SEMINAR | AMERICAN UNIVERSITY | 21 OCTOBER 2026

Biased Information, Biased Preferences

How Fragmented Media Environments Undermine Support for Redistribution

Oscar Barrera Rodriguez University College Dublin
Lora Pavlova ZEW Mannheim **Li Yang** ZEW Mannheim

Pre-registered at the AEA RCT Registry, AEARCTR-0014878

People no longer see the same facts

- Beliefs about the economy shape the demand for redistribution
  (Alesina & Glaeser 2004; Bénabou & Tirole 2006)
- Explanations for inequality have polarized along party lines
  (Suhay 2025)
- News is chosen, and ideologically segregated
  (Gentzkow & Shapiro 2011; Levy 2021)

Two literatures, studied apart

- **Information experiments**
  the effect of exposure on people who did not choose it
  (Kuziemko et al. 2015; Haaland et al. 2023)
- **Selective exposure**
  what people choose, not how the same people respond to the other side
  (Chopra et al. 2024)

This paper. Does choosing what to read polarize preferences and votes? And does forcing the other side on people backfire?

This paper in one slide

Design

- Pre-registered survey experiment, **4,524** German adults
- Fielded Dec 2024 – Feb 2025, around the federal election of 23 Feb 2025
- Two real narratives about the 2022–23 inflation surge
- Four arms: no article, **forced**, **optional**, **choice**

Findings

1. Choice is confirmatory.
2. Exposure moves preferences toward the narrative read, policy by policy.
3. Facts land, even on opponents; on one policy the most right-leaning move *against* the article.
4. Persuasion on policy does not carry voters to the left platforms.

A political-science lens: Receive, Accept, Sample

Receive

Exposure to political messages rises with **political awareness**.

→

Accept

People **resist** messages inconsistent with their predispositions, but only if they can **see the link**.

→

Sample

Opinions are built on the spot from the **considerations at the top of the head**.

Key implication. When elites divide, the most aware partisans polarize: they receive messages from both sides but accept only their own (Zaller 1992, chs. 3 and 6).

Our model in RAS terms, with two twists

	Zaller (1992)	Our model	What changes
Receive	exposure rises with awareness	respondents **choose** what to read: value of information against the cost of exposure	reception is chosen
Accept	resist inconsistent messages you can recognize	updating, with one weight on the **fact** and another on the **policy**	acceptance has two tiers
Sample	average of top-of-mind considerations	outcomes measured minutes after reading	answers follow what the article brings to mind

Stage 1: What to read

$$u_i(s) = \underbrace{\alpha_i \lambda_i \left(\hat{\theta}_s - \mu_i\right)^2}_{\text{value of information}} - \underbrace{\varphi_i \left(\hat{\theta}_s - \mu_i\right)^2}_{\text{cost of exposure}} + \varepsilon_{is}$$

Value of information: what you expect to learn, net of the cost of changing your mind. Larger for the article that *disagrees* more.

Cost of exposure: the discomfort of reading the other side, paid *even if you do not change your mind.*

A cost of changing your mind alone never makes you avoid an article: you can always read and not update. **Confirmatory choice requires** $\varphi_i > \alpha_i \lambda_i$, given that both articles are expected to be equally accurate (otherwise: Gentzkow & Shapiro 2006).

Stages 2–3: What to believe, and what to want

Fact "inflation hit the poor harder"

$$\mu_i^{\text{post}} = \mu_i + \lambda_i\,(\hat{\theta}_s - \mu_i)$$

with $\lambda_i = \alpha_i/(\alpha_i + \kappa_i)$

$g_i(\cdot)$: what the fact implies for the policy

Policy "…so, a basic income?"

