Polarization and Media Bias*

A. Arda Gitmez† Pooya Molavi‡

May 5, 2022

This paper presents a model of partisan media trying to persuade a sophisticated and heterogeneous audience. We base our analysis on a Bayesian persuasion framework where receivers have heterogeneous preferences and beliefs. We identify an intensive-versus-extensive-margin trade-off that drives the media's choice of slant: Biasing the news garners more support from the audience who follows the media but reduces the size of the audience. The media's slant and target audience are qualitatively different in polarized and unimodal (or non-polarized) societies. When the media's agenda becomes more popular, the media become more biased. When society becomes more polarized, the media become less biased. Thus, polarization may have an unexpected consequence: It may compel partisan media to be less biased and more informative.

*We would like to thank Daron Acemoglu, Ethan Bueno de Mesquita, Peter Buisseret, Mehmet Ekmekci, Drew Fudenberg, Emin Karagözoğlu, Anton Kolotilin, Parag Pathak, James Robinson, Mehdi Shadmehr, Iván Werning, Alex Wolitzky, David Yang, Adam Zelizer and Leifu Zhang for helpful discussions, and the participants at the MIT Theory Lunch, MIT Political Economy Lunch, NC State Theory Seminar, and Political Economy in the Chicago Area (PECA) Conference for their comments. An earlier version of this paper was previously circulated under the title "Media Capture: A Bayesian Persuasion Approach."

†Department of Economics, Bilkent University. E-mail: arda.gitmez@bilkent.edu.tr.

‡Northwestern University. E-mail: pmolavi@kellogg.northwestern.edu.

1 Introduction

American media tend to be partisan in their coverage of politics, and partisanship has increased in recent years.[1] Americans have started to take note: According to a Pew Research Survey in 2020, 79% of Americans believe “In presenting the news dealing with political and social issues... news organizations tend to favor one side.”[2] The partisanship has led to lower trust in media: The percentage of US adults who have at least some trust in the information from national news organizations has decreased from 76% in 2016 to 58% in 2021.[3]

Nevertheless, some citizens find it worthwhile to pay attention to the news—even in a landscape with biased media and low trust. Citizens might benefit from paying attention to a media source, as long as its slant is known and there is enough information in its coverage to potentially alter citizens’ choices. Whether to follow media is thus a strategic choice that depends on a variety of factors, among which are (i) the initial opinions of citizens, (ii) the citizens’ attitudes about the discussed policies, and (iii) the adopted slant of media sources.

Citizens are not the only strategic actors. Media, albeit partisan, are also strategic in their choice of slant. They consider the distribution of opinions and attitudes in society when deciding on their slant.[4] By their choice of slant, media effectively choose their audience. What is the partisan media’s adopted bias given the distribution of initial opinions and attitudes? Who pays attention to media, given the bias?

In this paper, we model the slant decision of partisan media and study how it changes with changes in society’s opinions and attitudes. The media face a sophisticated audience with diverse preferences and beliefs. The individuals in the audience choose whether to pay attention to the news given the media slant, and the media choose whom to target by choosing their editorial policies.

Our main findings are threefold. First, the optimal choice of slant is qualitatively different in *polarized* and *unimodal* (or non-polarized) societies. When society is unimodal in its opinions and attitudes, partisan media tailor their messaging to individuals already in the media’s camp. In contrast, partisan media attempt to convince skeptics when society is polarized. Second, when partisan media’s agenda becomes more popular among the citizens, the media become more biased.

Third, when the opinions and attitudes become more polarized, partisan media’s bias decreases. This finding suggests that polarization may have a silver lining: It may force partisan media to be more informative. If polarization is strong enough, the media may find it optimal

[1] See Groseclose and Milyo (2005), Larcinese, Puglisi and Snyder Jr (2011), Puglisi and Snyder Jr (2011), and Lott and Hassett (2014).

[2] https://www.pewresearch.org/pathways-2020/WATCHDOG_3/total_us_adults/us_adults

[3] https://www.pewresearch.org/fact-tank/2021/08/30/partisan-divides-in-media-trust-widen-driven-by-a-decline-among-republicans/

[4] There is an active literature focused on quantifying how exposure to partisan media shapes the opinions and attitudes of citizens (Levendusky, 2013; Prior, 2013) or political behavior (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017). Our approach in this paper is to take the opinions and attitudes as given and study how strategic media respond to the existing attitudes.

to reach beyond their base of support and tailor their messaging to the opinions and attitudes of their opponents. This result complements the central thesis of Oliveros and Várdy (2015) by offering a supply-side rationale for ideological moderation under political polarization.

At the model's heart is a trade-off faced by partisan media. The media want to persuade citizens to take a certain action (such as supporting a policy or voting for a candidate). Therefore, naturally, the media would like to spin the news in their desired direction (since such a spin increases the support among the citizens who pay attention to the news). However, the audience is sophisticated and realizes that the coverage has a spin. As a result, fewer people pay attention to partisan media because the bias makes the coverage less informative. The partisan media's choice is thus represented by an *intensive-versus-extensive-margin* trade-off. A positive spin leads to higher support among those who pay attention to the news (i.e., gains along the intensive margin) but fewer people paying attention to the news (i.e., losses along the extensive margin). The optimal policy balances these two effects.

The media's optimal policy is different in unimodal and polarized societies. A unimodal society has many citizens with moderate opinions and attitudes. Moderates are swayed by even small amounts of information, so they pay attention to the news even if it is significantly biased. As a result, the extensive margin is relatively less sensitive to changes in the spin. In equilibrium, the media place a positive spin on their coverage to the point where marginal gains along the intensive margin equal marginal losses along the extensive margin. Due to the positive spin of the news, negative news is less frequent and more informative. Consequently, even citizens initially supportive of the media's agenda pay attention to the news since the occasional news that goes against their prior changes their actions.

There are few moderates in highly polarized societies. Instead, there are two blocs of citizens with extreme opinions and attitudes: supporters who take the media's preferred action unless there is informative negative news and opponents who do not take the preferred action unless the news is positive and highly informative. The media must convince the opponents to listen to the news without alienating the supporters. The optimal strategy is to put a negative spin on the news so that the occurrence of positive news is an event rare enough to convince the opponents. Still, the negative news is abundant enough that they do not sway supporters' opinions. Partisan media thus end up reaching out to citizens with opposing views.

The trade-off between the two margins is also behind our comparative statics results. When media's platform becomes more popular, there are more supporters in society, and the audience of partisan media increases. As a result, improving the intensity of persuasion becomes more crucial, and the media place a more positive spin on their coverage. On the other hand, an increase in polarization in a unimodal society turns some moderates into radicals, who are not inclined to listen to the news unless if it is highly informative. Consequently, it becomes more important for the media to improve the extent of its reach by making the news informative enough for radicals. The outcome is a less positive spin and more informative media.

Related Literature. Our model is a Bayesian persuasion model à la Kamenica and Gentzkow (2011) with a heterogeneous audience. We incorporate both heterogeneous preferences (Wang, 2015; Alonso and Câmara, 2016a; Kolotilin, Mylovanov, Zapechelnyuk and Li, 2017; Bardhi and Guo, 2018; Chan, Gupta, Li and Wang, 2019; Arieli and Babichenko, 2019; Kerman, Herings and Karos, 2021; Sun, Schram and Sloof, 2022) and heterogeneous priors (Alonso and Câmara, 2016b; Laclau and Renou, 2017; Kosterina, 2021).[5] We present our theoretical results by introducing a new object, the *virtual density*, which summarizes the two dimensions of heterogeneity in a one-dimensional object. Our comparative statics results are based on identifying changes that lead to tractable changes in the virtual density. This approach is similar in spirit to the one in Kolotilin (2015) and Kolotilin, Mylovanov and Zapechelnyuk (2021). But unlike Kolotilin (2015), whose main focus is changes in welfare, we analyze how the optimal policy changes with parameters of the model. Our measure of popularity is complementary to that in Kolotilin, Mylovanov and Zapechelnyuk (2021), and our measure of polarization is a novel one. Sun, Schram and Sloof (2022) derive comparative statics results with respect to the sender's preferences and the voting rule in a voting environment with heterogeneous preferences. Our comparative statics result complement theirs by focusing on changes with respect to the audience's characteristics.

