may see even greater political fragmentation in information consumption as people get better tools for filtering the news based on their own reactions and reactions of other people like them.

In this paper, we report on the design and deployment of a widget designed to nudge its users to read balanced political viewpoints. This widget represented the aggregate political lean of users' weekly and all time reading behaviors to encourage those whose reading leaned one way or the other to read more balanced news.

Before discussing the current study, we will motivate it by briefly reviewing selective exposure theory, why exposure to diverse points of view is important, and the evidence from online environments for and against the type of reading behavior selective exposure theory would predict in an environment of choice in information. After the study results and discussion, we close with a discussion of a design space for such widgets.

Selective Exposure Theory

Long before the Internet, researchers observed that people prefer to access information that suits their own opinions and worldview and to avoid information that challenges these beliefs (Lazarsfeld et al. 1944). This behavior is one way that people can avoid the uncomfortable feeling of cognitive dissonance (Festinger 1957). By selecting confirmatory information, they need not reevaluate their existing information or stance and they can reaffirm that they have correct beliefs (Kastenmüller et al. 2010).

Selective exposure theory claims that people prefer to access supporting information and avoid challenging in-

formation. When people behave consistently with these preferences, the information they access supports their existing beliefs and preferences, and they are not prompted to go through the difficult process of reevaluating their views. Literature also describes some conditions under which people may *not* act according to that preference (Frey 1986). For example, people are less likely to prefer a diet of only agreeable information when they are particularly curious about the topic (Frey 1986), when potentially discordant information is expected to be particularly useful (Freedman 1965), or in response to a fairness norm (Sears 1965). The preferences described also predict how people will behave in information environments that offer choice between agreeable, disagreeable, and diverse information: given sufficient choice, they will read primarily agreeable information.

Selective Exposure to Political Information

This prediction alarms many political theorists, who see exposure to diverse information as a prerequisite for many positive outcomes. Society risks losing these positive outcomes if people prefer to access only agreeable information and are able to construct a political information diet free of discordant information. In this section, we review three positive outcomes that can only occur if people are exposed to diverse political news, opinions, and information. First, exposure to diverse views is a necessary ingredient in deliberative debate, which political theorists argue is necessary for a healthy democracy. Second, counter-attitudinal information is necessary for people to learn and for better problem solving. Finally, understanding the distribution of opinions is necessary for people to accept the legitimacy of decisions with which they may not agree. We discuss these in more detail below.

First, exposure to diverse viewpoints is a prerequisite for ideal deliberative debate (Habermas 1984). Democracy thrives in societies where political discussion is frequent and frequently approaches deliberative ideals (Cohen 1989). If people only expose themselves to agreeable political news or discussion, this debate will not occur, and negative consequences will result from its absence. Interaction with like-minded people leads to polarization: participants tend to end up with more extreme views than they started with (Brown 1986, Schkade et al. 2007, Sunstein 2002). Selective exposure to reinforcing news and opinion articles might similarly lead people to take more extreme positions, and the fragmentation of the audience to different, ideologically-agreeable spaces and sources may lead to on-line discussion of articles among homogenous groups that leads to even further polarization.

Second, exposure to diverse opinions promotes learning and better problem solving. When people hold inaccurate beliefs, they must be confronted with information that challenges these beliefs in order to correct their under-

standing (Frey 1986, Hart et al. 2009). Consideration of diverse opinions promotes divergent, out of the box thinking, which can improve problem solving (Nemeth 1986). Through exposure to multiple viewpoints, people become more aware of relevant information and are more able to think through all of the outcomes of a decision, and so societies will make better collective choices on important matters at all levels of government (Benhabib 1996).

Finally, there is a tendency for people, particularly those in the minority, to think that their own views are more broadly shared than they actually are (Sanders & Mullen 1982). This tendency is known as the false consensus effect: people can increase their self-esteem by seeing their views as normative (Ross et al. 1977). A better assessment of views' true popularity may help people accept the legitimacy of disagreeable outcomes in the political sphere, rather than clinging to conspiracy theories to explain how an imagined majority will was thwarted. Even when people reach different conclusions after hearing all of the sides, exposure to and consideration of different opinions persuade participants that the opponents' views have merit (Benhabib 1996). Achieving this goal, however, can be difficult. When people receive evidence that their views are not normative, they can continue to increase their self-esteem by perceiving people with contrary views as acting according to a situational constraint rather than according to their own attitudes (fundamental attribution error, Ross 1977) or as defective in some way (Ross et al. 1977).

The benefits, to the individual and to society, of considering diverse and challenging points of view and the dangers of selective exposure make it a reasonable goal for people to be exposed to viewpoints other than their own. Sunstein and others, however, raise alarms that the Internet's increased choice of news sources and better tools for filtering out disagreeable news will undermine the role of media in presenting people with diverse viewpoints (Pariser 2011, Sunstein 2001).

Selection exposure & the Internet

Even when people prefer agreeable information, environments with limited numbers of sources can prevent the expression of such preferences. People act on selective exposure preferences by selecting among sources rather than by selecting items from sources (Lowin 1967, Prior 2007). Though people might prefer a source with entirely agreeable items, if their only source choices all present some disagreeable information, they will not avoid reading altogether. Thus, in an environment where a modest number of media channels served a broad audience, mainstream channels could broadcast the news with enough diversity to appeal to a good-sized audience, without a risk of being avoided for that programming, and so their audience would have some exposure to informational programming that included challenging opinions. National

policy in the U.S. promoted a certain level of diversity among broadcast media, and a large portion of the news audience received these reasonably balanced broadcasts simply by leaving the television on after a favorite program, the so-called inadvertent audience (Bennett & Iyengar 2008, Prior 2007). Throughout the latter half of the 20th century, this limited the effects of selective exposure preferences in determining individuals' actual exposure.

