

Personalized Filters Yes: Bubbles No


Paul Resnick
Professor
University of Michigan



PERSONALIZATION IS UNDER ATTACK



I'm glad you came back for this last session. We have a problem, folks. Our field is under attack. And we need to respond, by doing better research and by engaging in the public debate.



Ideas worth
spreading

Themes


Speakers

Talks

Translations

TALKS


Eli Pariser: Beware online "filter bubbles"



The Filter Bubble

What the Internet is Hiding From You

Eli Pariser



March of this year, Eli Pariser gave a TED talk, as a way to start publicizing his book, *The Filter Bubble*.

I understand some reference was made to this yesterday morning at the conference.
Quick Poll: how many of you have seen the TED talk, or read the book, or otherwise become familiar with Pariser's argument?

Bubble is a metaphor; Pariser argues that personalized filters put us inside bubbles that we can't see out of. But he doesn't strictly stick to the metaphor. The word bubble actually stands in for a variety of concerns.

Individuals:

We won't see what we really want: no serendipity, no exploration of new stuff.

We won't see what we need: things that are in our long-term interest, things that challenge us and change our minds.

As a society:

We won't have common reference points, because we all see different stuff

We will fragment into polarized enclaves

THE CONCERNS AREN'T NEW

GroupLens: An Open Architecture for Collaborative Filtering of Netnews

Paul Resnick, Neophytos Iacovou**, Mitesh Suchak*, Peter Bergstrom**, John Riedl***

* MIT Center for Coordination Science

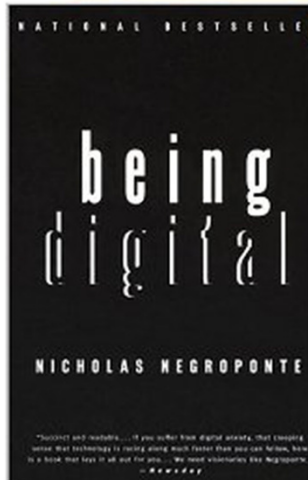
** University of Minnesota

Global Villages

Present newsgroups, like newspapers and local television shows before them, provide a shared history for their community of readers. With GroupLens, users may choose to read articles only from a small group with whom they share many common interests. Over time this could lead to a fracture of the global village into many small tribes, each forming a virtual community but nonetheless isolated from each other.

Some kind of fracture is inevitable and even desirable, because no user can keep up with the overwhelming volume of news produced each day. The question is whether the subgroups will be closed or permeable. One argument for prognosticating permeability is that many





154

serendipity

On Sunday afternoon, however, we may wish to experience the news with much more **serendipity**, learning about things we never knew we were interested in, being challenged by a cross-word puzzle, having a good laugh with Art Buchwald, and finding bargains in the ads. This is *The Daily Us*. The last thing you want on a rainy Sunday afternoon is a high-strung interface agent trying to remove the seemingly irrelevant material.

These are not two distinct states of being, black and white. We tend to move between them, and, depending on time available, time of day, and our mood, we will want lesser or greater degrees of personalization. Imagine a computer display of news stories with a knob that, like a volume control, allows you to crank personalization up or down. You could have many of



BUT THERE'S NEW RESONANCE



While the concerns mostly aren't new, there's a new public resonance for them.

republic.com
casssunstein

Customer Reviews

Republic.com

24 Reviews

5_star: (4)

4_star: (4)


3_star: (4)

2_star: (2)

1_star: (10)

A

SI

school of
information
University of Michigan

The reviews say things like, “not the Internet I know” and “he’s an authoritarian trying to interfere with free speech.”

The screenshot shows the TED website interface for a talk by Eli Pariser titled "Beware online 'filter bubbles'". The TED logo and tagline "Ideas worth spreading" are at the top left. Navigation links for Themes, Speakers, Talks, and Translations are at the top right. The talk title is prominently displayed, along with the view count "479,537 Views" and a comment count of "259 total comments". Below the title, there is a section for "Customer Reviews" showing 19 reviews with a star rating breakdown: 5 stars (15), 4 stars (3), 3 stars (0), 2 stars (1), and 1 star (0). The average customer review is shown as four stars. To the right of the reviews, there is a preview of the talk's content, titled "The Filter Bubble", with a snippet of text: "What [redacted] the [redacted] Internet [redacted] is [redacted] Hiding [redacted] From [redacted] You [redacted] Eli Pariser". At the bottom of the page, the "school of information" logo and "University of Michigan" are visible.

It has 479K views.

For comparison, Sergey Brin and Larry Page’s talk from 2007 has 219K views, and only 29 comments.

It got a mostly positive review in The New York Times, and a separate contributed op-ed article.

And it has rave reviews on Amazon.

The reviews have titles like, **“Thank God the Internet isn't hiding The Filter Bubble from me!”**, and **“An absolute must read for anyone who uses the Internet, left or right.”** and **“Shrewd, thoughtful, and well-executed Insight into Downside of Internet Personalization”**.

A New Parlor Game

The New York Times

By NATASHA SINGER
Published: May 28, 2011

“It’s totally creepy if you think about it,” said Tze Chun, [a filmmaker](#) who agreed to participate in a similar experiment at a recent dinner party we both attended in Brooklyn. Five of us used our phones to search for “Is Osama really dead?” a phrase Mr. Chun suggested.

