Breaking echo chambers with personalized news

Jiemai Wu

September 26, 2018

- New technologies enable contemporary media to collect and analyze personal data
- This means digital media (e.g. google and facebook) can and will provide different information customized for different readers.

Question: How does customized information depend on existing personal biases?

Popular (non-academic) theories claim that websites will design news feed that caters to individual's existing biases, thus creating:

- "Echo chambers" - Sunstein 2001, etc.
- "Filter bubbles" - Eli Pariser, 2010

Academic theories explain that traditional media (newspapers and cable TVs) indeed conform to population's existing biases to maximize subscription

- Survey: Gentzkow, Shapiro, Stone (2014)
- Gentzkow, Shapiro (2006, JPE)
- Suen (2004, Economic Journal)

New technology
$\Rightarrow$ more personalization
$\Rightarrow$ media can cater to personal biases more easily
$\Rightarrow$ more echo chambers and filter bubbles
???
Logical fallacy:

- Personalization $\nRightarrow$ personalization that conforms to the reader's existing bias
- Traditional media conform to readers' biases $\nRightarrow$ new media also conform to readers' biases

Previously:

- Gentzkow, Shapiro (2006, JPE) - reputation
- Suen (2004, Economic Journal) - delegation

- In these papers, media sell "headlines" - binary signals generated from their private information of the true state.
- Their objective is to maximize paid subscription to these headlines.
- Application: Newspapers, cable TVs
- Prediction: due to consumer demand, media generate headlines that **confirm** the reader's prior bias.

In my paper:

- A medium provides a **free** headline. The reader can click on it to learn its **associated content**.
- The medium's objective is to **maximize clicks**.
- Application: free-to-access websites with digital KPIs: "click-through rate", "watch time", "view-through rate"
- Prediction: medium provides headlines whose content **contradicts** the reader's prior bias.

Demand-side effect:

- Reader prefers to click on a news title that is contrary to her existing view.
- This is because a rational decision maker benefits only from information that can potentially change her action.

Supply-side effect:

- Website prefers to recommend more news that are contrary to the reader's current view.
- Doing so can make the reader feel less certain about the true state, hence encouraging clicks.

Remark

Although this paper gives the opposite prediction, it does not completely contradict Gentzkow-Shapiro (2006) or Suen (2004).

Model setup

- State of the world, $\omega \in \{L, H\}$
- In each state, there is a distribution of articles such that each article $a \in \mathbb{R}$ is an i.i.d. signal about the true state with the following conditional normal distribution

$$a \sim \begin{cases} N(\mu, \sigma) & \text{when } \omega = H \\ N(-\mu, \sigma) & \text{when } \omega = L \end{cases}$$

- Each article a is associated with a title t such that

$$t(a) = \begin{cases} t^{+} & \text{when } a \geq 0 \\ t^{-} & \text{when } a < 0 \end{cases}$$

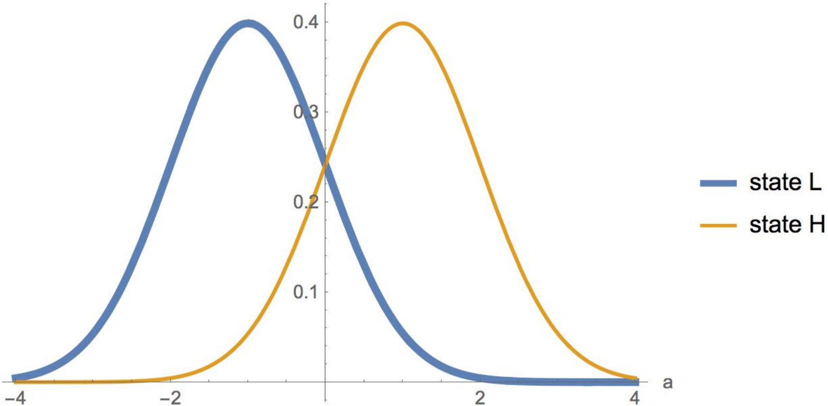
0.4
0.3
0.2
0.1
−4
−2
2
4
a
state L
state H

The website's strategy

- Design an algorithm that recommends either a positive or a negative article to the reader conditional on the true state ω, the reader's prior belief p_0, and the reader's preferences.
- If the website decides to recommend a positive (negative) article, it randomly draws an article from the distribution of articles with that title and send it to the reader.
- The reader initially sees only the title t; she must click on the title (which incurs a cost) in order to learn the associated article a.