Technologies that offer more choices and personalization, however, undermine this role of broadcast media. They allow for the delivery of more channels, which have economic incentives to cater to niche audiences (Mullainathan & Shleifer 2005). Considerable evidence shows that people self-segregate into ideological television viewing groups (Iyengar & Hahn 2009, Pew 2004), while broadcast news audiences have declined precipitously (Bennett & Iyengar 2008): nearly half of Americans watched one of three broadcast news programs in 1970; this had dropped to ten percent by 2007 (Prior 2007). Those who seek out news, rather than merely leaving the television on, can and do choose news programming that caters to ideological niches (Morris 2007). As people have greater choice in sources, they have greater potential to avoid disagreeable news sources or news sources altogether.

The Internet allows for even larger numbers of niche sources and for personalization of the sources that do have broad appeal. In *Being Digital*, Negroponte celebrates the potential for a "Daily Me" – a highly personalized and tailored news source, but for scholars concerned about negative societal consequences of political selective exposure (e.g., Sunstein 2001), the prospect of such news sources is alarming. Research presents conflicting evidence about the extent to which people prefer agreeable political information (e.g., Sears & Friedman 1967) and to which they are using the Internet act on those preferences.

- people use the Internet to access and engage with mostly agreeable political information (e.g., Purcell et al. 2010, Adamic & Glance 2005, Lawrence et al. 2010, Gilbert et al. 2009),
- people say they use the Internet to seek out a greater variety in sources (Stromer-Galley 2003), Internet users become more aware of political news and events (Horrigan et al. 2004), and USENET groups feature spirited disagreements (Kelly et al. 2005),
- people are indifferent to challenging information but seek reinforcing information (Garrett 2009a, 2009b).

Gentzkow and Shapiro (2011) analyzed traffic data from comScore's panel of over one million US Internet users and found that ideological isolation in people's online news sources was lower than many feared, and lower than their day-to-day face-to-face interactions, yet higher than the ideological segregation of consumption of offline news sources.

Prior work attempted to resolve which of these apparently contradictory theories describes people's preferences and behavior for accessing political news online. An online experiment found that some people prefer agreeable news while others prefer a mix of agreeable and disagreeable viewpoints (Munson & Resnick 2010).

These results, however, leave several important questions. Are there individual differences that predict to what extent people access agreeable political information while avoiding challenging material? Can people who use the Internet to read mostly agreeable news be nudged to read more diverse news? The second question – whether or not people can be nudged to read more diverse reading habits – has been partially answered in laboratory settings. Park et al. (2009) designed a system, *NewsCube*, which grouped articles on the same topic into "aspects" that reflected different sub-topics defined by the appearance of different words. During a laboratory experiment, subjects using NewsCube read more stories and explored more aspects on each topic, compared to Google News or an interface that grouped stories into aspects randomly. Could people be similarly nudged in the wild, and can they be nudged toward ideological diversity rather than just the topical aspect diversity found in the NewsCube study?

The Balancer Study

In this section, we report on a field study with the following goals:

1. Evaluate whether some individual characteristics – including their political preferences, demographics, and personality attributes – predict the political bias in an individual's online news-reading behavior.
2. See whether feedback about the aggregate bias or lean one's reading behavior altered that behavior, and if there are individual differences in responsiveness to that feedback.

We designed a widget that that reflected users' reading behavior for the week to date along with normative messaging. This widget shows an approximate histogram of the user's liberal and conservative pages visited (which we used as a proxy for reading) for the week to date. This was motivated by the idea that while many people might agree, in principle, with the normative goal of reading diverse news, they might not realize how skewed their own reading behavior actually is.

Communicating a norm can increase the persuasive power of information, particularly when an individual's present behavior deviates from that norm (Schultz et al. 2007). Feedback in the form of a simple histogram of counts of liberal and conservative items does not communicate any form of norm – just the individual's own behavior. We believed that communicating an injunctive norm (that is, what one *should* do (Cialdini et al. 1991))

might be more effective than a simple histogram. Apropos to the name of the research project – Balance – we chose to implement the widget as a character on a tightrope, with his balance affected by the histogram. If one's reading behavior is too skewed, the character is in peril of falling (Figure 1). We hoped this would encourage subjects to consider the norm of fairness and balance, one of the methods for counteracting selective exposure (Sears 1965), though it was also possible that they would treat the exercise of balancing the character as a bit of a game, which would also have the desired effect of nudging people toward more diverse exposure. The widget reflected history for the week-to-date, reset on Sundays.

Initial study

We initially focused on news aggregators where the main content is selected by users' votes and that have substantial politics sections, specifically *Digg* and *Reddit*. We created a user-installable extension to the Firefox web browser that augmented *Digg* and *Reddit* with the balancer widget.¹ It tracked click-throughs to news stories from those sites. Although 178 people installed the extension in the fall of 2011, and they had a median of 1.89 visits to *Digg* and/or *Reddit* that fall (mean = 5.88), few of them clicked through to news stories (median = 0.4/day, mean = 0.95). As a result, only six people read enough stories in three months to complete the series of stages in our experiment, which varied the forms of feedback shown to users.