Although our top 10 results included the same link — to Yahoo Canada answers — in first place, two of us also received a link to a post on [jewishjournal.com](#), a newspaper site. Meanwhile, Mr. Chun and two other filmmakers had links to more conspiratorial sites like [deadbodies.info](#).



There’s even a new parlor game, comparing search results on the same search.

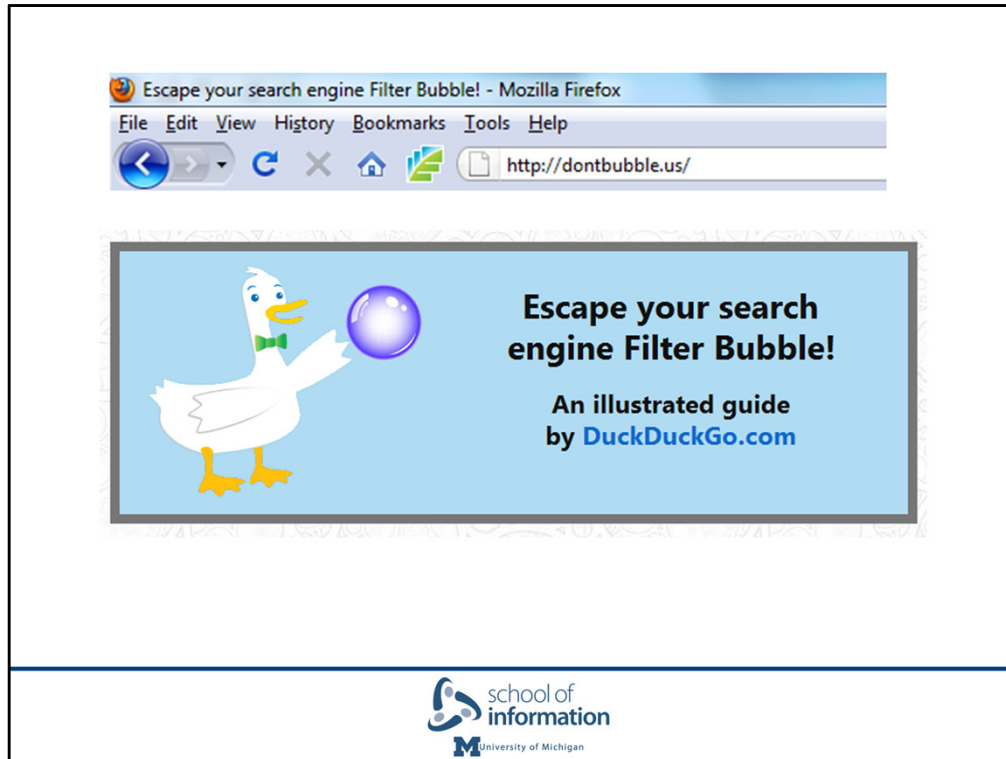
Pariser reports on playing the game: “In the spring of 2010, while the remains of the Deepwater Horizon oil rig were spewing crude oil into the Gulf of Mexico, I asked two friends to search for the term “BP.” They’re pretty similar—educated white left-leaning women who live in the Northeast. But the results they saw were quite different. One of my friends saw investment information about BP. The other saw news. For one, the first page of results contained links about the oil spill; for the other, there was nothing about it except for a promotional ad from BP.”

Students have sent emails to our all-school mailing list about their results. I think it has made the whole idea much less abstract for people. And caused them to ask all sorts of questions, like, what am I missing?

Pariser, Eli (2011). The Filter Bubble: What the Internet Is Hiding From You (Kindle Locations 93-96). The Penguin Press. Kindle Edition.



There's enough resonance that the ability to turn off personalization is being used as a marketing tool by the search engine DuckDuckGo.

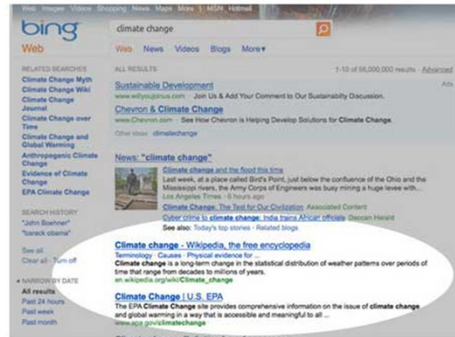


They have a little cartoon story at dontbubble.us that I've made a few excerpts from here.

search engines now show different results to different people.

Search for “Climate Change”

Informational Sites



vs. Climate Action Sites



Search for 'barack obama':

Ann gets MSNBC:



Elaine gets FOX News:





Of course, personalized filtering isn't just happening on searches. It's an integral part of the FB news feed.

FB doesn't reveal exactly how they're filtering, but it seems likely that they're tracking clicks, likes, and comments, and trying to "learn" a model of which people or which kind of content you're most likely to respond to.

Tell story of Eli Pariser's conservative friends' links not showing up in his feed.

PERSONALIZATION BAD?



So, is personalization really bad, inherently?

BAD PERSONALIZATION?



Or are people responding to experiences of badly executed personalization? Should we be looking to counteract the problems by turning off personalization or making it better? I'm going to argue that most of the problems are due to naïve personalization, that we, in this research community, have working on fixing, and should continue to do so.