The reader makes two choices sequentially.

1. choose whether to pay a cost c in order to click on the recommended title and obtain its associated signal a
2. Then, choose a binary action, either h or l, that best matches her posterior belief about the true state of the world.

Payoffs

- The website get payoff 1 if the reader chooses to click on the recommended title and 0 otherwise.
- The reader's payoff depends on her click, action, and the true state.

	$\omega = H$	$\omega = L$
choose h	0	$-v_1$
choose l	$-v_2$	0

if the reader does not click

	$\omega = H$	$\omega = L$
choose h	$-c$	$-v_1 - c$
choose l	$-v_2 - c$	$-c$

if the reader clicks

- the reader chooses action h if and only if $\Pr(H) \geq \frac{v_1}{v_1+v_2}$

Timeline

1. The websites designs an algorithm for its news feed.
2. Nature reveals the true state ω, the reader's prior belief p_0, and the reader's preferences (v_1, v_2, c).
3. Based on the algorithm, the website recommends a title t to the reader.
4. Observing the title, the reader updates her interim belief p_1 and decides whether to click on the title. If she clicks, she pays a cost c and learns the article a associated with the title.
5. The reader updates her posterior belief p_2 and chooses a binary action that best matches the posterior belief.
6. The website and the reader receive their payoff.

Reader's optimization:

Fix reader's interim belief p_1. Compare two scenarios:

- Reader clicks on a positive title $\Rightarrow$ expected utility $u(p_1, +)$.
- Reader clicks on a negative title $\Rightarrow$ expected utility $u(p_1, -)$.

- Reader clicks on a positive title $\Rightarrow u(p_1, +)$.
- Reader clicks on a negative title $\Rightarrow u(p_1, -)$.

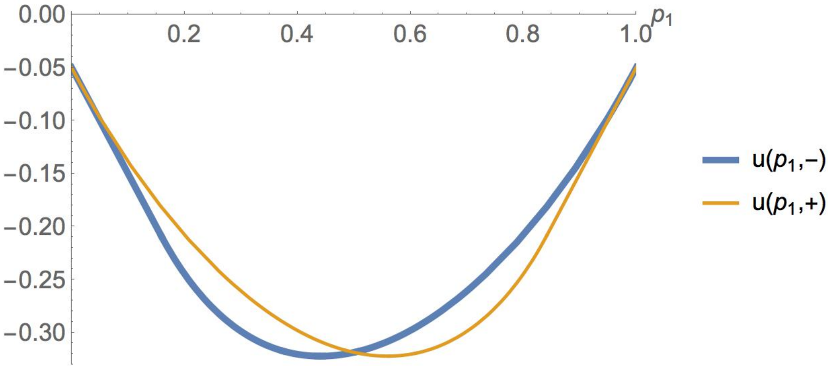

Figure: Example: $v_1 = v_2, \ \mu = \sigma = 1, \ c = 0.05$

Example

Consider the following two news articles. Which link would you click?

> *"Scientists have shown that eating too few vegetables is bad for your health."*

> *"Scientists have shown that eating too many vegetables is bad for your health."*