Perhaps skimming article abstracts on *Digg* or *Reddit* is sufficient news exposure for users of those sites, and so they rarely click through to news stories. It is also possible that our subjects were getting most of their political news elsewhere – directly from news websites, through other aggregators such as *Google News*, through mobile applications, and so on – and were getting something else from *Digg* and *Reddit*.

Revised study

Ideally, a study design improving on study 1's limitations would allow us to observe subjects' reading behavior across all of their political news sources, rather than just articles they arrive at through *Digg* and *Reddit*, be able to classify all of these observed items, and offer the user sufficient value that they would want to install it and participate in the study, including answering a questionnaire and sharing their information with us.

Observing news reading behavior

There are a number of possible technical implementations that would at least partially achieve these goals, including a system-wide proxy, an application that analyzes and



Figure 1. Example feedback. Left: No articles (beginning of week), center: unbalanced, right: balanced.

transmits a user's history once, and a browser extension that accesses a user's web history.

With a **system-wide proxy**, all of a subject's web traffic would pass through a proxy administered by researchers. This would allow tracking behavior from all web browsers and applications (unless they overrode the proxy), across multiple devices, including mobile. Workplace computers would likely still be excluded, since many companies use their own proxies and/or do not give employees the permissions required to set a proxy. Unfortunately, it is also somewhat resource-intensive to run a proxy and requires high uptime to keep subjects happy. Menchen-Trevino and Karr (2011) pursued this approach with *Roxy*, but rented a dedicated server to support their 46 study participants.

Another approach would be to write a **history analysis application** that subjects download and run. The analysis could be conducted locally (preserving more privacy) or remotely on our system. Such an application could read from the history directories or files for major web browsers. Variations would need to be written for different operating systems, and it might fail on non-standard configurations.

Similar to the original study design, we could also use a **browser extension** to monitor subjects' web browsing and/or access and classify their entire browsing history. As with the application, classification could occur locally or remotely. Unlike the application method, a single extension could be cross-operating system, and users may also be more willing to install a browser extension, with limited permissions, than a full application. Neither the browser extension nor the history analysis application would be able to observe traffic from mobile applications or non-browser desktop applications (such as an RSS reader), and, unless the user syncs their browsing history between their mobile and desktop browsers, would not be able to track articles read on the mobile browser. Despite these drawbacks, we felt that a browser extension was the best approach for this study.

Classifying articles

When we had only a limited set of articles – as in Study 1, where only articles receiving enough Digg or Reddit upvotes were included – it was possible to use Mechanical Turkers to classify articles in a reasonably short period of time. This approach was not suitable for classifying entire web histories for hundreds or even thousands of users.

¹ This initial browser extension also included another widget, not discussed here for space constraints.

Researchers have succeeded at using machine learning to classify articles as liberal or conservative, fairly reliably, over time (Dehghani et al. 2011, Oh et al. 2009). These classifiers would need up-to-date training data to reflect current political figures, positions, and topics, but classifiers using *Digg* articles and votes, or articles linked from classified blogs, similar to the approach developed by Zhou et al. (2011), could provide a steady source of classified articles to use as training data. These classification methods, however, require accessing the website, extracting the article, and then classifying it as political or not, and, if political, according to its bias. Doing so in a way that offered timely feedback was beyond our resources.

Our alternative approach was to classify simply based on an item's URL. This is how *Memeorandum Colors*² provides its visualization: it reflects the political leaning of sites that typically link to the source, though the individual item's bias may differ. There is also precedent for building hindsight widgets based on URL-based classification. In 2010, *Slate* released a tool that let people check the bias of their reading history against 112 websites that had been classified by Gentzkow and Shapiro (2011) and then showed them their isolation score and how they compared to other *Slate* readers (Matlin et al. 2010). We also classified sites using *Digg* votes (Munson et al. 2009). We assembled a "whitelist" of sites, consisting of the union of sites from those three sources. All three assigned scores in a range from conservative to most liberal. For our study, we normalized them to the range (-1, 1); when a site appeared in multiple data sets, we simply averaged the values.

Study Design

Apparatus

The *Balancer* Chrome extension shows a persistent iconic indicator of the subject's balance in the browser window and provides more detail when clicked (tightrope walkers in Figure 2). The persistent icon helps people to monitor their behavior and to make changes in response. Participants have responded positively to using persistent, glanceable visualizations for self-monitoring in studies in other domains, such as physical health (Consolvo et al. 2008) or personal transportation choices (Froehlich et al. 2009).

Based on users' visits to sites on our whitelist – including visiting the homepage of a site (e.g., cnn.com) or visiting an article on that site (e.g., cnn.com/date/article.html) – we calculated a Balance score, S_u . Call the site visits to whitelisted sites V_u . A user's aggregate balance score for any time period was calculated from the average whitelist scores, s_v , for sites in V_u : $\zeta_u = \sum_{v \in V_u} s_v / |V_u|$. These ag-




Figure 2 Chrome extension. The icon provides persistent indication of the reader's history for the week, while clicking shows a larger view.

gregate scores were normalized to the range (1,11), a choice made because we had 11 possible versions of the tightrope walker. Empirically, we found that $S_u = 6(\zeta_u + 1)$ fell almost entirely in the range [1,11]. Values below 1 or above 11 were treated as 1 and 11, respectively.

Users who were unbalanced for the week ($S < 5$ or $S > 7$) also received some recommendations of sites representing more neutral positions or the "other" side. These recommendations were hard-coded to top-level addresses for news and opinion sites.