But first, I need to unpack the concerns a little more clearly, and deal with them one at a time.

THE CONCERNS

- Trapped in the Old You



The first concern is that personalized filters will trap you in the old you. Whatever you express an interest in you get more of, and you get less and less of anything else.

So you won't get what you would have liked. You won't get the serendipitous encounters with things you would have liked. Moreover, you won't even realize what you're missing.

This, I think, is at the heart of the knee-jerk negative reaction people get to parlor game of seeing what's different about your results from everyone else's. I will argue, however, that is really a problem of bad personalization, not a problem of personalization. The most effective personalization will optimize the mix of exploration and exploitation, so that you will not be trapped in the old you. In fact, good personalization should give you more serendipitous experiences than bad personalization or no personalization.

- Trapped in the Old You
- Reinforcing Your Baser Instincts



The second concern is that you'll get what you want instead of what you really need.

People have short-term preferences that don't maximize their long-term utility. In the short-term, I'm lazy and don't want to exercise. But my long-term utility is maximized if I do exercise. Speaking of which, if anyone is sticking around Girona for the next couple of days and wants to go biking with me on the greenway from Olot to Girona, let me know.

Personalized filters, if they are based on user models that capture our short-term choices, our clicks, may reinforce our baser, shorter-term preferences, for entertainment rather than news, for example.

Here, too, I will argue that we just need better personalization, personalization that takes into account longer-term preferences.

- Trapped in the Old You
- Reinforcing Your Baser Instincts
- Fragmenting Society



The third concern is that if everyone gets what's good for them personally, it will be bad for society as a whole, because it will fragment into enclaves.

One version of this concern, relatively benign and again likely to be solved with “good personalization”, is that there will be no common points of reference. We won't be able to talk to each other over the water cooler because we won't have seen the same TV shows or news stories.

The other version of this concern is that society will fragment into subgroups, enclaves that start with different values and worldviews but grow farther apart over time as they are exposed to different facts and only reinforcing opinions. And people won't even be aware of the minority status of their own opinions, leading to mistrust and an inability to find common ground on important political decisions.

This is the issue of tribes and permeability that we raised in our original GroupLens paper.

It's not clear exactly how much cross-fertilization and permeability between groups is needed to prevent society from fragmenting. Perhaps a few brokers will be enough. But maybe not. This is the one concern that I think may not be solvable just with “better” personalization; we may need to have explicit perspective-taking features that help people understand what other people are seeing; and we may need to establish social norms of broad exposure and of using the perspective-taking features.

TRAPPED IN THE OLD YOU?

Not with Good User Models and Good Personalization



I think most of us have had serendipitous experiences using personalized recommender systems and personalized search engines like Google. Something comes up that is just perfect but isn't something we knew we were looking for.

In collaborative filtering based recommender systems, it comes from other people who share some sensibility having already discovered the good find. "I wasn't aware of this book/movie/restaurant, but people who liked other things I like did like it, so I'll try it." That's serendipity.

Multi-dimensional Preferences

- Many dimensions to match on
 - Topic
 - Location
 - Viewpoint/ideology
 - Audience
- Better match on one dimension may yield diversity on another



But I think it works similarly with filters on other dimensions. If I become obsessed with the war in Libya, and consume all the news I can find about it, eventually I'm going to find something that offers a different perspective than the "Qaddafi is evil and that trumps almost everything else" perspective that I started with.

It's only if the filter works by intersecting all the criteria, and there are plenty of things at that intersection, that I get narrowing on all the dimensions. Even then, as I will argue in a minute, this is not an inherent problem with personalization but an artifact of naïve implementations of personalization that fail to properly navigate the explore/exploit tradeoff or the declining marginal value of another similar item.

The Explore-Exploit Tradeoff

- Goal: optimize long-term click-throughs
- Can recommend from stream A or B
- Click-through rate from stream A has been higher
- Should you sometimes show from B? How many times?



The explore-exploit tradeoff is that sometimes you want to exploit the things that are known to be good, and sometimes it's better to explore, to try things that have lower expected value, in case they might turn out to be good and point the way towards a bunch of other things the user will like.

As a simplification, suppose you were trying to optimize the total click-through rate on recommended news articles over a long period of time. And suppose you have two black-box algorithms available, A and B. You track the click-through rate when you show items you show from stream A and from stream B. At some point in time, A has had a higher click-through rate so far than B. Should you still show items from B sometimes, if you're trying to maximize long-term click-throughs. The answer is yes, because, by exploring, you might learn that Bs are actually better than As. And it's worth a little expected short-term loss to find out.

...

The Explore-Exploit Tradeoff



Photo credit: [Yamaguchi先生](#)



Statisticians and theoretical computer scientist have formalized this problem as a multi-armed bandit problem.

Think of it as trying a bunch of slot machines in a Casino to try to figure out which one has the best average payoffs. (By the way, don't try this in a real casino, since they **all** have negative average payoffs.)

If the true click-through rates, or average payoffs, for streams A and B are unknown, but fixed, it turns out you should keep sampling, on occasion, forever. There's always a small chance that you'll learn your previous estimate was wrong, and that you just had an unlucky streak from it. If it does turn out you were wrong, you'll benefit from the better choice forever. With discounted payoffs or a finite horizon, there are circumstances where you would stop experimenting, and it starts getting more complicated from there.