Our findings contribute to the theory of media bias (see Gentzkow, Shapiro and Stone (2015) for a survey). Unlike Mullainathan and Shleifer (2005) and Bernhardt, Krasa and Polborn (2008), in our model, citizens choose media sources purely on informational grounds. Gentzkow and Shapiro (2006), Burke (2008) and Chan and Suen (2008) propose models where slant can arise when the media are not inherently biased towards an outcome. As in Baron (2006) and Duggan and Martinelli (2011), we study an environment where the media are inherently biased, with their main objective the persuasion of citizens. Our focus is on how the intensity and direction of media bias respond to changes in the audience's characteristics.

A large recent literature has focused on quantifying the extent of polarization (DiMaggio, Evans and Bryson, 1996; Glaeser and Ward, 2006; Ansolabehere, Rodden and Snyder Jr, 2006; McCarty, Poole and Rosenthal, 2006; Fiorina and Abrams, 2008; Abramowitz and Saunders, 2008; Gentzkow, 2016). The causes and consequences of polarization have also gathered attention in the popular press (Sunstein, 2009; Klein, 2020). But relatively little attention has been paid to the question of how media respond to polarization. Another active research area studies how the changing media landscape—partisan or not—affects the patterns of polarization in society

[5] Also related is the literature on information design, which studies the optimal information structure in a game to be played among multiple players (Bergemann and Morris, 2019; Taneva, 2019; Mathevet, Perego and Taneva, 2020; Inostroza and Pavan, 2022).

(Campante and Hojman, 2013; Flaxman et al., 2016; Bail et al., 2018). Our focus here is the opposite: understanding how the media landscape is affected by polarization.

A possible interpretation of our model is as one in which the media truthfully report the news, but politicians suppress or overturn them. Our findings can then be interpreted as follows: (i) as the politician becomes more popular, she chooses to suppress more news, (ii) as society becomes more polarized, the politician allows for more informative media, and (iii) in a highly polarized society, the politician allows negative news to be published, in order to convince her opponents without alienating her supporters. Through this lens, our model contributes to the literature on media capture (Besley and Prat, 2006; Corneo, 2006; Petrova, 2008; Prat, 2015), information manipulation by autocratic regimes (Edmond, 2013; Shadmehr and Bernhardt, 2015; Guriev and Treisman, 2020), and media freedom (Egorov, Guriev and Sonin, 2009; Gehlbach and Sonin, 2014; Boleslavsky, Shadmehr and Sonin, 2021).

The crucial assumption in the Bayesian persuasion literature is commitment by the sender. In our model, partisan media can commit to a strategy, which is observable by all receivers. This assumption can be defended on several grounds. First of all, in our setup, persuasion satisfies the credibability assumption of Lin and Liu (2021).[6] Second, the media's chosen strategy can viewed as an "editorial policy," which describes the general attitude of a media source, with the details of the coverage to be decided by reporters and editors (Gehlbach and Sonin, 2014). Finally, the outcome under commitment can be seen as a benchmark, which describes the best-case scenario for the sender. Under this interpretation, our results characterize an "ideal media landscape" for a politician in a heterogeneous society. Our results show that, in a highly polarized society, the politician may indeed benefit from media that frequently publish negative news about the politician. Intuitively, this is the only way the politician can garner the opponents' support. The rare, but convincing, occurrence of positive news about politician is the most effective way for politicians to convince skeptical citizens (Chiang and Knight, 2011).

2 Setup

2.1 The Model

There are two types of agents: a sender (female) and a unit measure of receivers (males), indexed by $r \in [0,1]$. The sender wants to persuade the receivers to support a policy she is proposing. Before persuading the receivers, the sender learns whether the policy is good for the receivers. But the receivers do not learn this information until after they have decided on whether to support the policy.

[6]In particular, we can allow for undetectable deviations by the sender. Since the sender's payoff in our model is additively separable, there is no profitable deviation that gives the same message distribution as the optimal policy. It should be noted that with heterogeneous priors, the set of undetectable deviations is different for each agent. In such a case, one has to define the undetectable deviations from the perspective of the sender.

There is an underlying state of the world: $\theta \in \{0, 1\}$. We call the $\theta = 1$ state the "good" state. When the state is good, the sender's proposed policy is beneficial to the receivers. The $\theta = 0$ state is the "bad" state where the policy is not beneficial to the receivers. Each receiver has to decide on the extent to which to support the policy. We denote by $a_r \in [0, 1]$ receiver r's degree of support for the policy. Receiver r's payoff when he chooses action a_r and the state is θ is given by

$$u_r(a_r, \theta) = a_r(\theta - c_r), \tag{1}$$

where $c_r \in [0, 1]$ is receiver r's cost of supporting the policy. If receiver r knew the state, he would give the policy his full support (i.e., $a_r = 1$) in the good state and no support (i.e., $a_r = 0$) in the bad state.

The sender wants to maximize the support from the receivers. Her payoff when receiver r provides support a_r and the state is θ is given by

$$u_s(\{a_r\}_r) = \int_0^1 a_r dr. \tag{2}$$

When the state is good, the sender and receivers have common interests, whereas when the state is bad, their interests are opposed. We denote the sender's prior that the state is good by $p_s \in (0, 1)$.

Since the receivers do not observe the state, they can only act based on their beliefs. The sender can influence those beliefs (and the resulting actions) by sending informative messages. To simplify the analysis, we assume that the sender can commit to a public communication strategy $\sigma : \{0, 1\} \to \Delta(M)$, where $\sigma(\theta)[m]$ is the probability that public message $m \in M$ is generated when the state is θ. The communication strategy represents the editorial policies of a collection of partisan *media* controlled by the sender and used by her to influence the views of the receiver population.

The receivers are *heterogeneous* both in their preferences and their prior beliefs. The heterogeneity of priors captures the idea that even people with identical payoffs may have different perspectives about the likelihood that a given policy will succeed. We let p_r denote receiver r's prior that the state is good and, let $f(c, p)$ denote the joint density of costs and priors in the population of receivers. We take f as a primitive of the model and study how changing the distribution affects the sender's optimal policy. We assume that f is common knowledge and continuously differentiable and bounded over its support.

We can partition the set of receivers into *ex-ante supporters* and *ex-ante opponents* of the proposed policy. Any receiver r with $p_r \geq c_r$ is an ex-ante supporter, who would support the proposed policy without additional information. Ex-ante supporters are in the bottom right half of Figure 1. Conversely, the receivers in the top left half of Figure 1 are ex-ante opponents, who have high costs and unfavorable priors and would not support the policy without additional information.

The heterogeneity of perspectives poses a challenge for a sender who wants to garner broad support for her proposed policy. Convincing different receivers with different preferences and

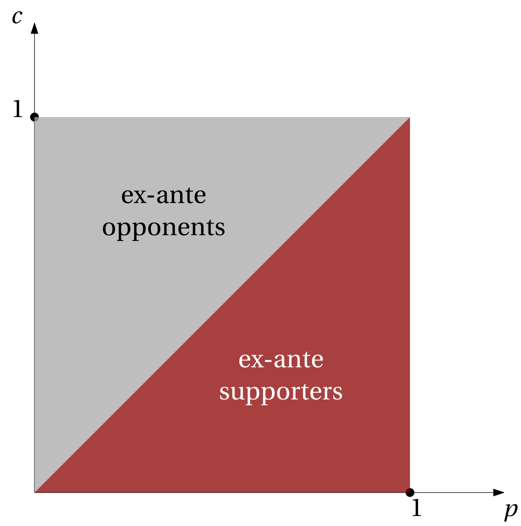

Figure 1. The ex-ante supporters and opponents of the policy.

beliefs requires different communication strategies. Yet, communication is public, so the sender cannot tailor her messaging strategy to the diverse perspectives of receivers. The main characterization result of the paper concerns the optimal way of resolving the inherent tension in convincing different segments of the population.