$$y_i = y_i^0 + \lambda_i^{\text{pol}}\left[g_i(\mu_i^{\text{post}}) - y_i^0\right]$$

with $\lambda_i^{\text{pol}} = \alpha_i/(\alpha_i + \kappa_i^{\text{pol}})$

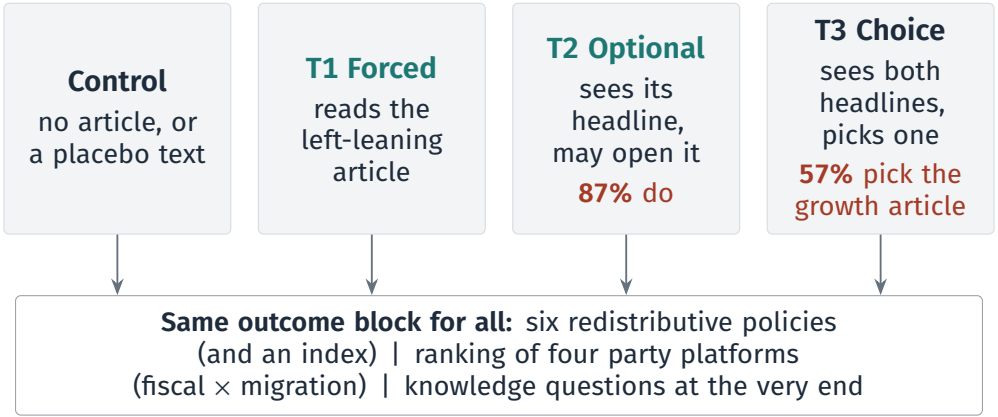
A cost of revision can **slow** movement toward a message, never **reverse** it.

Backfire needs one more ingredient

on the fact	discount the source as propaganda ($\rho > 1$)
on the policy	**identity**: stand apart from the position the source advocates ($\eta_i > \alpha_i$)
on the policy	**a model of the economy** in which the fact argues against the policy ($g_i' < 0$)

Only the first touches the fact, and we measure both tiers.

Design: four arms



Qualtrics online panel, quotas on age, gender and education; block randomization within strata, about 1,100 per arm; fielded 12 Dec 2024 – 25 Feb 2025.

Two real narratives about the same inflation

LEFT-LEANING

“Social gap in inflation at its highest level”

Low-income singles: **10.7%** inflation, against 7.8% for the wealthy (Nov. 2022)

Hardship example: a family of four on the **citizen’s allowance**

Proposes an excess-profits tax, a wealth tax and a higher minimum wage

Sources: TAZ, DGB, Hans-Böckler Foundation

RIGHT-LEANING

“Handouts won’t save Germany”

Growth, not “ever-increasing social benefits”

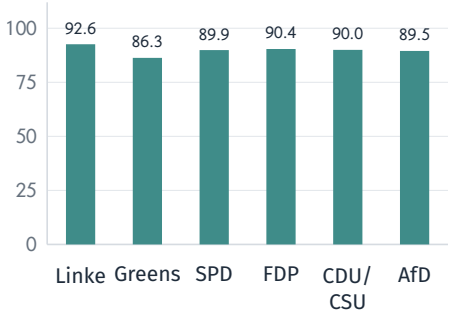
The **citizen’s allowance** weakens work incentives; “40 percent of recipients are foreigners”

“Excessive spending risks fueling inflation”; proposes corporate tax cuts

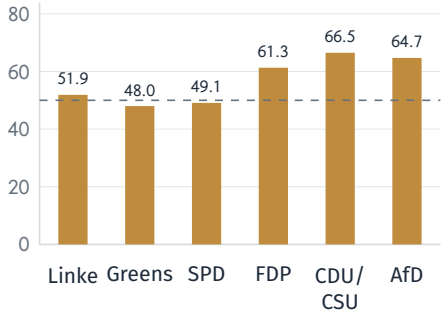
Sources: Welt, Focus, DW, Financial Times

Result 1: Whether to read is not ideological; what to read is

T2 Optional: % who open the left article



T3 Choice: % who choose the growth article



By 2021 vote. Overall, 87% open the left article and 57% choose the growth article; the share choosing growth rises by **5 pp per step** of left–right self-placement ($p < 0.001$).