The Explore-Exploit Tradeoff



Photo credit: [Yamaguchi先生](#)



The point is that, even if we assume people's preferences aren't changing, and we're just trying to learn those preferences in order to maximize click-throughs, the optimal personalized recommender should **not** always recommend its best guess for you. It should do some exploration of other options, in order to better learn about your tastes.

If the user's tastes may change over time, there is even more reason for the personalization algorithms to be doing continual exploration of user tastes, not just exploitation of their best model of user tastes.

Moreover, optimal personalization algorithms will do more exploration for another reason: when you give people more of the same, eventually they get bored and their click-through rates will go down. So optimal personalization algorithms have to keep exploring.

Conventional Item Preference Models

- Utility on a per-item basis
- Independent of other items in set/portfolio



Not accounting for changing utility is actually a weakness in our typical user models, and we're going to need to get more sophisticated.

Typically, we model user preferences as if the utility of an item of some type can't be determined independent of the consumption of any other items.

We know that this is just an approximation. Eventually, anyone will get bored and have lower utility for the 7,000th episode of a TV show than they had for the 7th, no matter how good the show.

The Need for Portfolio Preferences

- Presenting Diverse Political Opinions: How and How Much
- CHI 2010
- with Sean Munson

The screenshot shows a web interface for a political opinion aggregator. It displays a list of news articles with headlines and brief summaries. The articles include:

- Why the public school matters (Paul Krugman - NYT.com)**: Headline about public schools and the role of the state.
- It's come to this August vote? A rebuke for as at all says (The Hill)**: Headline about the August vote and the role of the state.
- Krugman: Obama's health system worked 'very well' for Barack's cancer survivor without insurance (Think Progress)**: Headline about Obama's health system and the role of the state.
- Gay Marriage: Rags (Maggie Gallagher - Townhall.com)**: Headline about gay marriage and the role of the state.
- Party meeting and of the day (The Hill)**: Headline about the party meeting and the role of the state.

Below the list of articles, there is a survey question: "Suppose this was the front page of a political opinion aggregator. How would you feel about the viewpoints represented in it?"

The survey options are:

- ☐ Very interested
- ☐ Somewhat interested
- ☐ Neither satisfied nor dissatisfied
- ☐ Somewhat dissatisfied
- ☐ Very dissatisfied

Below the survey options, there is a text input field for the user to provide feedback: "Please tell us a bit about why you responded the way you did".

But that user model doesn't always hold. Sometimes utility of an item depends on what other items are shown.

We recruited a bunch of Mechanical Turkers and gave them a set of news articles every day. (next slide)

The Need for Portfolio Preferences

- Presenting Diverse Political Opinions: How and How Much
- CHI 2010
- with Sean Munson

The screenshot shows a survey form with a list of political articles and a rating scale. The articles are:

- **Why the public school matters (Paul Krugman - NYT.com)**
- **It's time to stop August 2010's a setback for us all, says Obama (The Hill)**
- **Krugman: Obama's health system worked 'very well' for bankrupt cancer survivors (Think Progress)**
- **Gay Marriage Raps (Maggie Gallagher - Townhall.com)**
- **Party meeting and attendance (David S. Broder - Washington Post)**

Below the list is a rating scale:

Suppose this was the front page of a political opinion aggregator. How would you feel about the viewpoints represented in it?

- ☐ Very dissatisfied
- ☐ Somewhat dissatisfied
- ☐ Neither satisfied nor dissatisfied
- ☐ Somewhat satisfied
- ☐ Very satisfied

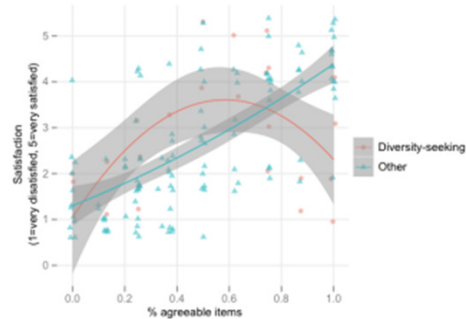
Below the rating scale is a text box for a comment:

Please help a bit about why you responded the way you did

After each set, we asked them to rate their satisfaction with the set, on a five-point scale. We systematically varied the number of articles that were matched or mismatched with their own political position as liberals or conservatives.

The Need for Portfolio Preferences

- Presenting Diverse Political Opinions: How and How Much
- CHI 2010
- with Sean Munson

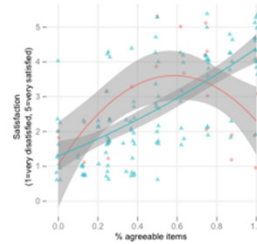


Some of the people, the blue curve, seemed to have per-item preferences, and to prefer articles that reinforced their viewpoint. On the x-axis is the percentage of agreeable items in the set, and the y-axis is expressed satisfaction with the set.

But some people clearly had preferences over the whole set, that they preferred a mix of reinforcing and challenging items. Their peak satisfaction came with sets that were about 60% reinforcing items.

Models of Portfolio Preferences

- Conventional Item preference models
 - Utility on a per-item basis
 - Independent of other items in set/portfolio
- Set/Portfolio preference models
 - Utility for complete sets
 - Or for items contingent on...
 - other items in set
 - recency of exposure



In addition to preferences that are contingent on other items in a set, there is often simply declining marginal utility: if you've seen an item of one type recently (not necessarily in the current set) you will get a little less utility for the next one of that type than you got for the previous one.