Timing. The timing of the communication game is as follows:

1. The prior and cost of each receiver is drawn, and each receiver r observes (p_r, c_r).
2. The sender commits to a strategy σ, which is observed by each receiver.
3. The state is realized, and the sender sends the message drawn according to σ.
4. Each receiver r updates his prior and chooses an action a_r.
5. Payoffs are realized.

The solution concept we adopt is the Perfect Bayesian Equilibrium.

2.2 An Equivalent Representative-Receiver Problem

The fact that the sender is communicating with a population of heterogeneous receivers complicates her problem. However, the sender's optimal strategy can be found by solving a related persuasion problem with a *representative receiver* whose prior coincides with the sender's prior.

The key simplification comes from Proposition 1 of Alonso and Câmara (2016b). Consider receivers r and r' with priors p_r and $p_{r'} = p_s$. Since the two receivers observe the same (public) message, their posteriors are related through the following expression:

$$\mu_r = \frac{\mu_{r'} \frac{p_r}{p_{r'}}}{\mu_{r'} \frac{p_r}{p_{r'}} + (1 - \mu_{r'}) \frac{1-p_r}{1-p_{r'}}}, \tag{3}$$

where μ_r and $\mu_{r'}$ denote the posteriors of r and r', respectively.[7] This coupling of posteriors holds regardless of the communication strategy employed by the sender. It uniquely pins down the posterior μ_r of *every* receiver r as a function of the posterior of receiver r'—who will be our representative receiver.

Receiver r takes action $a_r = 1$ if and only if his posterior that the state is good is at least as large as his cost of action; that is, if and only if

$$c_r \leq c(\mu_s, p_r) \equiv \frac{\mu_s \frac{p_r}{p_s}}{\mu_s \frac{p_r}{p_s} + (1-\mu_s)\frac{1-p_r}{1-p_s}}, \tag{4}$$

where μ_s denotes the posterior of the representative receiver (who has the same prior as the sender). The payoff to the sender is given by the share of receivers who take the good action:

$$v(\mu_s) = \int_0^1 \int_0^{c(\mu_s,p)} f(p,c)\,dc\,dp. \tag{5}$$

The sender's problem is thus equivalent to a standard Bayesian persuasion problem with a representative receiver. The sender and the receiver share the common prior p_s that the state is good. The payoff to the sender when she induces a posterior of μ_s for the representative receiver is given by $v(\mu_s)$, defined in equation (5). Following Kamenica and Gentzkow (2011), we refer to $v(\mu_s)$ as the sender's *value function*. Whenever there is no risk of confusion, we drop the s subscript and simply write $v(\mu)$ for the value to the sender of inducing posterior μ for the representative receiver.

Figure 2 shows a useful graphical representation of the value function. The value to the sender from inducing posterior μ is given by the measure of receivers r who take the $a_r = 1$ action (the area shaded in red in the figure). This measure depends on the distribution of costs and priors in the population as well as the induced posterior μ. As μ increases, more and more receivers support the proposed policy, expanding the shaded area in the figure and increasing the sender's payoff.

The value function has several useful properties. First, $v(\mu)$ is increasing in μ. Inducing a higher posterior for the representative receiver results in a higher posterior for every receiver, thus increasing the share of receivers who take the $a = 1$ action. Second, $v(0) = 0$ and $v(1) = 1$. When the representative receiver is certain that the state is bad, so is every other receiver. Therefore, every receiver takes the $a = 0$ action. Likewise, when the representative receiver is certain that the state is good, every other receiver is also certain that the state is good and takes the $a = 1$ action. Finally, $v(\mu)$ is differentiable in μ due to the differentiability of f.

The value function can thus be seen as a differentiable cumulative distribution function. We let $h(\mu) \equiv v'(\mu)$ denote the corresponding density and refer to it as the *virtual density* of the persuasion problem with heterogeneous receivers. The virtual density has an intuitive interpretation: $h(\mu)$ is the density of receivers who are indifferent between taking the two actions whenever the representative receiver's posterior is equal to μ.

[7] Throughout the paper, we use posterior to mean subjective posterior probability of state $\theta = 1$ given an agent's information.

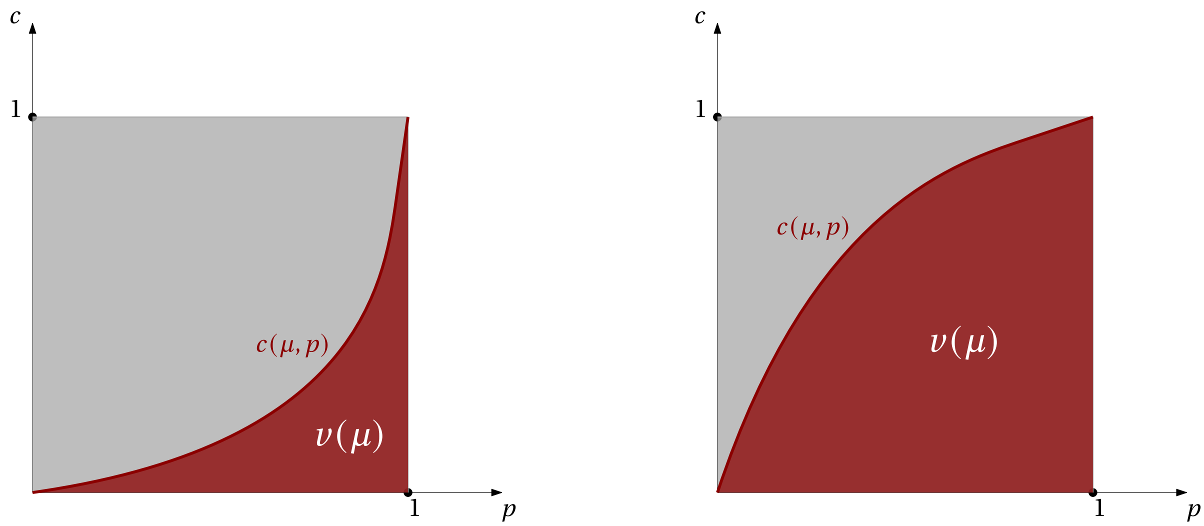

Figure 2. The value function: $\mu < p_s$ in the left panel and $\mu > p_s$ in the right panel.

2.3 Single-Peaked Distributions

The solution to the sender's optimal persuasion problem takes a particularly simple form when the distribution of receiver types satisfies the following condition:

Definition 1 (single-peaked density)**.** The virtual density $h(\mu)$ is *single-peaked* if there exists some $\tilde{\mu} \in [0, 1]$ such that $h'(\mu) > 0$ for all $\mu < \tilde{\mu}$ and $h'(\mu) < 0$ for all $\mu > \tilde{\mu}$.

Single-peakedness is an assumption on the joint distribution of receivers' costs and prior beliefs. It requires a large share of receivers to have moderate preferences and beliefs, with fewer and fewer people having extreme preferences or beliefs. We thus consider single-peaked virtual densities to be representative of *unimodal* societies.

The significance of Definition 1 rests on the following observation: When the virtual density is single-peaked, the sender's value function is first convex and then concave. Figure 3 illustrates the value function in this case. Corollary 2 of Kamenica and Gentzkow (2011) implies that, under

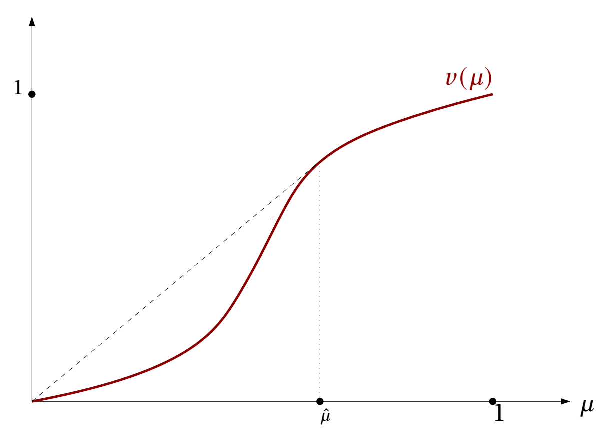

Figure 3. The value function under the assumption that the virtual density is single-peaked.

the sender's optimal strategy, the representative receiver's posterior only takes two values. In particular, we have the following characterization of the optimal strategy when the virtual density is single-peaked:

Theorem 1. *If the virtual density is single-peaked, the optimal strategy uses only two messages.*

We maintain the assumption of single-peakedness in the next two sections. We do so in part for tractability. However, single-peaked distributions also constitute a natural and widely used class of distribution functions. In Section 5, we show that optimal strategy in the case where the virtual density is instead single-dipped is the mirror image of the optimal strategy in the single-peaked case.