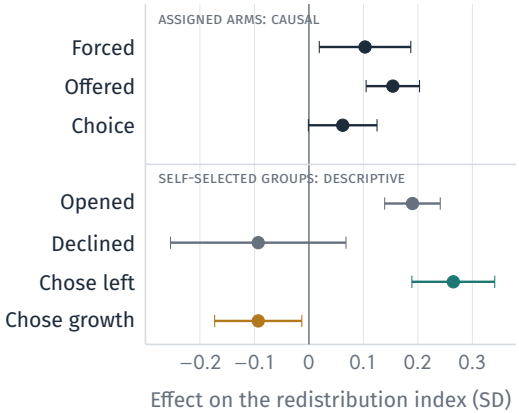
Result 2: Preferences move, policy by policy

ITT, SD units	Wealth tax	Inheritance tax	Corporate tax	Minimum wage	Top income tax	**Basic income**
T1 Forced	0.05	0.09	**0.17**	0.05	0.00	−0.02
T2 Optional	**0.15**	**0.13**	**0.17**	0.05	0.02	−0.04
T3 Choice	0.09	0.11	0.06	−0.01	0.00	**−0.10**†

Red: $p < 0.10$ under both the wild-cluster bootstrap ($G = 16$ states) and the Romano–Wolf adjustment for six outcomes. †Bootstrap $p = 0.053$; Romano–Wolf $p = 0.27$.

The left article moves **taxing the rich**, not **giving to the poor** (Cavaillé & Trump 2015). A basic income moves only where the growth article is read: −0.10 in the choice arm, −0.20 among those who chose it (self-selected).

Chosen versus imposed: the effect, or the people?



Those who chose the left article end **0.27 SD** above control; forced readers **0.10**.

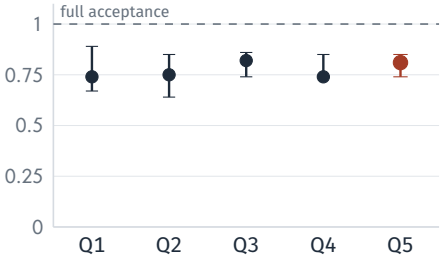
But choosers **started elsewhere**: the gap mixes movement with selection. Offering and imposing do **not** differ on the index (0.05, $p = 0.30$).

> What a fragmented environment *delivers*, not the causal effect of choosing.

95% CIs, clustered by state. Under the wild-cluster bootstrap the forced effect on the index is only suggestive ($p = 0.14$).

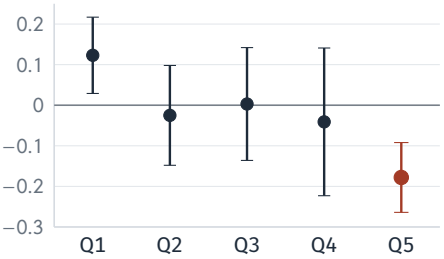
Result 3: Facts land; one policy moves against

Belief: share of the way to the article's 10.7%



Q1 most left-leaning · · · Q5 most right-leaning

Policy: effect on support for a basic income (SD)



Q1 most left-leaning · · · Q5 most right-leaning

Forced readers, by quintile of right-leaning propensity. The most right-leaning fifth accept the figure as fully as anyone (prior 7.0%, posterior 10.0%), yet move against a basic income. No other policy, and not the index, moves against the article.

Where does the resistance sit?

Where	Mechanism	Verdict
The **fact**	discount the article as propaganda, $\rho > 1$ (Nyhan & Reifler 2010)	✗ Q5 update as much as anyone
The **policy**: identity	a basic income marks the other side: conceding a fact is cheap, a policy is not (Kahan 2013)	✓ fact intact, one policy
The **policy**: a model of the economy	"transfers feed inflation", as the growth article argues (Andre et al. 2026)	✓ fact intact, one policy
The **answer**: what comes to mind	the article brings migration to mind; "universal" reads as "also for migrants"	✓ fact intact, one policy

Not the **factual backfire** that is rare (Wood & Porter 2019), but an **attitudinal backlash** (Guess & Coppock 2020), confined to one policy. This wave cannot tell the last three readings apart.

A Sample-stage reading: what the article brings to mind

The left article's example of hardship: a family on the **citizen's allowance**

↓

Right-leaning reporting ties the allowance to **foreign recipients**; our growth article: "40 percent of recipients are foreigners"

↓

A **universal** basic income then reads as transfers that also reach migrants

↓

Support falls among the most right-leaning

Why the article, not the word? The control group sees the same word, "universal".