REINFORCING YOUR BASER INSTINCTS?



How about the second concern?

Delayed Indicators of Long-term Value



If we want to avoid the trap of building a model based on the short-term preferences user's reveal, we need to know about not immediate value, but value in retrospect.

When I read an Amazon review of a product, if the review starts, "Just got the item and I'm so excited," I ignore the rest of the review.

Story of student reactions to 501 at end of semester vs. end of next summer

Ideas for how to collect these in personalization interfaces?

- monitor dwell time, not just clicks; links; copy and paste of phrases; phrases that appear in something you write
- ask for explicit feedback, retrospectively after some time

Nudges Toward Your Better Self



Sometimes we need a little help to stick with our long-term preferences, our better selves. For example, many people wish that they read a balanced news diet, they think it's a good idea, but they don't actually do it when confronted with a particular set of choices.

Nudges Toward Your Better Self

The screenshot shows the Digg website interface. At the top is a navigation bar with links for 'Submit Link', 'My News', 'Top News', 'Upcoming', 'Join Digg!', 'Login', and a search bar. Below this is a secondary navigation bar with categories: 'All Topics', 'Business', 'Entertainment', 'Gaming', 'Lifestyle', 'Offbeat', 'Politics', 'Science', 'Sports', 'Technology', and 'World News'. The main content area is titled 'Upcoming Politics News' and features a list of four stories, each with a Digg score, title, source, and a brief description. The stories are: 1. 'The Inside Story of Marcus Bachmann's 'Ungay' Therapy' (148 Digg), 2. 'Police charge mother in Nashville airport altercation' (71 Digg), 3. 'Obama's Drug Czar Rejects Medical Marijuana' (86 Digg), and 4. 'Michele Bachmann's Church: The Pope Is The Anti-Christ' (91 Digg). To the right of the stories is a 'Balance' widget with a cartoon illustration of a person on a seesaw and text indicating the user needs to read 12 more red articles for balance. At the bottom of the page is a logo for the 'school of information' at the 'University of Michigan'.

Upcoming Politics News Trending ▾ All Stories ▾

148
Digg ↑ **The Inside Story of Marcus Bachmann's 'Ungay' Therapy**
blogs.citypages.com — When Andrew Ramirez was 17 years old, he came out to his family. His mom and sisters accepted the news, but his stepfather didn't. He forced ... 15 hr 44 min ago
via bukowsky 30 Comments 498 Save Bury

71
Digg ↑ **Police charge mother in Nashville airport altercation**
tennessean.com — A 41-year-old Clarksville woman was arrested after Nashville airport authorities say she was belligerent and verbally abusive to security officers.... 17 hr 24 min ago
via ferretman 13 Comments 701 Save Bury

86
Digg ↑ **Obama's Drug Czar Rejects Medical Marijuana**
blogs.sfweekly.com — If you needed more proof that Obama's promises on medical marijuana were false, look no further than his Drug Czar's most recent report, ... 9 hr 52 min ago
via beshirhappy 9 Comments 28 Save Bury

91
Digg ↑ **Michele Bachmann's Church: The Pope Is The Anti-Christ**
tpmdc.talkingpointsmemo.com — Josh Green is out with a story that could rock the religious-conservative foundations of Rep. Michele Bachmann's (R-MN) presidential campaign.... 9 hr 50 min ago

Balance Freebie About Stats

You need to read 12 more red articles for balance this week.
My History

Hot Stories From Digg

school of information
University of Michigan

Foresight and hindsight widgets

(next slide, opportunity to reflect on history)

Nudges Toward Your Better Self

digg+balance

Your reading history:

[Back to Digg](#)

Jul 14 -

these 30 articles:

all articles:

22

blue

22

blue

8

red

8

red

Read



How Many Congressmen Does It
Take To Screw Up a Light Bulb?

Read



A New Bid to Make 'South California'
the 51st State

FRAGMENTING SOCIETY



How about the problem of fragmenting society? There were actually two sub-problems, a lack of common reference points, and division into tribes, where each tribe's members talks among themselves, in echo chambers.

Common Reference Points



I'm not too worried about the lack of common reference points. There's still plenty of pop culture. When big sporting events like the World Cup happen, people realize they are interesting enough to enough people that they individually prefer to learn something about them.

But I might be wrong. A Michigan undergrad just wrote to me with an idea for a website that would have one news story per day, so that all students would read the same story on campus, and she'd be able to talk about it with other students before and after class.

In any case, I think it's pretty natural for the user model to include some utility for seeing the same things that others see.

Common Reference Points



“Negroponte wrote, in the 1990s, of the Daily Me, ‘Imagine a computer display of news stories with a knob that, like a volume control, allows you to crank personalization up or down.’ He wrote of it as a way to generate serendipity. But, as I’ve argued, I think you get plenty of serendipity from good personalization. Instead, we can think of it as a way to get more common reference points with the population as a whole.

I’m willing to leave this one up to individual users to set the utility they place on it (or to our personalization methods to infer that utility function). I am not convinced that there is a great societal externality, beyond what individuals already individually perceive, of having some common reference points. So I don’t think we need to nudge or impose this on them.