Whether the virtual density is single-peaked only depends on the distribution of types, f, and the sender's prior, p_s. In the remainder of this subsection, we find a set of easy-to-check sufficient conditions for the virtual density to the single-peaked. If receivers have a common prior which coincides with the sender's prior, then the single-peakedness of the virtual density is equivalent to the single-peakedness of the density of costs:

Proposition 1. *Suppose $p_r = p_s$ for all r. The virtual density $h(\mu)$ is single-peaked in μ if and only if the density of costs $f(c)$ is single-peaked in c.*

If receivers have a common cost, on the other hand, then the single-peakedness of the virtual density is implied by a condition that is weaker than the log-concavity of the density of priors:

Proposition 2. *Suppose $c_r = c \in (0,1)$ for all r. The virtual density $h(\mu)$ is single-peaked if the density of priors $f(p)$ is strictly positive for all $p \in (0,1)$ and satisfies*

$$\frac{d^2}{dp^2}\log f(p) < 2(\gamma-1)^2 \min\left\{1, \frac{1}{\gamma^2}\right\} \quad \text{for all } p \in (0,1), \tag{6}$$

where $\gamma \equiv \frac{1-c}{c}\frac{1-p_s}{p_s} \geq 0$.

The following corollary of Proposition 2 is a straightforward consequence of the facts that the left-hand side of equation (6) is negative if $f(p)$ is log-concave, while its right-hand side is always non-negative:[8]

Corollary 1. *Suppose $c_r = c$ for all r. The virtual density $h(\mu)$ is single-peaked in μ if the density of priors $f(p)$ is strictly log-concave in p.*

3 Optimal Policy

3.1 Never-supporters, Always-supporters, and Compliers

In light of Theorem 1, we can assume without loss that the sender uses only two messages. We label the messages $m \in M = \{0,1\}$, with $m = 1$ the "good" message, which is suggestive of $\theta = 1$, and

[8]See Bagnoli and Bergstrom (2005) for a list of well-known distributions satisfying log-concavity.

$m = 0$ the "bad" message, which is suggestive of $\theta = 0$. The sender's strategy can be represented by a pair of numbers:

$$\sigma = (\sigma^0, \sigma^1) \in [0,1]^2,$$

where $\sigma^\theta \equiv \sigma(\theta)[m = 1]$ is the probability of sending the good message in state $\theta \in \{0,1\}$. Throughout, we assume without loss of generality that $\sigma^1 \geq \sigma^0$.

The media's optimal strategy thus partitions the set of receivers into three groups:

1. The *never-supporters*, who choose $a = 0$ regardless of the message m.

2. The *compliers*, who choose $a = m$.

3. The *always-supporters*, who choose $a = 1$ regardless of the message m.

The never-supporters set $a = 0$ even if they receive the good message: They are not convinced by the media because their initial beliefs are too pessimistic relative to their costs. The always-supporters set $a = 1$ even if they receive the bad message because of their optimistic priors and low costs. The compliers are the the most interesting group. They pay attention to the news and adjust their actions in response to what they learn.

Our next result characterizes the set of never-supporters, compliers, and always-supporters as a function of the strategy followed by the sender:

Proposition 3. *Given (σ^0, σ^1), where $\sigma^1 \geq \sigma^0$, let:*

$$\underline{p}(c) \equiv \frac{c\sigma^0}{c\sigma^0 + (1-c)\sigma^1}, \tag{7}$$

$$\overline{p}(c) \equiv \frac{c(1-\sigma^0)}{c(1-\sigma^0) + (1-c)(1-\sigma^1)}. \tag{8}$$

Receiver r is a never-supporter if $p_r < \underline{p}(c_r)$, a complier if $p_r \in [\underline{p}(c_r), \overline{p}(c_r))$, and an always-supporter if $p_r \geq \overline{p}(c_r)$.

Figure 4 illustrates the partition of receivers. When the media are fully uninformative (i.e., $\sigma^1 = \sigma^0$), then $\underline{p}(c) = \overline{p}(c) = c$ for all $c \in [0,1]$, and no receiver is a complier. The sets of always-supporters and never-supporters then coincide with the sets of ex-ante supporters and ex-ante opponents, respectively. The lightly shaded region of Figure 4 then disappears, and the figure reduces to Figure 1. As the media become more informative, the set of compliers grows at the expense of the always-supporters and never-supporters. When the media are fully informative (i.e., $\sigma^1 = 1$ and $\sigma^0 = 0$), every receiver is a complier, and the lightly shaded area comprises the entirety of the unit square.

3.2 Intensive and Extensive Margins

By choosing its strategy, the sender effectively chooses who is a complier and how frequently the compliers see the good message. The sender wants to turn as large a share of the receivers as

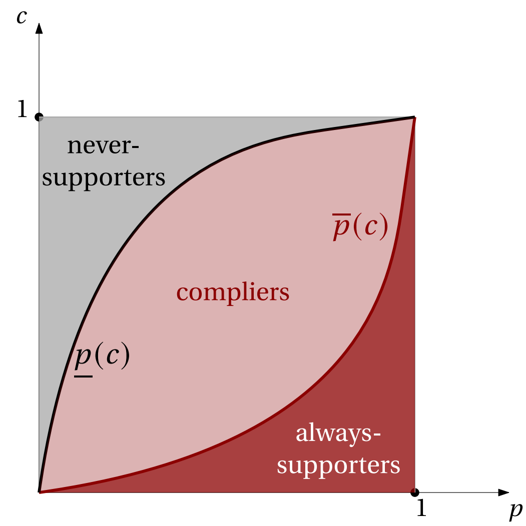

Figure 4. The set of never-supporters, compliers, and always-supporters given policy (σ^0, σ^1).

possible into compliers (and always-supporters) and send the good message as often as possible. However, she faces a trade-off between these two objectives.

One can view the sender's trade-off as one involving improvements along intensive and extensive margins. The sender's expected payoff from following strategy (σ^0, σ^1) is given by

$$\underbrace{\left(\int_0^1 \int_{\underline{p}(c)}^{\overline{p}(c)} f(p,c)dpdc\right)}_{\text{measure of compliers}} \underbrace{\left(p_s\sigma^1 + (1-p_s)\sigma^0\right)}_{\text{likelihood of the good message}} + \underbrace{\int_0^1 \int_{\overline{p}(c)}^1 f(p,c)dpdc}_{\text{measure of always-supporters}}, \tag{9}$$

where $\underline{p}$ and $\overline{p}$ are given by (7) and (8), respectively. The sender wants to increase the likelihood of the good message. Doing so increases the probability that the compliers support the proposed policy, thus allowing the sender to gain support along the *intensive margin*. But increasing the likelihood of the good message makes the media less informative. The decrease in the informativeness of the media turns some compliers into never-supporters and always-supporters. If many more compliers become never-supporters than always-supporters, the sender loses support along the *extensive margin*. Figure 5 illustrates this trade-off.

3.3 Who Follows the Media?

We now turn to the question of how the sender optimally resolves her trade-off.