Recall

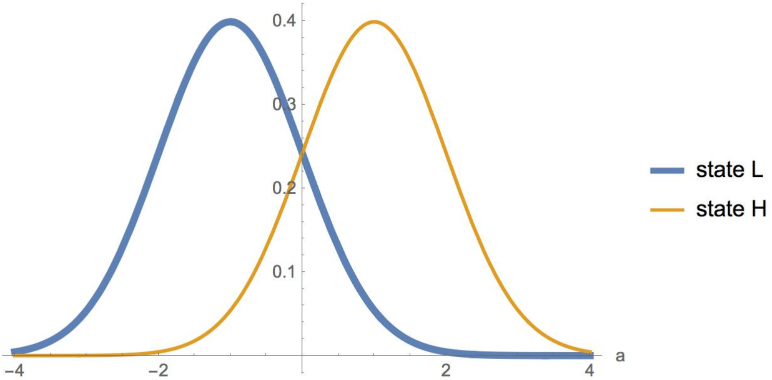

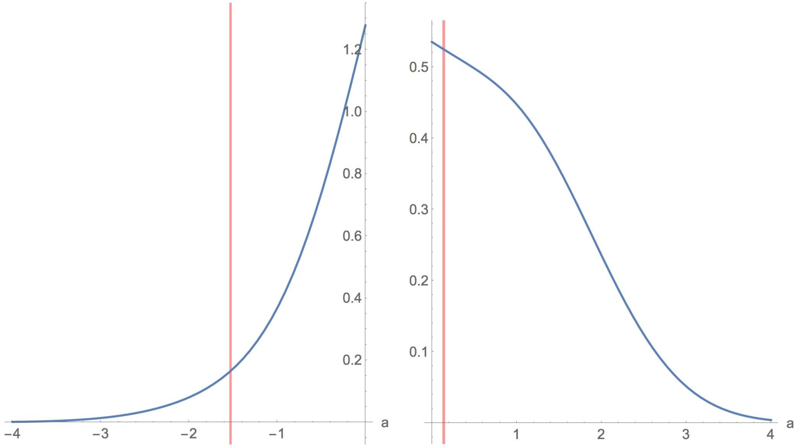

Figure: Left: expected distribution of negative articles. Right: expected distribution of positive articles. Red lines: values of negative/positive article that makes reader indifferent between actions l and h. $p_1 = 0.8$, $v_1 = v_2$, $\mu = \sigma = 1$.

Proposition

$$u(p_1,+) - u(p_1,-) \begin{cases} = 0 & when\ p_1 \in \left\{0, \frac{v_1}{v_1+v_2}, 1\right\} \\ > 0 & when\ p_1 \in \left(0, \frac{v_1}{v_1+v_2}\right) \\ < 0 & when\ p_1 \in \left(\frac{v_1}{v_1+v_2}, 1\right) \end{cases}$$

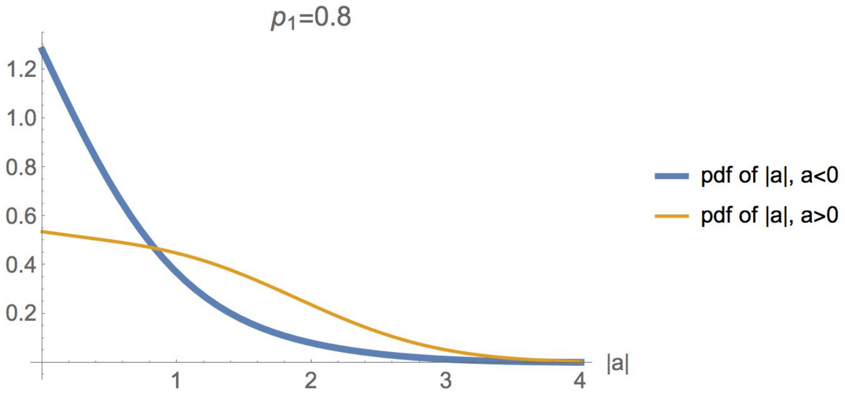

p_1=0.8
1.2
1.0
0.8
0.6
0.4
0.2
1
2
3
4
|a|
pdf of |a|, a<0
pdf of |a|, a>0

Which article do you expect to be more credible?

> *"Scientists have shown that eating too few vegetables is bad for your health."*

> *"Scientists have shown that eating too many vegetables is bad for your health."*

This paper does not contradict Gentzkow, Shapiro (2006, JPE).

- If the reader previous believes that vegetables are excellent for health, she will indeed believe that the information in the first link is more trustworthy.
- However, she is simply not interested in spending time on trustworthy information that she already knew.

Proposition

Let $v_1 = v_2$. Given p_1, let F^+ be the expected cdf of $|a|$ given $a > 0$ and F^- be the expected cdf of $|a|$ given $a < 0$. Then, if $p_1 > 0.5$, F^+ FOSD F^-; if $p_1 < 0.5$, F^- FOSD F^+.

Proposition 1+2: reader prefers an article opposite to her current belief, even though she thinks this article is less accurate.

Given p_1, let $u(p_1, \varnothing)$ denote the expected utility of the reader when she does not click on the recommended title.