Echo Chambers



As for retreating into echo chambers, the real question, as we suggested back in our original paper, is how much cross-fertilization there will be. And that depends a lot on user preferences.

Challenge Aversion?

- Everyone likes reinforcement of pre-existing opinions



What are people's preferences, especially with regard to politically challenging information?

There's a long history of research on what's called "selective exposure". There's pretty clear evidence that people like to see information that reinforces their existing opinions.

Challenge Aversion?

- Everyone likes reinforcement of pre-existing opinions
- But...



Sometimes that is described as an aversion to challenge, which seems like an equivalent formulation. But it isn't really the same, it turns out.

Challenge Aversion? Not So Much

- Everyone likes reinforcement of pre-existing opinions
- But...
 - ...challenge is only mildly aversive on average



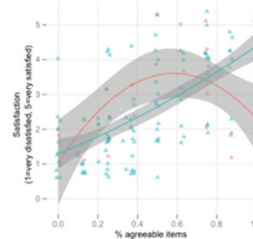
My former student, Kelly Garrett, in his dissertation research, since published in Communication Studies journals, found that people are much more likely click on news stories that they expect to have reinforcing information, but that expecting it to contain some challenging information has only a slight negative impact on probability of clicking, and, once you click, you spend more time reading the article.

Moreover, he found that, in a nationally representative survey, people who were more partisan were more likely to be aware of arguments *against* their positions and candidates than were people with milder preferences.

Reviewing other studies as well, Garrett concludes that there appears not to be a strong aversion to challenging information, at best a mild one. Most of the observed selective exposure results can be chalked up seeking reinforcement, not aversion to challenge.

Challenge Aversion? Not So Much

- Everyone likes reinforcement
- But...
- ...challenge is only mildly aversive on average
 - Some people prefer a mix of reinforcement and challenge



Also, recall the study I told you about before, where there are individual differences, some preferring more of a mix, and others preferring as much reinforcement as possible.

Challenge Aversion? Not So Much

- Everyone likes reinforcement
- But...
- ...challenge is only mildly aversive on average
 - Some people prefer a mix of reinforcement and challenge
 - Everyone has other criteria that matter



And even when people want reinforcement or want to avoid challenge, there are always other factors that matter. If it's a better fit on relevance, or better written, people will accept a little less fit on ideology.

So, contrary to what some others who are arguing that better personalization will inevitably lead to fragmentation into tribes and polarization among those tribes, I think it could turn out the other way, because there is not a strong universal preference to avoid challenge, there are other forces that will naturally lead to some exposure to challenge for everyone, and a lot for some people. And that may be enough to keep society from fragmenting.

Perspective-Taking Features



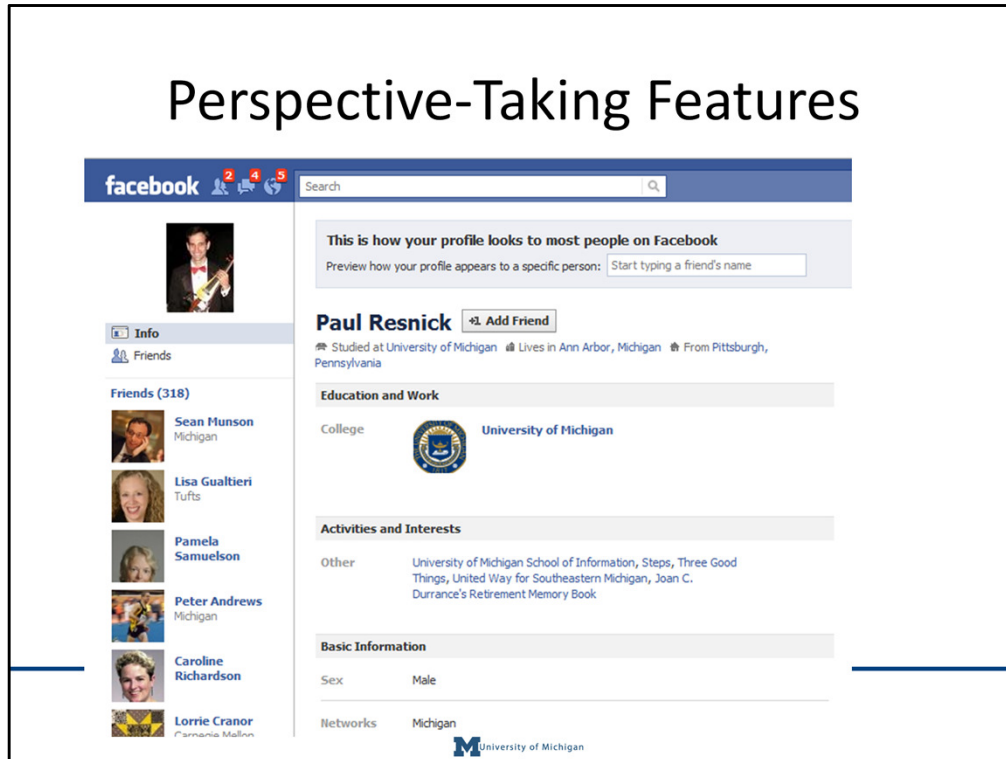
I do think that some awareness of what others are seeing is a public good.