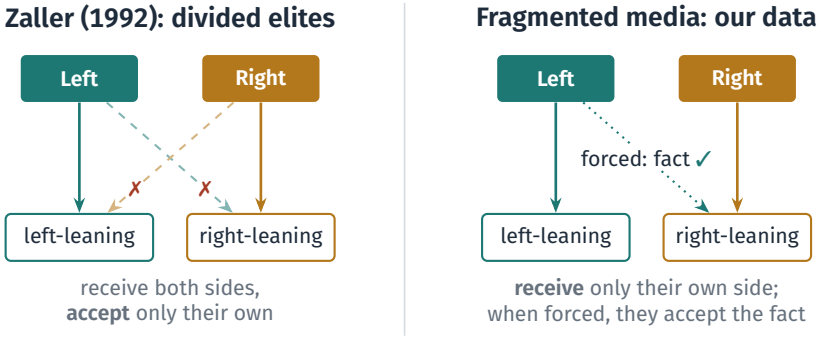
Consistent with

- backfire only on the policy whose wording invites the link
- chose the growth article: basic income −0.20 SD (self-selected)
- priming immigration lowers support for redistribution (Alesina, Miano & Stantcheva 2023)

Tests

- post hoc: AfD voters (migration) or the more educated (inflation)?
- wave 2: "universal" vs. citizens-only wording

Zaller revisited: polarization moves upstream



Zaller's condition for resistance, *seeing the link to one's predispositions*, explains the two tiers. An inflation statistic contradicts no one's predispositions; a universal basic income, in Germany in early 2025, does.

Votes: policy moves left, politics does not

Four platforms, fiscal × migration

		control
Left	open borders, radical redistribution	11.5%
Center-Left	inclusive migration, progressive taxation	21.8%
Center-Right	selective skilled migration, moderate tax cuts	31.8%
Far-Right	strict migration controls, corporate tax cuts	25.6%

Aligned by design: the more redistributive, the more open on migration. Control: share ranking each platform first.

Ranked first: change vs. control (pp)

Forced, left	Far-Right	+2.5**
Chose left	Center-Right	−4.0**
	Left	+2.3
	Center-Left	+1.8
Chose growth	Far-Right	+3.8***
	Center-Left	−2.4*

A fiscal shift moves votes only where the migration dimension allows (Roemer 1998). It **cannot** raise the far-right.

Taking stock: which reading of the model survives?

	Accuracy only	Dissonance only	**Hybrid**
Choice is confirmatory	✗[a]	✓	✓
The fact is accepted by opponents	✓	✗	✓
Forced exposure moves preferences (weakly)	✓	✓	✓
Backfire on one policy, for one group	✓[b]	✓[b]	✓

[a] Given equal expected accuracy. [b] Only via a model in which the fact argues against the policy.

Accuracy alone fails at the **choice**, dissonance alone at the **belief**: the cost of exposure governs what people read, accuracy what they believe.

Limits, and the next wave

This wave

- Priors not elicited before treatment: updating is estimated at the group level
- Unpaid knowledge questions: stated beliefs (Bullock et al. 2015; Prior et al. 2015)
- The readings of the backfire are not separable
- Platforms bundle fiscal and migration positions

Wave 2, pre-registered before fielding

- Incentivized priors and posteriors, with confidence
- Willingness to pay to read or avoid each article (prices φ)
- Expected accuracy of each headline, before choosing
- Backfire: what a policy does to inflation (model); random outlet attribution (identity); "universal" vs. citizens-only (migration)
- 2×2 platforms: fiscal $\times$ migration

Takeaways

1. In a fragmented environment, the first decision, **what to read**, does much of the polarizing work.
2. **Facts land**, even on opponents: no backfire on the fact.
3. Resistance sits at the step from **fact to policy**, and only on the most contested policy.
4. Persuasion on economics does not carry voters to left platforms that bundle redistribution with open migration.