Proposition 4. *If the virtual density is single-peaked, under the optimal strategy:*

1. *The bad message fully reveals the bad state.*

2. *The ex-ante supporters are all compliers.*

The argument for the Proposition can be best understood by studying Figure 3. When the virtual density is single-peaked, optimal persuasion involves moving the representative receiver's

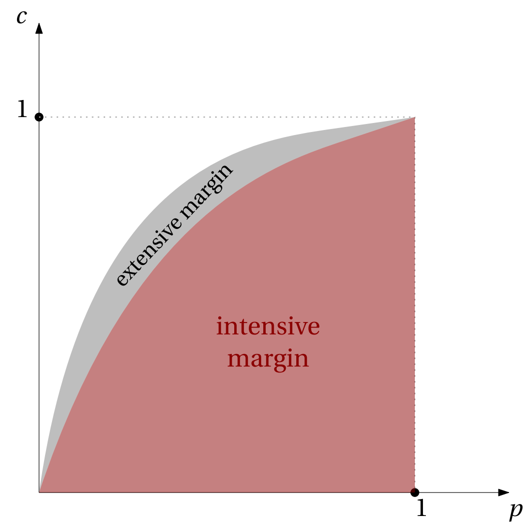

Figure 5. The trade-off between gaining support along the intensive and extensive margins for the $\sigma^1 = 1$ case. Choosing a higher σ^0 leads to a higher likelihood of the good message (gain along the intensive margin) but fewer compliers (loss along the extensive margin).

posterior to one of two points: $\mu = 0$ and $\mu = \hat{\mu}$.[9] Inducing the $\mu = 0$ posterior requires the bad message to fully reveal the bad state. But then even the most ardent ex-ante supporters find it worthwhile to follow the media on the off chance that the state is revealed to be bad. Therefore, every ex-ante supporter is a complier, and there are no always-supporters. Figure 6 illustrates the partition of receivers under the optimal persuasion strategy.

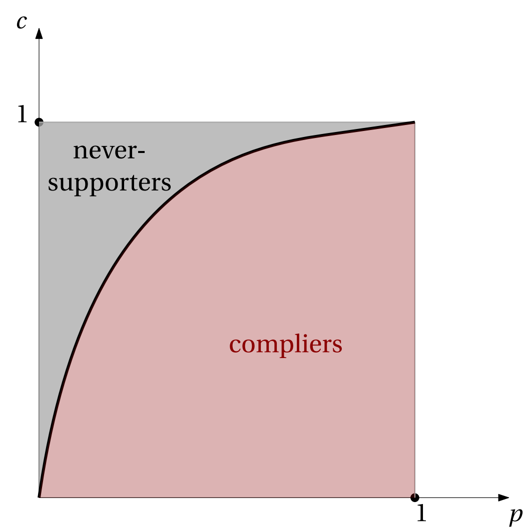

Figure 6. The set of never-supporter and compliers when the virtual density is single-peaked.

The intuition behind the result is as follows. Moderate receivers with middling costs and priors are the ones most inclined to follow the media since their behavior is sensitive even to messages with little information about the state. When the virtual density is single-peaked, there are many

[9]When the sender's prior is high enough that $p_s > \hat{\mu}$ in Figure 3, the optimal strategy reveals no information, and any strategy that satisfies $\sigma^0 = \sigma^1$ is optimal. In such a case, we choose the strategy $\sigma^0 = \sigma^1 = 1$ and assume that a receiver's posterior following the (zero probability) bad message is given by $\mu = 0$.

more moderates than are receivers with extreme preferences or beliefs. Therefore, the sender has much more to gain by increasing the likelihood of the good message and the probability of support from the moderates than it has to lose from turning ex-ante opponents off. Figure 5 depicts the trade-off faced by the sender in this case. Under the optimal policy, the marginal gain from having more compliers—the measure of receivers in the grey sliver—equals the expected loss of support from compliers—the likelihood that receivers in the red region support the policy.

4 Media Bias in Unimodal Societies

The media are *biased* if they send the good message when the state is bad or send the bad message when the state is good. Recall that, when the virtual density is single-peaked, optimal persuasion entails sending the good message when the state is good. Therefore, media bias is conveniently summarized in the single-peaked case by the probability σ^0 of sending the good message when the state is bad. We use the following notion of media bias in this case:

Definition 2. Consider single-peaked virtual densities h_1 and h_2 with the corresponding optimal strategies $\sigma_1 = (\sigma_1^0, \sigma_1^1)$ and $\sigma_2 = (\sigma_2^0, \sigma_2^1)$ for the sender. The media are *more biased* given h_1 than given h_2 if $\sigma_1^0 \geq \sigma_2^0$.

In this section, we characterize how changes in the primitives of the model affect the extent of media bias.

4.1 Changes in Popularity

We first study how shifts in the distribution of costs and priors affect the sender's optimal strategy and the resulting media bias. We use the following partial orders on the set of distributions:

Definition 3. Consider probability density functions f_1 and f_2 with the corresponding cumulative distribution functions F_1 and F_2. We say f_1 is larger than f_2 in the *hazard rate order* if

$$\frac{f_1(x)}{1 - F_1(x)} \leq \frac{f_2(x)}{1 - F_2(x)} \qquad \text{for all } x. \tag{10}$$

Definition 4. Consider probability density functions f_1 and f_2 with the corresponding cumulative distribution functions F_1 and F_2. We say f_1 is larger than f_2 in the *reversed hazard rate order* if

$$\frac{f_1(x)}{F_1(x)} \geq \frac{f_2(x)}{F_2(x)} \qquad \text{for all } x. \tag{11}$$

Hazard rate and reversed hazard rate orders are related to two well-known partial orders on distributions: They are more complete than the monotone-likelihood ratio property (MLRP) but less complete than first-order stochastic dominance (FOSD).[10] If f_1 is larger than f_2 in the hazard rate or reversed hazard rate orders, then it also first-order stochastically dominates f_2.

Our next result establishes that an increase in the support for the policy increases media bias:

[10] See, for instance, Shaked and Shanthikumar (2007, Section 1.B).

Theorem 2. *Let h_1 and h_2 be two single-peaked virtual densities. If h_1 is smaller than h_2 in the reversed hazard rate order, the media are more biased given h_1 than h_2.*

The result states that media are more biased when the policy has a larger ex-ante support. A *smaller* virtual density indicates a *more popular* policy. With a more popular policy, a larger fraction of receivers support the policy absent any persuasion by the sender. Therefore, the optimal strategy involves less information transmission and more reliance on the receivers' priors, and the media are thus less informative and more biased under the sender's optimal strategy.

Why is the reversed hazard rate order the right notion? The answer is best seen by examining the trade-off the sender faces between improving along the intensive and extensive margins. Suppose the sender reduces the probability of the good message when the state is bad by $d\sigma^0$. The result is a reduction in the posterior of the representative receiver following the observation of the good signal by some $d\mu$ (which is related to $d\sigma^0$). The change makes the media more informative and increases the measure of compliers by $h(\mu)d\mu$, an improvement along the extensive margin. But this gain comes at the expense of reducing the probability that the inframarginal compliers support the policy by $H(\mu)d\sigma^0$, a loss along the intensive margin. Therefore, $h(\mu)/H(\mu)$ is a measure of the net benefit from making the media more informative. When $h_1(\mu)/H_1(\mu) \leq h_2(\mu)/H_2(\mu)$, the sender gains relatively more by making the media more informative when the virtual density is h_2 than when it is h_1. Consequently, the media are more biased when the virtual density is h_1.

We next study two special cases of Theorem 2, where receivers have common costs or priors. With a common prior, the reversed hazard rate order on virtual densities reduces to the reversed hazard rate order on the distributions of costs:

Corollary 2. *Suppose $p_r = p_s$ for all r, and consider two single-peaked densities $f_1(c)$ and $f_2(c)$ for the receivers' cost of support. If f_1 is smaller than f_2 in the reversed hazard rate order, the media are more biased given f_1 than f_2.*

Intuitively, the condition identified in Corollary 2 means that the costs are lower under f_1 than they are under f_2. This translates into a more popular policy. But when the proposed policy is more popular, the gains to the sender from persuasion are lower. Therefore, the media are less informative and more biased when receivers have lower costs of support.