Social psychologists have found that there is a false consensus effect. People who have an attribute think the attribute is more common than people who don't have the attribute. In a 1983 article in the *European Journal of Social Psychology*, Sanders and Mullen found that the perceptions of those in the minority on the attribute were less accurate: indeed, on almost every attribute they thought at least half the population shared the attribute, even though in most cases less than a third did.

You can imagine that, if this carries over to political opinions, and people with unusual views incorrectly think that most people share their views, they could become quite alienated from the political process when their views are not acted on, and will start to concoct conspiracy theories.

Sanders, G.S. and Mullen, B. (1982). "Accuracy in perceptions of consensus: Differential tendencies of people with majority and minority positions," *European Journal of Social Psychology* 13(1): 57-70.

Perspective-Taking Features



For personalized filtering, we can take inspiration from a feature in Facebook that lets you see a page from the perspective of different people.

Here's how my profile page looks to strangers.

Perspective-Taking Features



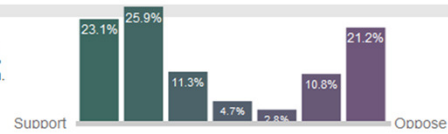
And here's how it looks to one of my students, Sean Munson. He can see that I'm married, and see various photos people have tagged me in.

Perspective-Taking Features

– Living Voters Guide, Kriplean et al

Explore what other voters think about 1098.

Click on a bar to see some of the key points that the group believes in.



Travis Kriplean, a PhD student at the University of Washington, has taken this same idea and applied it to a Living Voters Guide that lets people explore Pro and Con arguments for ballot initiatives in Washington State last year.

First, you can see the histogram of the positions on the ballot measure declared by all the other registered users of the site.

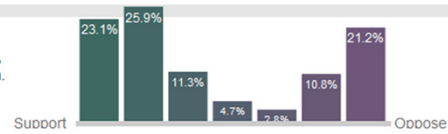
If you click on the leftmost bar, you can see the pro and con arguments most often listed by strong supporters of the ballot measure. (click)

Perspective-Taking Features

– Living Voters Guide, Kriplean et al

Explore what other voters think about 1098.

Click on a bar to see some of the key points that the group believes in.



Key points for those who strongly support 1098

Points often included in the positions of those who strongly support Initiative 1098.

show overall key points
show your pros and cons

Pros



by Brian Glanz

This will lower taxes for almost everyone, and raise taxes for a tiny few who can afford it. Most states tax this way, it's simple fairness.

discuss



by Katie Schmidt

Given recent statewide budget cuts, this is a great way to reduce cuts to essential services without burdening those who cannot afford it.

Cons



by Josh Beauregard

Allows legislature to increase the tax with a simple majority rather than a 2/3rds majority.

discuss (1)



by Akash Saini

Local companies like Bartell Drugs will be forced to pay more state taxes than larger national corporations like Walmart. The difference? 9%

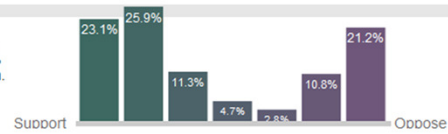
Even the strong supporters do list some Con arguments, such as the proposed tax would allow legislators to increase the tax with a simple majority rather than a 2/3 majority.

Perspective-Taking Features

– Living Voters Guide, Kriplean et al

Explore what other voters think about 1098.

Click on a bar to see some of the key points that the group believes in.



Key points for those who oppose 1098

Points often included in the positions of those who oppose Initiative 1098.

[show overall key points](#)
[show your pros and cons](#)

Pros



by Andrew Hamada

This taxing scheme should help alleviate the burden on the middle class, particularly first-time home buyers.

[read more](#) | [discuss](#)



by Rod Brooks

The state budget will be infused with a significant new source of revenue.

[discuss](#)

Cons



by Michelle Munneke

Legislators have routinely raided "dedicated accounts" in the past 10 years. Olympia has raided 74 "dedicated budget accounts," ...

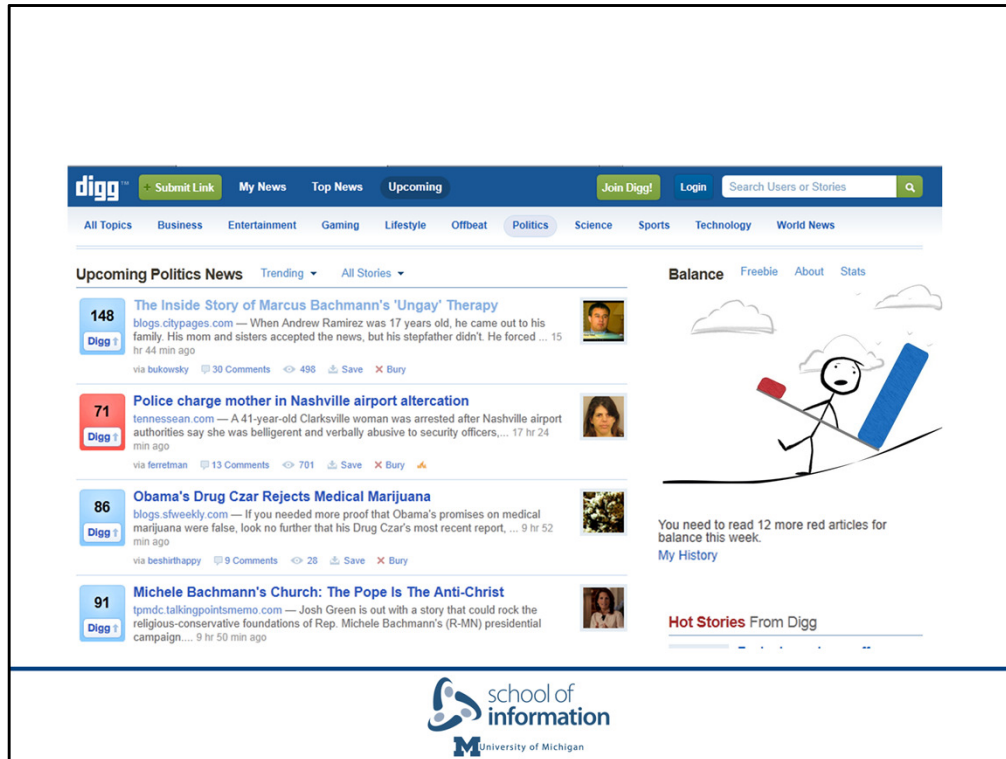
[read more](#) | [discuss](#)