> Correcting the facts inside each reality changes beliefs; it does not change the positions that define the sides.

oscar.barrerarodriguez@ucd.ie

Backup slides

Inference | updating by arm | selection | value of information | identity | platforms | references

Backup: registered tests, with few-cluster inference

	Wealth tax	Basic income	Min. wage	Top income tax	Inheritance tax	Corporate tax
T1 Forced	0.051	−0.022	0.049	0.003	0.087	0.165
p: bootstrap / R–W	.265 / .428	.385 / .670	.106 / .345	.963 / .955	.063 / .278	**.006 / .062**
T2 Optional	0.150	−0.037	0.054	0.016	0.129	0.173
p: bootstrap / R–W	**.001 / .005**	.189 / .272	.069 / .219	.588 / .578	**.006 / .009**	**.000 / .009**
T3 Choice	0.091	−0.096	−0.010	0.004	0.106	0.057
p: bootstrap / R–W	.010 / .213	.053 / .271	.813 / .959	.924 / .959	.038 / .175	.217 / .515
T2 − T1 (offering vs. imposing)	0.099	−0.015	0.005	0.013	0.042	0.008
p: bootstrap / R–W	.018 / .251	.689 / .971	.845 / .994	.823 / .994	.304 / .786	.844 / .994

Intent-to-treat effects in control-group SD units, $N = 4{,}524$; standard errors clustered by federal state ($G = 16$). Wild-cluster bootstrap-t: Webb weights, null imposed, 4,999 replications. Romano–Wolf (R–W) step-down across the six outcomes within each arm. Bold: both $p < 0.10$. Controls include strata, demographics and the 2021 vote. Index, T2 − T1: 0.051, bootstrap $p = 0.30$.

Backup: belief updating by arm

	N	Prior	Posterior	Signal	$\hat{\lambda}$	95% CI
Inflation for low-income singles (left article: 10.7%)						
T1 Forced	1,150	7.9	10.0	10.7	0.75	[0.74, 0.79]
T2 Read	1,002	7.9	10.0	10.7	0.75	[0.74, 0.82]
T2 Skipped (placebo)	149	7.9	7.9	10.7	0.00	[−0.43, 0.04]
T3 Chose left	490	7.9	10.0	10.7	0.75	[0.74, 0.82]
T3 Chose right (placebo)	643	7.9	8.0	10.7	0.04	[0.00, 0.21]

$\hat{\lambda}$ = (posterior − prior)/(signal − prior), medians; prior = control-group median; cluster-bootstrap CIs. The shortfall from 1 is rounding: the modal answer is 10. In the control group the most right-leaning fifth put low-income inflation at 7.0%, essentially overall inflation (6.9%); after forced reading they report 10.0%, like everyone else. For self-selected groups the control prior carries a selection term, so the argument rests on forced exposure by quintile.

Backup: chosen versus forced is not a causal contrast

For a group defined by its own choice c, e.g. those who chose the left article:

$$\beta_c = \underbrace{\mathbb{E}[y_i^0 \mid c] - \mathbb{E}[y_i^0]}_{\text{where the group started}} + \underbrace{\mathbb{E}[y_i - y_i^0 \mid c]}_{\text{how far it moved}}$$

- If the cost of exposure sorts respondents toward congenial articles, the first term is **positive** for left-choosers and **negative** for right-choosers.
- Given the sorting we observe, **any** updating rule produces "chosen > read > forced".
- Those who decline receive no message: their coefficient is the selection term alone.
- Controls, including the 2021 vote, remove selection on observables, not within them.

Backup: why choice needs a cost of exposure

Prior 8%, a trusted article says 10%: distance $d = 2$. With $\alpha = 0.75$ and $\kappa = 0.25$, $\lambda = 0.75$.

	Belief after	Cost of being wrong, $\alpha\,(\cdot)^2$	Cost of changing mind, $\kappa\,(\cdot)^2$	Total cost
Skip, or read and ignore	8	$0.75 \times 4 = 3.00$	0	**3.00**
Move halfway	9	0.75	0.25	1.00
Move $\lambda = \frac{3}{4}$ of the way	9.5	0.19	0.56	**0.75**
Move all the way	10	0	1.00	1.00

- Value of reading $= 3.00 - 0.75 = \alpha\lambda d^2 = 2.25$: reading is **never** worse than skipping.
- An article 4 points away is worth $\alpha\lambda \cdot 16 = 9$: the one that disagrees more has more to teach.
- Only a cost paid on reading itself makes the congenial article win: $(\alpha\lambda - \varphi)\, d^2$.