If the receivers have common costs, on the other hand, a larger distribution of priors in the hazard rate order leads to a smaller virtual density in the reversed hazard rate order. The intuition for why we need the densities of priors to be ordered in the hazard rate order—and not the *reversed* hazard rate order as in Corollary 2—is as follows: *Higher* priors act like *lower* costs, making the policy more popular and lowering the gains from persuasion and the resulting media bias. The following corollary to Theorem 2 formalizes this intuition:

Corollary 3. *Suppose $c_r = c$ for all r, and consider two distributions of priors $f_1(p)$ and $f_2(p)$ that are both strictly positive for all $p \in (0,1)$ and satisfy condition (6). If f_1 is larger than f_2 in the hazard rate order, the media are more biased given f_1 than f_2.*

4.2 Polarization

We next examine how increased polarization changes media bias. We measure the extent of polarization using a novel partial order on probability distributions:

Definition 5. Consider single-peaked densities f_1 and f_2 supported on a common compact set and satisfying

$$f_2(x) = \frac{(f_1(x))^\alpha}{\kappa} \qquad \text{for all } x, \tag{12}$$

some $\alpha > 0$, and a normalization constant $\kappa > 0$. If $0 < \alpha \leq 1$, then f_2 is *more polarized* than f_1. If $\alpha \geq 1$, then f_2 is *less polarized* than f_1.

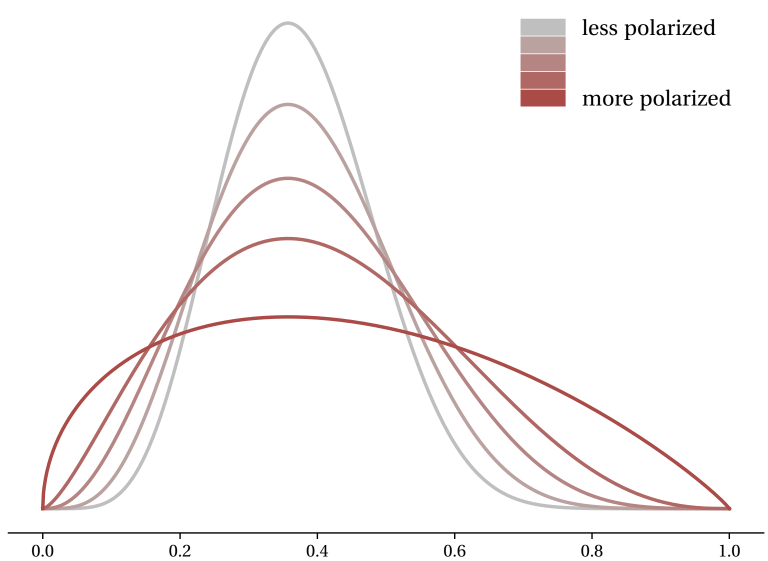

Figure 7. The polarization order on single-peaked densities.

The partial order has an intuitive interpretation. Consider densities f_1 and f_2 satisfying (12) for some $\alpha > 1$. Going from f_1 to f_2 moves mass from parts of the distribution that initially have smaller mass to parts with larger initial mass. In other words, f_2 looks like f_1, but with higher peaks and deeper troughs. On the other hand, since f_1 is single-peaked, most of its mass is concentrated around its peak. Therefore, f_2 has even more mass in the center and even less mass in the periphery than f_1; that is, f_2 is less polarized than f_1. Figure 7 illustrates the probability density functions for a set of single-peaked Beta distributions that are ranked in the polarization order.

It is instructive to consider the extremes of equation (12). In the $\alpha \to \infty$ limit, f_2 becomes a point mass at the mode of f_1. Therefore, for any distribution f_1 with a unique mode, the degenerate distribution f_2 with a point mass on the mode of f_1 is less polarized than f_1. Conversely, in the $\alpha \to 0$ limit, f_2 becomes the uniform distribution on the support of f_1. Thus, the uniform distribution is more polarized than any single-peaked distribution with the same support.

Members of many parametric families of distributions can be ordered in the polarization order. Two examples follow:

Example 1. Consider two single-peaked Beta distributions:

$$f_1 = \text{Beta}(\alpha_1, \beta_1),$$
$$f_2 = \text{Beta}(\alpha_2, \beta_2),$$

where $\frac{\alpha_2 - 1}{\alpha_1 - 1} = \frac{\beta_2 - 1}{\beta_1 - 1} = \alpha$ for some $\alpha \geq 0$. If $\alpha \leq 1$, then f_2 is more polarized than f_1, while if $\alpha \geq 1$, then f_2 is less polarized than f_1. In particular, any two symmetric and single-peaked Beta distributions are ranked according to the polarization partial order.

Example 2. Consider the following truncated normal distributions on $[0, 1]$:

$$f_1 = \text{TruncatedNormal}(\mu, \sigma_1^2),$$
$$f_2 = \text{TruncatedNormal}(\mu, \sigma_2^2).$$

If $\sigma_2^2 \geq \sigma_1^2$, then f_2 is more polarized than f_1, while if $\sigma_2^2 \leq \sigma_1^2$, then f_2 is less polarized than f_1.

Our next result establishes that polarization decreases media bias:

Theorem 3. *Let h_1 and h_2 be two single-peaked virtual densities. If h_1 is more polarized than h_2, then media are less biased given h_1 than h_2.*

Theorem 3 suggests a striking consequence of polarization: It tends to reduce the partisan media's bias. Intuitively, an increase in polarization turns some ex-ante moderates into ex-ante extremists. Under the optimal policy, then, there are fewer compliers and more never-supporters. As a result, the intensive margin becomes less important. The change in the trade-off incentivizes partisan media to reduce its bias and increase its outreach. Indeed, one may expect the segregation in news consumption to *decrease* as polarization increases, even with the existence of partisan media.

This result suggests a potential explanation for (Gentzkow and Shapiro, 2011, p.1819)'s finding that "There is no evidence that ideological segregation on the Internet has increased. If anything, segregation has declined as the Internet news audience has grown." It is also consistent with Flaxman, Goel and Rao (2016)'s finding that social networks and search engines are "associated with an increase in an individual's exposure to material from his or her less preferred side of the political spectrum."

Theorem 3 describes the consequences of polarization for media bias while maintaining the assumption that the society is unimodal, and so, the virtual density is single-peaked. In the next section, we study persuasion in highly polarized societies in which there are more people in the extremes of preference and belief distribution than are at its center.

5 Highly Polarized Societies

Throughout this section, we study the properties of the optimal persuasion strategy when the virtual density is the polar opposite of single-peaked:

Definition 6 (single-dipped density)**.** The virtual density $h(\mu)$ is *single-dipped* if there exists some $\tilde{\mu} \in [0,1]$ such that $h'(\mu) < 0$ for all $\mu < \tilde{\mu}$ and $h'(\mu) > 0$ for all $\mu > \tilde{\mu}$.

In a society with a single-dipped virtual density, there are fewer moderates than those with extreme preferences or beliefs. We therefore consider single-dipped virtual densities to be representative of *highly polarized* societies (DiMaggio, Evans and Bryson, 1996; Fiorina and Abrams, 2008).

When the virtual density is single-dipped, the sender's value function is first concave and then convex. Our next result shows that, as a result, the optimal persuasion strategy is qualitatively different in a highly polarized society compared to a unimodal one:

Proposition 5. *If the virtual density is single-dipped, the optimal strategy uses only two messages. Under the optimal strategy:*

1. *The good message perfectly reveals the good state.*
2. *The ex-ante opponents are all compliers.*

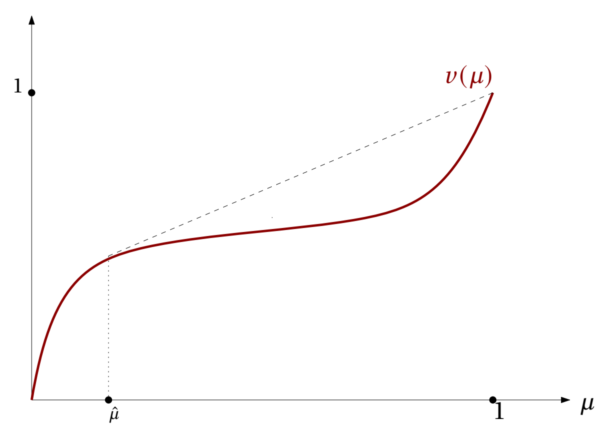

Figure 8. The value function under the assumption that the virtual density is single-dipped.