by Michelle Munneke

Why more money to Education? State Legislature has increased K through HS spending 70% the past 12 years. Yet, test scores are declining!

But those who strongly oppose the measure listed different con arguments as their most popular, such as a claim that legislators have raided dedicated accounts in the past.



In the annotated version of Digg that we've developed so far, we don't actually filter anything. But we're working on a version that does. It will set a different popularity threshold for items that are liked by your side or liked by the opposition. When we do, a nice perspective-taking feature would be to give users the option to see what the news stream looks like to someone who has a different political filter on it.

CONCLUDING THOUGHTS

Filtering Happens

- Any sorting will hide some stuff
- There is no “natural” default sort order

Non-Personalization is Worse

- Worse individual matches
- Fewer things get seen by someone
 - Less chance for viral effects to counteract the machine's mistakes



When things are seen by someone, human agency can intervene to get things noticed even if they are not heavily promoted.

Better Personalization

- Multi-dimensional preferences
- Optimize the explore-exploit tradeoff
- Portfolio preferences
- Delayed Indicators of Preferences
- Nudges Toward Long-term Preferences
- Common Reference Point Features
- Perspective-Taking Features

Filterer's Responsibilities



I'm not entirely unsympathetic to Pariser's argument that we have some responsibilities as designers, beyond simply making personalization work better. While I reject his claim that personalization itself will make us narrower in our exposure or fragment society into non-interacting tribes, I do think we, as the technology, have some new-found power and with that comes some responsibility.

I don't think that the natural forces of personalized filtering, done well, will have bad effects on society. But I do think there are some unnatural, deliberate manipulations that could have big impacts.

For example, FaceBook teamed with some academics during the last U.S. congressional election to give people reminders to go out and vote. (In the U.S., turnout is often very low). Some people got a version that mentioned specific friends who said they had voted; others got something more generic. Apparently, the more personalized version was a little more effective at getting people to vote. What if one political party was in control of FaceBook, and showed the more effective version to people it had profiled as being supporters, and the less effective version, or no version at all, to others. In a close election, that might be enough to swing the outcome.

Clearly, we need to have some public responsibility for sites that people are expecting to be non-partisan to actually be non-partisan. (next slide for the actual code)

Filterer's Responsibilities

- No Hidden Personalization
- No Partisan Manipulation of Personalization



Intuitively, my code of ethics would be no hidden personalization and no partisan manipulation of personalization. I wouldn't go further than that. I wouldn't establish any universal affirmative responsibilities for filters, to make people have diverse exposure. Personally, I'm working on making news aggregators that nudge people toward diverse exposure, but I don't think all news aggregators should be morally or legally bound to do the same. The burden is on me to make diverse news more attractive than anything else people can get.

Summary

- Personalization is not inherently evil
- Accurate personalization is less evil than unpersonalized or poorly personalized
- Filterers have power, and therefore some public responsibilities
- We need to engage the public debate



[Narrate slide, then...]

...and by “we” I mean everyone in this room. I think there’s a real danger that the general public, or at least liberal political bloc, is going to lock in on a consensus that personalized filters are harmful to individuals and society. I think that consensus would be a bad outcome for the debate.

I have to admit that I haven’t even made a blog post or or posted a review on Amazon or responded to those emails floating around on the School of Information email list. I hereby resolve that I will engage beyond the research community on the issue of the social implications of personalized filters. I hope that many of you will, too.

Better Personalization Anthem

User models set me free
as you build the Daily Me
Yes exploit, but please explore
could just be that I'll want more
Broaden what my models know
UMAP scholars make it so

Music by Wolfgang Amadeus Mozart
Lyrics by Paul Resnick and Joe Konstan



At the conference banquet the night before my talk, attendees from different countries were invited to find their compatriots and choose a song to sing for everyone else. (The five Americans sang, “This Land is Your Land”).

Inspired by that, I decided to compose a song we could all sing together to close the talk and the conference, and which would reinforce some themes of my talk.

The melody is “Twinkle, Twinkle, Little Star”. The conference venue was a converted church, and the acoustics were great. Many people sang along.