Backup: identity in the model

A wish to stand apart from the position the article argues for, $y_s = g_i \equiv g_i(\mu_i^{\text{post}})$:

$$\min_y \ \alpha_i (y - g_i)^2 + \kappa_i^{\text{pol}} (y - y_i^0)^2 - \eta_i (y - y_s)^2 \quad \Rightarrow \quad w_i = \frac{\alpha_i - \eta_i}{\alpha_i + \kappa_i^{\text{pol}} - \eta_i}$$

- w_i, the weight on the article's position, is **negative exactly when** $\eta_i > \alpha_i$ (given $\alpha_i + \kappa_i^{\text{pol}} > \eta_i$): identity outweighs accuracy.
- As a source discount: $\rho_i = \dfrac{\eta_i \kappa_i^{\text{pol}}}{\alpha_i (\alpha_i + \kappa_i^{\text{pol}} - \eta_i)} > 1$ exactly when $\eta_i > \alpha_i$.
- $\eta_i < 0$, affinity with the source, raises the weight: the most left-leaning fifth move *toward* a basic income.
- The data identify only the sign of $\eta_i - \alpha_i$, as the choice data identify only the sign of $\varphi_i - \alpha_i \lambda_i$.

Backup: a two-dimensional model of platform choice

Platform j has fiscal position f_j and migration position m_j; respondent i has ideal point (f_i, m_i):

$$U_{ij} = -w_f\,(f_j - f_i)^2 - w_m\,(m_j - m_i)^2 + \varepsilon_{ij}, \qquad \Delta U_{ij} - \Delta U_{ik} = 2\,w_f\,\delta\,(f_j - f_k)$$

1. A fiscal shift δ reorders platforms by their **fiscal positions alone**.
2. It changes a first choice only near indifference; with aligned platforms, restrictive voters rarely cross from right to left (Roemer 1998).
3. Under a leftward shift every platform gains on the far-right: **no one switches to it** from another platform.

So the far-right gain under forced exposure needs another channel: protest, or the salience of migration.

References (1/2)

Alesina, A., & Glaeser, E. (2004). *Fighting Poverty in the US and Europe: A World of Difference*. Oxford University Press.

Alesina, A., Miano, A., & Stantcheva, S. (2023). Immigration and redistribution. *Review of Economic Studies*, 90(1), 1–39.

Andre, P., Haaland, I., Roth, C., Wiederholt, M., & Wohlfart, J. (2026). Narratives about the macroeconomy. *Review of Economic Studies*.

Bénabou, R., & Tirole, J. (2006). Belief in a just world and redistributive politics. *Quarterly Journal of Economics*, 121(2), 699–746.

Bullock, J., Gerber, A., Hill, S., & Huber, G. (2015). Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, 10(4), 519–578.

Cavaillé, C., & Trump, K.-S. (2015). The two facets of social policy preferences. *Journal of Politics*, 77(1), 146–160.

Chopra, F., Haaland, I., & Roth, C. (2024). The demand for news: Accuracy concerns versus belief confirmation motives. *Economic Journal*, 134(661), 1806–1834.

Gentzkow, M., & Shapiro, J. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2), 280–316.

Gentzkow, M., & Shapiro, J. (2011). Ideological segregation online and offline. *Quarterly Journal of Economics*, 126(4), 1799–1839.

Guess, A., & Coppock, A. (2020). Does counter-attitudinal information cause backlash? Results from three large survey experiments. *British Journal of Political Science*, 50(4), 1497–1515.

Haaland, I., Roth, C., & Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, 61(1), 3–40.

References (2/2)

Kahan, D. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4), 407–424.

Kuziemko, I., Norton, M., Saez, E., & Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, 105(4), 1478–1508.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3), 831–870.

Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32(2), 303–330.

Prior, M., Sood, G., & Khanna, K. (2015). You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions. *Quarterly Journal of Political Science*, 10(4), 489–518.

Roemer, J. (1998). Why the poor do not expropriate the rich: An old argument in new garb. *Journal of Public Economics*, 70(3), 399–424.

Suhay, E. (2025). *Debating the American Dream: How Explanations for Inequality Polarize Politics*. Russell Sage Foundation.

Wood, T., & Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior*, 41(1), 135–163.

Zaller, J. R. (1992). *The Nature and Origins of Mass Opinion*. Cambridge University Press.