The argument for the proposition is easiest to see by examining the sender's value function. Figure 8 illustrates the value function in this case. When the distribution is single-dipped, the optimal policy induces only two values for the posterior of the representative receiver, one of which is

$\mu = 1$.[11] For the good message to induce the $\mu = 1$ posterior, it must perfectly reveal the good state. But since the good message is fully revealing, even the most vehement ex-ante opponents find the good message informative enough to sway their decision. Therefore, all ex-ante opponents are compliers, and there are no never-supporters. Figure 9 illustrates the partition of receivers under the optimal policy.

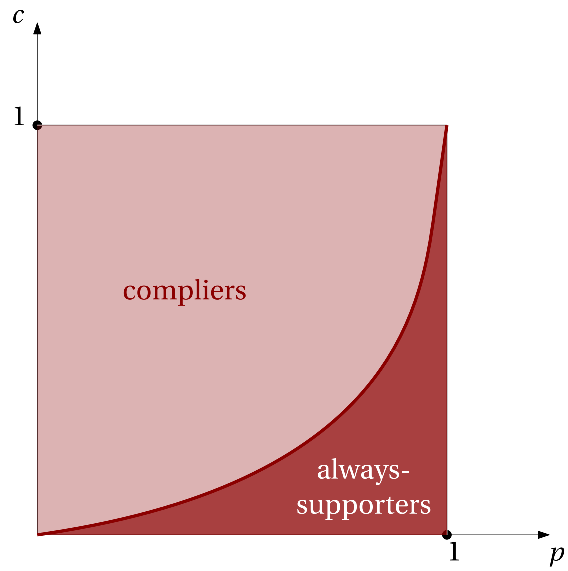

Figure 9. Receivers under the optimal policy when the virtual density is single-dipped.

We next introduce a set of sufficient conditions for the virtual density to be single-dipped. If receivers have a common prior which coincides with the sender's prior, the single-dippedness of the virtual density is equivalent to the single-dippedness of the density of costs:

Proposition 6. *Suppose $p_r = p_s$ for all r. The virtual density $h(\mu)$ is single-dipped in μ if and only if the density of costs $f(c)$ is single-dipped in c.*

If receivers have a common cost, the single-dippedness of the virtual density is implied by a condition that is slightly stronger than the log-convexity of the density of priors:

Proposition 7. *Suppose $c_r = c$ for all r. The virtual density $h(\mu)$ is single-dipped if*

$$\frac{\partial^2}{\partial p^2} \log f(p) > 2(\gamma - 1)^2 \max\left\{1, \frac{1}{\gamma^2}\right\} \quad \text{for all } p \in [0,1], \tag{13}$$

where $\gamma \equiv \frac{1-c}{c}\frac{1-p_s}{p_s} \geq 0$.

Note that the right-hand side of equation (13) is positive, and the left-hand side is positive if $f(p)$ is log-convex. Therefore, condition (13) can be interpreted as "$f(p)$ being sufficiently log-convex." If $p_s + c = 1$, then $\gamma = 1$, and condition (13) reduces to the log-convexity of the distribution of priors.

[11] When the sender's prior is low enough that $p_s < \hat{\mu}$ in Figure 8, the optimal policy reveals no information, and any policy that satisfies $\sigma^0 = \sigma^1$ is optimal. In that case, we choose the policy according to which $\sigma^0 = \sigma^1 = 0$ and set the receivers' posterior following the (zero probability) good message to $\mu = 1$.

When the virtual density is single-dipped, the media's optimal strategy entails sending the bad message when the state is bad. Then, media bias is summarized by the probability σ^1 of sending the good message when the state is good:

Definition 7. Consider single-dipped virtual densities h_1 and h_2 with the corresponding optimal strategies $\sigma_1 = (\sigma_1^0, \sigma_1^1)$ and $\sigma_2 = (\sigma_2^0, \sigma_2^1)$ for the sender. The media are *more biased* given h_1 than given h_2 if $\sigma_1^1 \leq \sigma_2^1$.

In the remainder of this section, we characterize how changes in the primitives of the model affect the extent of media bias in the single-dipped case.

5.1 Changes in Popularity

We begin by studying how shifts in the distribution of costs and priors affect the sender's optimal strategy and the resulting media bias. Our results are expressed in terms of the partial orders introduced in Section 4.1. Not surprisingly, as in the single-peaked case, an increase in the support for the policy increases media bias:

Theorem 4. *Let h_1 and h_2 be two single-dipped virtual densities. If h_1 is smaller than h_2 in the hazard rate order, the media are more biased given h_1 than h_2.*

When the virtual density is single-dipped, the right notion of popularity is the hazard rate order. Intuitively, because the ex-ante opponents are all compliers, the measure of inframarginal compliers is $1 - H(\mu)$. Consequently, $h(\mu)/(1 - H(\mu))$ informs the intensive-versus-extensive-margin trade-off.

With a common prior, the hazard rate order on virtual densities reduces to the hazard rate order on the distributions of costs:

Corollary 4. *Suppose $p_r = p_s$ for all r, and consider two single-dipped densities $f_1(c)$ and $f_2(c)$ for the receivers' cost of support. If f_1 is smaller than f_2 in the hazard rate order, the media are more biased given f_1 than f_2.*

If the receivers have common costs, a larger distribution of priors in the reversed hazard rate order leads to a smaller virtual density in the hazard rate order, and so, we have the following corollary to Theorem 4:

Corollary 5. *Suppose $c_r = c$ for all r, and consider two distributions of priors $f_1(p)$ and $f_2(p)$ that are both strictly positive for all $p \in (0, 1)$ and satisfy condition (13). If f_1 is larger than f_2 in the reversed hazard rate order, the media are more biased given f_1 than f_2.*

5.2 Polarization

We next examine the impact of increased polarization on media bias. The following partial order is the appropriate adaptation of the partial order defined in Section 4.2 for single-peaked densities to the set of single-dipped virtual densities:

Definition 8. Consider single-dipped densities f_1 and f_2 supported on a common compact set and satisfying

$$f_2(x) = \frac{(f_1(x))^\alpha}{\kappa} \qquad \text{for all } x, \tag{14}$$

some $\alpha > 0$, and a normalization constant $\kappa > 0$. If $\alpha \geq 1$, then f_2 is *more polarized* than f_1. If $0 < \alpha \leq 1$, then f_2 is *less polarized* than f_1.

Figure 10 illustrates the polarization partial order on a set of a single-dipped Beta distributions. As the distribution becomes more polarized, mass is moved from the center of the distribution to its extremes.

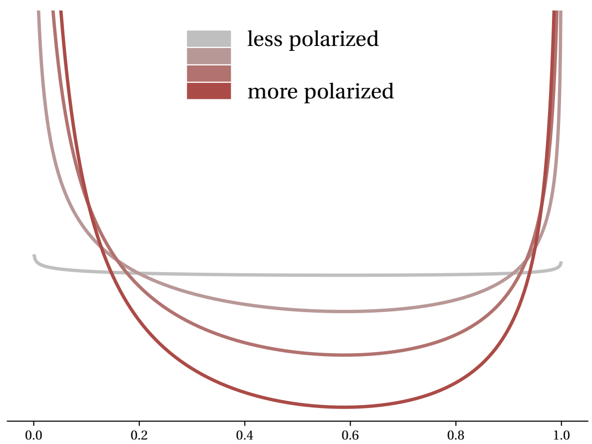

Figure 10. The polarization order on single-dipped densities.

It is, once again, instructive to consider the limits of equation (14). In the $\alpha \to \infty$ limit, f_2 becomes two point masses at the bounds of the support. Our measure identifies such distributions as extremely polarized. Conversely, in the $\alpha \to 0$ limit, f_2 becomes the uniform distribution. Thus, the uniform distribution is less polarized than any single-dipped distribution with the same support.

Our next result establishes that, as in the single-peaked case, polarization decreases media bias in the single-dipped case:

Theorem 5. *Let h_1 and h_2 be two single-dipped virtual densities. If h_1 is more polarized than h_2, then media are less biased given h_1 than h_2.*

Theorem 5 shows that the main message of Theorem 3 continues to apply for single-dipped virtual densities: Polarization tends to reduce the partisan media's bias. Intuitively, an increase in polarization turns some ex-ante moderates into ex-ante extremists. Under the optimal policy, then, there are fewer compliers and more always-supporters. As a result, partisan media find it beneficial to reduce their bias and increase their outreach.