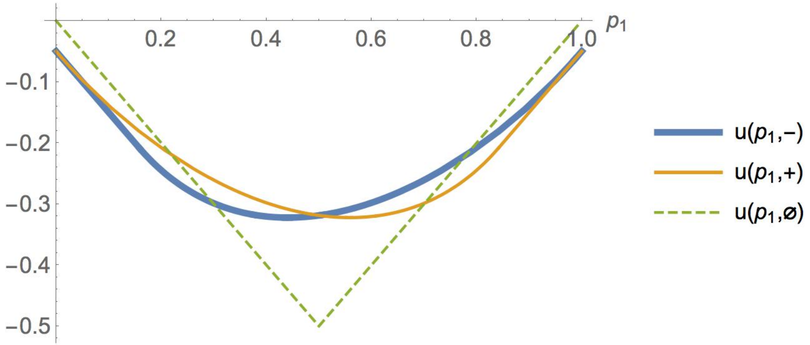

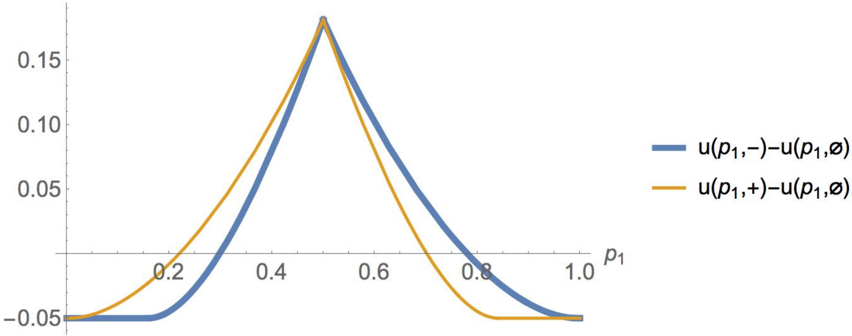
0.15
0.10
0.05
-0.05
0.2
0.4
0.6
0.8
1.0
p_1
u(p_1,−)−u(p_1,∅)
u(p_1,+)−u(p_1,∅)

Theorem

there exist thresholds $0 < \underline{\underline{p}} < \underline{p} < \frac{v_1}{v_1+v_2} < \bar{p} < \bar{\bar{p}} < 1$ *such that*

1. If the recommended title is t^+, the reader will click on it if and only if $p_1 \in \left[\underline{\underline{p}}, \bar{p}\right]$.
2. If the recommended title is t^-, the reader will click on it if and only if $p_1 \in \left[\underline{p}, \bar{\bar{p}}\right]$.

Website's optimization

Here is an optimal state-independent algorithm.

Theorem

When $p_0 \in \left[\underline{\underline{p}}, \bar{\bar{p}}\right]$, *it is an optimal algorithm for the website to always recommend*

$$\begin{cases} t^+ & when\ p_0 \in \left[\underline{\underline{p}}, \frac{v_1}{v_1+v_2}\right) \\ t^- & when\ p_0 \in \left[\frac{v_1}{v_1+v_2}, \bar{\bar{p}}\right] \end{cases}$$

The above algorithm is the uniquely optimal state-independent algorithm when $p_0 \in \left(\underline{\underline{p}}, \underline{p}\right) \cup (\bar{p}, \bar{\bar{p}})$.
i.e., it's the simplest one that works.

When p_0 is too extreme, the website uses the following optimal algorithm. [Kamenica and Gentzkow, 2011]

When $p_0 \in \left(0, \underline{\underline{p}}\right)$,

- the website always recommends t^+ with probability 1 in state H and probability $q_1 \in (0,1)$ in state L.
- q_1 is chosen so that the reader's interim belief is exactly $p_1 = \underline{\underline{p}}$ when t^+ is recommended.

When $p_0 \in (\bar{\bar{p}}, 1)$,

- the website always recommends t^- with probability 1 in state L and probability $q_2 \in (0,1)$ in state H.
- q_2 is chosen so that the reader's interim belief is exactly $p_1 = \bar{\bar{p}}$ when t^- is recommended.