Procedure

After installation, the extension transmitted to our server the user's previous month's history of visits to our whitelist of news sites. Before being able to see our analysis of their browsing history, the subject was prompted to complete a brief questionnaire, consisting of a twenty-item version of the Big Five personality index (Goldberg 1990), age, gender, and two questions about their political preferences. The inclusion of the Big Five personality inventory was somewhat exploratory, but we hypothesized that some personality characteristics might predict reading behavior. For example, people who are more open might read news representing more viewpoints. Prior work also suggests that these personality traits can predict how individuals respond to persuasive techniques (Halko & Kientz 2010, Nov et al. 2013).

In the treatment condition, the extension ran in the background as they browsed, updating the visualization and transmitting their visits to whitelisted sites. This design allowed us to collect and retain users' browsing history, including their history before they installed the extension and their browsing activity after installation. Privacy was a concern in this design but was mitigated by: (a) transmitting the data over a secure connection, (b) not collecting the subject's email address or identifiable information as part of the registration process, (c) limiting collection to a whitelist of URLs of news sites, and (d) disabling IP-logging on the web server.

² http://waxy.org/2008/10/memeorandum_colors/

Experimental conditions. A control group completed the same enrollment process. Instead of immediately receiving feedback, though, these subjects were informed they would have to wait 28 days for their feedback and saw a countdown of the remaining days until they would receive feedback. A control group was necessary: simply comparing subjects' history post-intervention to their history pre-intervention (within-subject), would conflate the effect of being in a study (and being prompted to attend to one's reading balance) with any effect of the widget. It is also vulnerable to history effects: if a major news story broke during the study and were covered more by one side or another, it would affect the balance of subjects' reading. With a control treatment, we were able to compare the change in reading behavior pre-intervention to post- between the treatment and control groups.

Recruitment. We recruited participants by advertising the extension on Google's advertising network, in the Google Extension directory, and on our own homepages and social network feeds. The extension also attracted some media attention, including coverage on the NBC News site, a technology blog, and a number of local blogs and online magazines. We did not incentivize participation. By offering only a compelling application, though, we could attract participants who are more likely to represent the actual users of such a tool, and who are more likely to explore its features and make use of it. It does, however, mean that any conclusions we draw about browsing behavior are limited to this group.

By 18 November 2012, 990 people had completed the post-installation survey. Balancer users were asked to identify themselves according to political lean (1= strongly liberal; 7 = strongly conservative) and political partisanship (1 = strong democrat, 7 = strong republican) as well as report demographic information. Their mean political lean was 2.89 (stdev = 1.37), mean partisanship 2.93 (stdev = 1.48), and mean age 36.7 years (stdev = 13.8). We recruited 736 men and 181 women (73 declined to specify). We collapsed political lean and partisanship into an ideology score by averaging them. A further 155 subjects installed the extension but skipped the survey entirely; we included these individuals in analyses that did not require demographic information.

Table 1. OLS model for agreeable reading R_A , prior to installation. $F=26.09$ on 8 and 743 degrees of freedom, $p<0.0001$; adjusted $R^2=0.21$.

Predictor	Coefficient	Std Err	p
Intercept	1.567	0.335	< 0.0001
Ideology (1-7)	-0.335	0.025	< 0.0001
Age (years)	0.002	0.002	ns
Female	0.040	0.085	ns
Extraversion	-0.058	0.036	ns
Agreeableness	0.009	0.047	ns
Conscientiousness	-0.035	0.042	ns
Intellect	0.032	0.064	ns
Neuroticism	0.077	0.049	0.0597

Results: Browsing history

Unsurprisingly, users' pre-installation balance scores, S_{pre} , were correlated with their ideology ($\rho=0.28$, $t=-8.1731$, $df=859$, $p<0.0001$). To analyze individual differences in reading behavior before installation of our browsing extension among political ideologues, we constructed a dependent variable, R_A , that represented how ideologically agreeable a subject's reading was, according to the following:

$$R_A = \begin{cases} S_{pre} - 6 & \text{if liberal (ideology < 4)} \\ 6 - S_{pre} & \text{if conservative (ideology > 4)} \end{cases}$$

For this analysis, we excluded 76 moderates (those with ideology scores of exactly 4); sites classified as moderate were not "agreeable" to moderate readers in the same way that liberal or conservative sites were to liberal and conservative readers. For example, a site classified as moderate might reflect a moderate view, views from the left and right, or report news and events relatively free of any editorial viewpoint. A further 171 subjects were removed for either not having preinstall browsing history or having incomplete demographic information, leaving 743 subjects. We also cleaned the data: we excluded duplicate visits to a page (the exact same URL) within five minutes, and we excluded all visits where the "transition type" in the history was "reload."

We conducted a regression analysis predicting R_A based on ideology, age, gender, and big five indicators (Table 1). Ideology was a significant predictor of the amount of agreeable news that people read: in our sample, conservatives read less ideologically agreeable news than liberals. Scoring higher on neuroticism may be correlated (it is weakly significant) with reading slightly more agreeable news.

Results: Behavior Change

The study design consists of both a between subjects and a within subjects component. This allows for an analysis that compares the differences between subjects' pre- and post-intervention reading behavior across the treatment groups. The Balancer extension was designed to influence users' consumption of online political news, so we expected the following:

- Users whose news consumption leans liberal (pre-installation Balance score, $S_{pre} \geq 6.5$) and who receive feedback should decrease their Balance scores compared to users whose news consumption also leans liberal but receive no feedback.
- For users whose news consumption is balanced ($5.5 < S_{pre} < 6.5$), there should be no difference between the control and treatment groups for $S_{pre} - S_{post}$.
- Users whose news consumption leans conservative ($S_{pre} < 5.5$) and who receive feedback should increase their Balance scores compared to users whose news consumption also leans conservative but receive no feedback.

Our analysis is based on all subjects who have pre- and post-Balancer installation browsing history, clipped to 28 days before and 28 days after. Figure 3 summarizes the conditions and their needed reading for balance. We also performed t -tests to determine if change in score ($S_{\text{post}} - S_{\text{pre}}$) is different between the treatment and control groups.

The majority of users were unbalanced, reading “too many” liberal (blue) articles (first two boxplots from left). The control group did not appear to have changed their reading behavior substantially (mean $S_{\text{post}} - S_{\text{pre}} = -0.15$, $\text{sd} = 0.59$) although there are some outliers who appear to increase their visits to conservative or neutral sites. The treatment group, who received immediate graphical feedback and recommendations for red and gray sites, also appear to have changed their reading behaviors very little (mean $S_{\text{post}} - S_{\text{pre}} = -0.28$, $\text{sd} = 0.79$), but there are many outliers who increase their consumption of conservative news post-feedback. In a t -test, this difference between the control and treatment groups is significant ($t = 2.446$, $\text{df} = 600$, $p = 0.015$). For the median “needs red” user – with $S_{\text{pre}} = 7.30$ and visiting about 187 sites matching our white list every four weeks – this translates to approximately four new monthly visits to a site classified as extremely right-leaning, such as *The National Review* ($s = -0.81$) or *Fox News* ($s = -0.77$), or 20 new monthly visits to a site classified as fairly neutral, such as *ABC News* ($s = -0.02$). Recall that S is based on the mean score of pages visited, so extra visits to neutral sites move S toward a neutral score of 6.

The two boxplots in the middle are users who are considered “balanced” based on the pre-Balancer installation browsing history. Neither the control nor treatment groups appear to have changed their reading habits greatly (control mean $S_{\text{post}} - S_{\text{pre}} = 0.16$, $\text{sd} = 0.65$; treatment mean $S_{\text{post}} - S_{\text{pre}} = -0.02$, $\text{sd} = 0.84$; $t = 1.24$, $\text{df} = 103$, $p = 0.218$). This is to be expected: subjects in the treatment group received feedback that they were already reading consistent with a goal of balance.

The two boxplots on the right are users who read too many red articles pre-installation. These are both smaller groups, making statistical significance harder to find ($t = -0.425$, $\text{df} = 40$, $p = 0.67$), but the differences suggest that both control and treatment groups increased their consumption of liberal news greatly (control mean $S_{\text{post}} - S_{\text{pre}} = 0.83$, $\text{sd} = 1.21$; treatment mean $S_{\text{post}} - S_{\text{pre}} = 0.99$, $\text{sd} = 1.29$). For the median “needs blue” subject ($S_{\text{pre}} = 4.79$; 241 monthly page views), the difference in differences would translate to about seven additional monthly visits to a site such as *Think Progress* ($s = 0.89$) or 40 new monthly visits to a neutral site like *CNN* ($s = -0.01$).

Individual differences. The result that *some* subjects appeared to shift their reading behavior while others did not heightened our curiosity about whether there were predictors for individuals who would be more receptive to the behavior change feedback and thus increase the diversity




Figure 3. Summary of change in Balance scores ($S_{\text{post}} - S_{\text{pre}}$) by condition and “needed” types of articles. Width of the boxplots is related to the total number of people in each group.

of their reading more. We conducted a regression analysis predicting $R_{A \text{ Pre}} - R_{A \text{ Pre}}$, among people in the treatment group, again excluding moderates ($\text{ideology} = 4$), using the Big Five personality scores, ideology, gender, and age as predictors. We identified no statistically or clinically significant differences, nor was the model significant.

Discussion

First, our study supports the results that online news reading is somewhat fragmented, with ideology corresponding to the sites people visit. Among our participants, we saw differences in the amount of disagreeable news that subjects read, predicted by ideology: conservatives using our extension read less agreeable news than liberals using our extension. We caution against reading too much into this result, however. Our subjects were not a representative sample of the general population, but a group who self-selected to use the extension after learning about it in one of the places it was publicized. Many of the news outlets on which *Balancer* was publicized, by our own classifications, have a left-leaning audience, so conservatives who learned about the extension and installed it may be predisposed to consider sites outside of their ideological niche.

We found that a browser extension showing aggregate history of the lean of the sites people visit, along with a character visibly uncomfortable with too much lean, nudged *some* users to shift to more balanced reading behavior. This is, to our knowledge, the first study to show that nudges to read more balanced news can work in the wild. There is, however, a broader design space, with possibly more effective interventions, to explore.

Design space

Features of the design space for interventions to increase the diversity of political news to which people are exposed include the time at which the information is communicated and the type of information communicated. Most of the visualization approaches we discuss are designed to take

advantage of one of two ways of combatting selective exposure: priming a norm of balance or fairness or increasing curiosity about the different views on a story.

When information is communicated

A system can present the reader with information about an item or list of items at three times: before a user elects to read items (which we call *foresight*), as the user reads items (*currentsight*), or after a user reads items, usually in aggregate (*hindsight*) such as in our study.

Based on the principle of *kairos* – intervening at the right time and place (Fogg 2002) – foresight is likely most effective for shaping a user’s behavior as it occurs, since the reader can consider the information when they make their choice. By contrast, currentsight may only increase awareness of what they are doing now. This awareness may affect future decisions or their perception of an item’s claims, but the choice to read the item has already been made. Finally, a hindsight widget, such as the one used in the study, can show a user patterns of which they may not be aware. It can increase awareness and understanding of behavior over time and prompt reflection (Li et al. 2010).