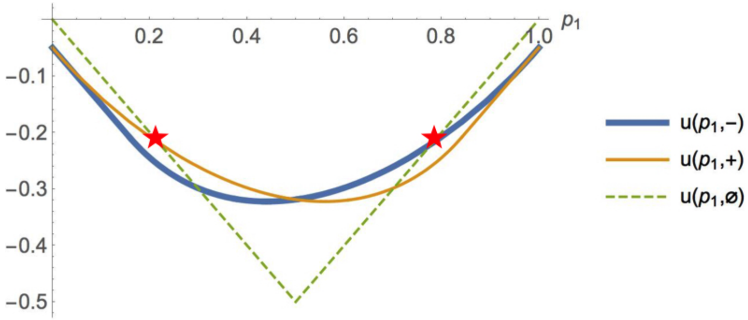

The optimal state-dependent algorithm that maximizes the reader's ex-ante expected utility when $p_0 \in \left[\underline{\underline{p}}, \bar{\bar{p}}\right]$:

- With probability $\frac{p_0 - \underline{\underline{p}}}{\bar{\bar{p}} - \underline{\underline{p}}}$, the website induces $p_1 = \bar{\bar{p}}$ and recommends t^-;
- With probability $\frac{\bar{\bar{p}} - p_0}{\bar{\bar{p}} - \underline{\underline{p}}}$, the website induces $p_1 = \underline{\underline{p}}$ and recommends t^+;

Suen (2004):

- Media are simply selling algorithms that map the true state to a binary headline
- Reader prefers an algorithm that has a headline that most often says "You are correct!"
- Implication: Traditional media that profit only from subscription tend to generate headlines that confirm the reader's bias.

This paper:

- Media use binary headlines as means to induce clicks
- In the Pareto optimal algorithm, a medium most often says: "You are correct! Therefore, you should read this opposite article!"
- Implication: Websites that profit from clicks tend to recommend articles that contradict the reader's bias.

Extension: biased taste for news

"I just like reading positive news more than negative news."

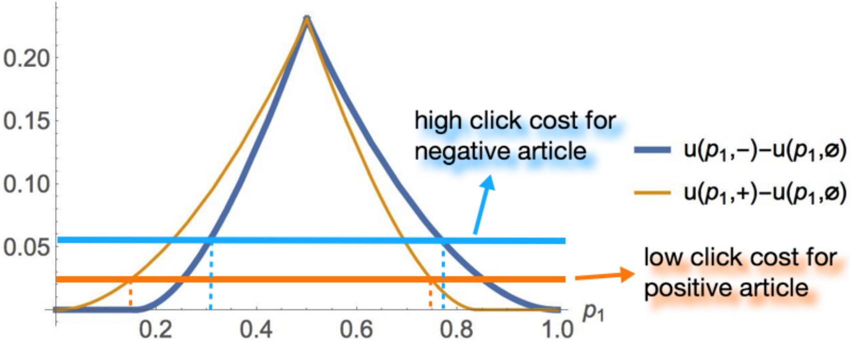

"I dislike news that contradict my current view."

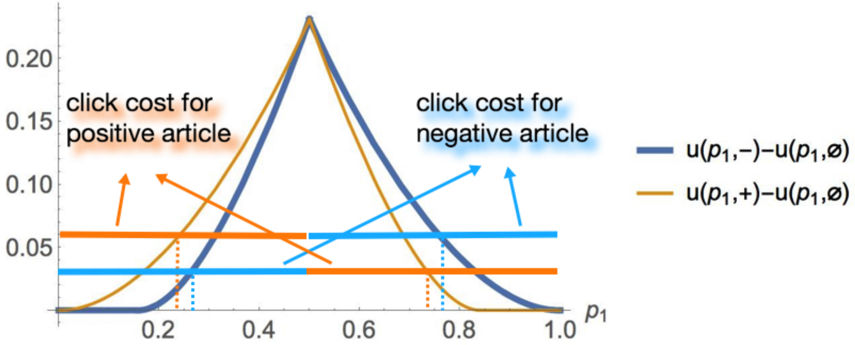

Conclusion

Contrary to popular beliefs and traditional theories, it is optimal for a click-maximizing website to recommend news articles that contradicts its reader's current bias. This is because

- the reader benefits more from information that can potentially change her mind
- the website can strategically recommend the opposite news to shatter reader's confidence about the true state, thus increasing her demand for information.