The different presentation times are not mutually exclusive and likely complement each other. For example, a hindsight visualization might make a reader aware of an imbalance in their reading history, and a foresight visualization could then help them identify articles to read that would help to balance their information diet.

What information is communicated

The design of representations of viewpoint diversity of a reader’s information diet also involves many choices about what information to communicate. A visualization can be descriptive of the user’s behavior or choices, or of other people’s choices. Alternatively, it could be more prescriptive. It could show norms about desirable behavior. It could allow them to set goals and then monitor their progress against targets. Or it could provide recommendations.

Descriptive Information. Feedback might include only descriptive information about what the user could read, is reading, or has read (such as the predicted bias of an item or a histogram of the reading behavior). It might also include comparative information, such as how many and which friends have read an item, how their reading history compares to that of friends or other visitors to a site. Such information could serve as social proof of the value of other stories or reading patterns, or it could simply make people curious about what others are seeing.

Normative information. The visualization could also add explicitly normative information according to progress against a designer’s goal or a commitment the user had made. Solely informational messaging about what others do, such as a comparison of one’s energy to similar homes, can have a “boomerang effect” on individuals who are performing better than their peers. Adding a normative mes-

sage, such as a sad emoticon for poor-performing individuals and a happy emoticon for high-performing individuals, can eliminate the boomerang effect for people who are already performing well while still persuading those who are not performing as well (Schultz et al. 2007).

People appear to believe in a norm of accessing diverse and balanced views. Fox News, generally regarded as one of the most ideological television news sources in America (Kull et al. 2003, Rendall et al. 2001) finds value in marketing itself as “fair and balanced.” At the University of Michigan’s 2010 commencement, President Obama implored the audience to read more diverse news. Overall, people reacted positively to this message (Garrett & Resnick 2011, as well as comments in the *Ann Arbor Chronicle*). This is possibly because most people think the other side(s) should listen to their own perspective more, but also possibly because most people agree with the norm even if they do not choose it in the moment or are not aware of their own behavior. If the latter explanation is true, tools that remind people of a norm of fairness and that highlight when they are behaving inconsistently with this norm may be particularly effective.

Goals and Targets. The choice to add normative information about progress relative to a goal implies a design choice: should a goal be presented and if so, should it be user-set, set by social information (what others do), or set by the system’s designer? The goal-setting literature provides a number of techniques for making goals more effective. In general, they should be specific and challenging, yet achievable, and people should have a way of monitoring their progress toward them (Locke & Latham 2005). The individual should be committed to the goal, and techniques such as asking them to record or state their commitment or explaining why the goal is important can increase their commitment (Cialdini 2009, Locke & Latham 2005).

Recommendations. Systems can also present recommendations from a variety of sources. Here, too, there are many options. Stories on a particular topic could be clustered, such as in *NewsCube*’s presentation (Park et al. 2009), possibly increasing readers’ curiosity about the different clusters. Many news sites now recommend stories that are popular, both among the general audience and among a reader’s friends. The system could also show recommended items that would help a reader meet goals, such as conservative items to balance out a liberally slanted reading history. Providing people with convenient recommendations to content supporting potentially disagreeable opinions may be particularly effective, given research showing that while people seek sources that reinforce their opinions, they do not avoid clicking on links to articles with challenging opinions when presented with these links (Garrett 2009b).

With whom the information is shared

In addition to presenting information back to the reader, the information can also be shared with others. For example, an extension might pair up individuals or establish teams from different political ideologies, with the understanding that they will, out mutual reciprocity, read, and possibly discuss, ideas supporting viewpoints favored by others. Or an individual who holds different political views from friends or family might use the tool to try to read more about their viewpoints, but also use their own balanced reading to show that they are trying to understand them. Alternatively, a user's avatar on a comment board on a political or news website might indicate the balance of their political reading, providing some context about who is commenting. The news site *Newsvine* uses a similar concept, where a user's "vine" image represents their history and tenure with the site. With the increased adoption of commenting systems, such as *Disqus*, that span activity on many sites, there is potential reach across many sites with ideologically different audience groups, using just one software widget.

Individual differences

For pre-install reading behavior, we observed a potential small effect for neuroticism: individuals scoring higher on neuroticism may read slightly more agreeable news. We found no predictors for individual differences in efficacy of the *Balancer* feedback.

Given that past work has found individual differences in preferences for diversity and predicts individual differences in responsiveness to persuasive technology, we expected to find variable prior reading behavior and variable responsiveness to the *Balancer* intervention. We observed variations in prior reading behavior and individual variations in efficacy. We did not find strong predictors for these differences, however. We hope that future work will find ways of distinguish between diversity seeking and challenge averse individuals, and between people who are and are not responsive to these persuasive techniques.

Future Work

Interviews or surveys of *Balancer* users might help us to better understand for whom the extension worked and did not work, as well as how to improve the tool for future deployments. Such data can also inform priorities for exploring the design space described above.

We also note that balanced news does not always achieve the goals that motivated our work: sometimes it can lead to more polarized or unbalanced views (Lord et al. 1979). As *Balancer*, *NewsCube* and other interventions are developed and tested, researchers should return to study whether the more balanced exposure decreases polarization, and increases learning and understanding of other

viewpoints. It is possible that different types of widgets or interventions will promote different processing of challenging views, even if they lead to comparable changes in reading behavior.

Conclusion

Scholars argue that if people use the Internet to read political news and opinion in ways consistent with selective exposure preferences, it lead to ideological fragmentation and polarization, with harmful societal outcomes. In this paper, we describe a browser widget, *Balancer*, which shows users whether their reading history is consistent with a norm of balanced reading. In a field study, we found that its use caused unbalanced readers to make small but real improvements in the balance of their reading. Finally, we present a design space for other widgets that might achieve these goals and argue for evaluation of widgets representing different points within that space.

Acknowledgements

Erica Willar, Catherine Hong Le, Jeremy Canfield, Peter Andrews, Brian Ford, and Daniel Xiaodan Zhou all contributed to early designs of the extension. Thank you to reviewers and Souneil Park for comments that improved this paper. This work was funded by National Science Foundation award #IIS-0916099.

References

- Adamic, L; Glance, N. 2005. The Political Blogosphere and the 2004 US Election: Divided They Blog, *Proceedings of the 3rd international workshop on Link discovery*: 36-43.
- Benhabib, S. 1996. Toward a Deliberative Model of Democratic Legitimacy, in *Democracy and Difference: Contesting the Boundaries of the Political*, 67-94. Princeton University Press.
- Bennett, WL; Iyengar, S. 2008. A New Era of Minimal Effects? The Changing Foundations of Political Communication, *J Commun* 58: 707-731.
- Brown, R. 1986. *Social Psychology* (2nd ed). NY: Free Press.
- Cialdini, RB. 2009. *Influence: Science and Practice* (5th ed). Boston: Pearson.
- Cialdini, RB; Kallgren, CA; Reno, RR. 1991. A Focus Theory of Normative Conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. *Adv Exp Soc Psychol* 24: 201-231.
- Cohen, J. 1989. Deliberation and Democratic Legitimacy, In *The Good Polity: Normative Analysis of the State*, edited by A Hamlin & P Petit. New York: Blackwell.
- Consolvo, S; McDonald, DW; Toscos, T; Chen, M; Froehlich, JE; Harrison, B; Klasnja, P; LaMarca, A; LeGrand, L; Libby, R; Smith, I; Landay, JA. 2008. Activity Sensing in the Wild: A Field Trial of UbiFit Garden, *CHI* 2008.
- Dehghani, M; Gratch, J; Sachdeva, S; Sagae, K. 2011. Analyzing Conservative and Liberal Blogs Related to the Construction of the Ground Zero Mosque, *CogSci 2011*: 1853-1858.
- Donnellan, MB; Oswald, FL; Baird, BM; Lucas, RE. 2006. The mini-IPIP scales: Tiny-yet-effective measures of the Big Five factors of personality. *Psychol Assessment* 18: 192-203.
- Festinger, L. 1957. *A theory of cognitive dissonance*. Evanston: Row & Peterson.
- Fogg, BJ. 2002. *Persuasive Technology: Using Computers to Change What We Think and Do*, San Francisco: Morgan Kaufmann.
- Freedman, JL. 1965. Confidence, utility, and selective exposure, *J Exp Soc Psychol* 2: 778-780.

- Frey, D. 1986. Recent Research on Selective Exposure to Information, *Adv Exp Soc Psychol* 19: 41-80.
- Froehlich, JE; Dillahun, T; Klasnja, P; Mankoff, J; Consolvo, S; Harrison, B; Landay, JA. UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits, *CHI* 2009.
- Garrett, RK. 2009a. Echo chambers online?: Politically motivated selective exposure among Internet news users, *J Comput-Mediat Comm* 14(2): 265-285.
- Garrett, RK. 2009b. Politically Motivated Reinforcement Seeking: Reframing the Selective Exposure Debate. *J Communication* 59(4): 676-699.
- Garrett, RK; Resnick, P. 2011. Resisting Political Fragmentation on the Internet, *Daedalus* 140(4): 108-120.
- Gentzkow, M; Shapiro, JM. 2011. Ideological Separation Online and Offline, *Q J Econ* 126: 1799-1839.
- Gilbert, E; Bergstrom, T; Karahalios, K. 2009. Blogs Are Echo Chambers: Blogs Are Echo Chambers, *HICSS* 2009.
- Goldberg, LR. 1990. An alternative "description of personality": The Big-Five factor structure, *J Pers Soc Psychol* 59(6): 1216-1229.
- Hart, W; Albarracin, D; Eagly, AH; Brechan, I; Lindberg, MJ; Merrill, Lisa. 2009. Feeling Validated versus Correct: A meta-analysis, *Psychol Bull* 135(4): 555-588.
- Habermas, J. 1962. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Translated by T Burger (1989). Cambridge: MIT Press.
- Habermas, J. 1984. *The Theory of Communicative Action*. Translated by Thomas McCarthy. Boston: Beacon Press.
- Halko, S; Kientz, J. 2010. Personality and persuasive technology: An exploratory study on health-promoting mobile applications. *Persuasive Technology*: 150-161.
- Horrigan, J; Garret, RK; Resnick, P. 2004. The Internet and Democratic Debate: Wired Americans hear more points of view. Pew Internet & American Life Project.
- Iyengar, S; Hahn, KS. 2009. Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use, *J Commun* 59: 19-39.
- Kastenmüller, A; Greitemeyer, T; Jonas, E; Fischer, P; Frey, D. 2010. Selective exposure: The impact of collectivism and individualism, *Brit J Soc Psychol*, 49(4): 745-763.
- Kelly, J; Smith, M; Fisher, D. 2005. Opinion Diversity in Online Political Discussion Networks, *Online Deliberation* 2005.
- Kull, S; Clay, R; Lewis, E. 2003. Misperceptions, Media, and the Iraq War, *Polit Sci Q* 118(4): 569-598.
- Lawrence, E; Sides, J; Farrell, H. 2010. Self-Segregation or Deliberation? Blog Readership, Participation, and Polarization in American Politics, *Persp on Politics* 8: 141-157.
- Lazarsfeld, PF; Berelson, B; Gaudet, H. 1944. *The people's choice*, New York: Columbia University Press.
- Li, I; Dey, A; Forlizi, J. 2010. A stage-based model of personal informatics systems, *CHI* 2010, 557-566.
- Locke, EA; Latham, GP. 2002. Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *Am Psychol* 57(9): 705-717.
- Locke, EA; Latham, GP. 2005. Goal Setting Theory. In KG Smith, & MA Hitt, *Great minds in management: the process of theory development*, Oxford University Press.
- Lord, CG; Ross, L; Lepper MR. 1979. Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence. *J Pers Soc Psychol* 37(11): 2098-2109.
- Lowin, A. 1967. Approach and avoidance: alternative modes of selective exposure to information, *J Pers Soc Psychol* 6(1): 1-9.
- Matlin, C; Singer-Vine, J; Wilson, C. 29 April 2010. Escape from the Echo Chamber, *Slate*. <http://www.slate.com/id/2252247/>.
- Menchen-Trevino, E; Karr, C. 2011. "Researching real-world Web use with Roxy: Collecting observational Web data with informed consent," *J Inf Tech & Politics Conference on The Future of Computational Social Science*.
- Morris, JS. 2007. Slanted Objectivity? Perceived Media Bias, Cable News Exposure, and Political Attitudes, *Soc Sci Q* 88(3): 707-728.
- Mullainathan, S; Shleifer, A. 2005. The market for news, *Am Econ Rev* 95: 1031-1053.
- Munson, SA; Zhou, DX; Resnick, P. 2009. Sidelines: An Algorithm for Increasing Diversity in News and Opinion Aggregators, *ICWSM* 2009.
- Munson, SA; Resnick, P. 2010. Presenting Diverse Political Opinions: How and How Much, *CHI* 2010.
- Munson, SA; Resnick, P. 2011. The Prevalence of Political Discourse in Non-Political Blogs, *ICWSM* 2011.
- Negroponte, N. 1995. *Being Digital*. New York: Alfred A Knopf.
- Nemeth, CJ. 1986. Differential contributions to majority and minority influence, *Psychol Rev* 93(1): 23-32.
- Nov, O; Arazy, O; Lopez, C. and Brusilovsky, P. 2013. Exploring personality-targeted UI design in online social participation systems. *CHI* 2013.
- Obama, B. 2010. University of Michigan Commencement Address, as transcribed by the *Ann Arbor Chronicle*. 1 May 2010.
- Oh, A; Lee, H; Kim, Y. 2009. User Evaluation of a System for Classifying and Displaying Political Viewpoints of Weblogs, *ICWSM* 2009.
- Park, S; Kang, S; Chung, S; Song, J. 2009. NewsCube: delivering multiple aspects of news to mitigate media bias, *CHI* 2009.
- Pariser, E. 2011. *The Filter Bubble: What the Internet is Hiding from You*. Penguin Press.
- Pew Research Center for the People and Press. 2004. *Online news audience larger, more diverse: News audience increasingly polarized*. Washington, DC: Pew Research Center.
- Prior, M. 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*, Cambridge University Press.
- Purcell, K; Rainie, L; Mitchell, A; Rosentiel, T; Olmstead, K. 2010. *Understanding the participatory news consumer*. Pew Research Center.
- Rendall, S. 2001. Fox's Slanted Sources: Conservatives, Republicans far outnumber others, Fairness and Accuracy in Reporting.
- Ross, L. 1977. The intuitive psychologist and his shortcomings: Distortions in the attribution process, in *Adv Exp Soc Psychol*, edited by L. Berkowitz. New York: Academic Press.
- Ross, L; Greene, D; House, P. 1977. The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes. *J Exp Soc Psychol* 13: 279-301.
- 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.
- Schultz, PW; Nolan, JM; Cialdini, RB; Goldstein, NJ; Griskevicius, V. 2007. The Constructive, Deconstructive, and Reconstructive Power of Social Norms, *Psychol Sci* 18(5): 429-434.
- Sears, DO. 1965. Biased indoctrination and selectivity of exposure to new information, *Sociometry* 28: 363-376.
- Sears, DO; Friedman, JL. 1967. Selective Exposure to Information: A Critical Review, *Public Opin Q* 31(2): 194-213.
- Schkade, D; Sunstein, CR; Hastie, R. 2007. What Happened on Deliberation Day? *California Law Review* 95(3): 915-940.
- Stroud, NJ. Media Use and Political Predispositions: Revisiting the Concept of Selective Exposure, *Political Behavior* 3(3): 341-366.
- Stromer-Galley, J. 2003. Diversity of Political Opinion on the Internet: Users' Perspectives. *J Comput-Mediat Comm* 8(3).
- Sunstein, C. 2001. *Republic.com*. Princeton University Press.
- Sunstein, C. 2002. The Law of Group Polarization, *J Polit Philos* 10(2): 175-195.
- Zhou, XD; Resnick, P; Mei, Q. 2011. Classifying the Political Leaning of News Articles and Users from User Votes, *ICWSM* 2011.