5.3 Polarization and Media Bias

Theorems 3 and 5 paint a consistent picture across the board: Polarization reduces media bias. This observation can be succinctly summarized in a single figure by extending the polarization partial order.

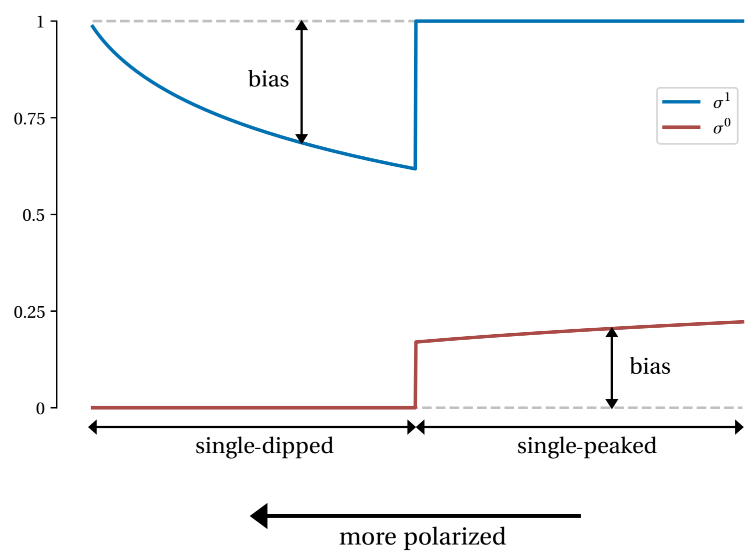

Figure 11. Partisan media's bias as a function of polarization in society.

Take a single-peaked virtual density h_1 supported on $[0,1]$, and consider the parametric family of distributions $\{h_\alpha\}_\alpha$ parameterized by the scalar $\alpha \in \mathbb{R}$:

$$h_\alpha(x) = \frac{(h_1(x))^\alpha}{\kappa(\alpha)} \qquad \text{for all } x,$$

where $\kappa(\alpha)$ is a normalization constant. For positive values of α, h_α is a single-peaked distribution. It becomes more polarized as α decreases to zero. As $\alpha \to 0$, h_α converges to the uniform distribution, which is more polarized than any single-peaked distribution. For negative values of α, h_α is a single-dipped density, which is more polarized than the uniform distribution. It becomes more polarized as α becomes more negative. The upshot is that a lower value of α—be it positive or negative—corresponds to a more polarized society.

Figure 11 illustrates the effect of polarization on media bias. The virtual density is a (symmetric) Beta$(1+\alpha, 1+\alpha)$ distribution, and the sender's prior is given by $p_s = 0.4$. The figure plots how the media's optimal strategy changes as α ranges from -1 to $+1$. In the right half of the figure, $\alpha > 0$, the distribution is single-peaked, and so, by Proposition 4, the optimal policy has the form $(\sigma^0_\alpha,$

$\sigma_\alpha^1) = (\sigma_\alpha^0, 1)$. As polarization increases, by Theorem 3, σ_α^0 decreases and the media become less biased. On the left half of the figure, $\alpha < 0$, the distribution is single-dipped, the optimal policy has the form $(\sigma_\alpha^0, \sigma_\alpha^1) = (0, \sigma_\alpha^1)$ (by Proposition 5), and σ_α^1 increases and media bias decreases with polarization (by Theorem 5).

Transitioning from a single-peaked to a single-dipped virtual density changes the nature of partisan media's optimal policy. This makes it hard to compare the extent of bias in the single-peaked and single-dipped cases. We continue using our simple measure of bias while acknowledging that any such measure will be imperfect. The change in the nature of optimal policy as we transition from a single-peaked to a single-dipped density manifests itself in a possibly discontinuous change in our measure of media bias, as can be seen in Figure 11.

6 Conclusion

Mass polarization has increased in recent decades (Abramowitz and Saunders, 2008; Gentzkow, 2016; Alesina, Miano and Stantcheva, 2020). This has led to concerns about the negative consequences of polarization. Perhaps chief among them is the concern that increased exposure to partisan media may further polarize their audience, creating a vicious cycle (Iyengar and Hahn, 2009; Stroud, 2010; Levendusky, 2013). This problem is especially grave when individuals have preferences for like-minded news (Gentzkow and Shapiro, 2010; Gentzkow, Shapiro and Stone, 2015; Chopra, Haaland and Roth, 2021; Fowler and Kim, 2022; Herrera and Sethi, 2022).

This paper identifies a force in the opposite direction: When media are strategic, and citizens are sophisticated, polarization reduces media bias. Polarization increases the number of citizens with extreme opinions and attitudes relative to the number of moderates. Partisan media thus have more to gain by reaching out to citizens on the fringes of society. The very fact that such citizens are harder to convince compels the media to become more informative and reduce their bias. This theoretical channel is a potential explanation for Prior (2013)'s finding that "There is no firm evidence that partisan media are making ordinary Americans more partisan."

In this paper, we considered a single media source. While our theoretical results maintain their relevance as a study of a politician's "ideal media landscape," the question of how they change with competition among several media sources (Chen and Suen, 2021; Perego and Yuksel, 2022; Sun, Schram and Sloof, 2022) is a fruitful avenue for future research.

Appendices

A Proofs for Section 2

Because f is continuously differentiable over its support and bounded, $h'(\mu)$ exists and is continuous. We begin by noting that the virtual density is single-peaked if and only if $h'(\mu)$ satisfies the *strict single-crossing-from-above property*. The strict single-crossing property is adapted from (Milgrom and Shannon, 1994, p.160) and is as follows:

$$\text{If } h'(\mu) \geq 0 \text{ for some } \mu \in [0,1]\text{, then } h'(\tilde{\mu}) > 0 \text{ for all } \tilde{\mu} < \mu.$$

In our proofs, we rely on the equivalence of this condition with single-peakedness of h.

Proof of Theorem 1. If $h'(\mu)$ satisfies the strict single-crossing-from-above condition, by definition, so does $v''(\mu)$. Therefore, whenever $v(\mu)$ is convex at μ, it is strictly convex at any $\hat{\mu} < \mu$. This means that $v(\mu)$ is first strictly convex and then strictly concave. Therefore, the set where the concave closure of $v(\mu)$—call it $V(\mu)$—coincides with $v(\mu)$ has the following form:

$$\{\mu \in [0,1] : V(\mu) = v(\mu)\} = \{0\} \cup [\hat{\mu}, 1],$$

for some $\hat{\mu} \in [0,1]$.

When $p_s < \hat{\mu}$, by Corollary 2 of Kamenica and Gentzkow (2011), the optimal policy generates two posteriors: $\mu \in \{0, \hat{\mu}\}$. This is achieved by two messages.

When $p_s \geq \hat{\mu}$, the optimal policy is not revealing any information. This can also be achieved by two messages, $m \in \{0,1\}$, and an information structure where $\Pr(m = 1|\theta = 0) = \Pr(m = 1|\theta = 1)$. □

Proof of Proposition 1. Since $p_r = p_s$ for all r, equation (5) simplifies to

$$v(\mu) = \int_0^{c(\mu, p_s)} f(c)dc.$$

On the other hand, by definition, $c(\mu, p_s) = \mu$ for all μ. Therefore, $h(\mu) = v'(\mu) = f(c(\mu, p_s)) = f(\mu)$, and so, h is single-peaked if and only if f is single-peaked. □

Proof of Proposition 2. When $c_r = c$ for all r, equation (5) simplifies to

$$v(\mu) = \int_{p(\mu,c)}^{1} f(p)dp, \tag{15}$$

where

$$p(\mu, c) \equiv \frac{1-\mu}{(1-\mu) + \mu\frac{1-c}{c}\frac{1-p_s}{p_s}}. \tag{16}$$

The virtual density $h(\mu)$ is then given by

$$h(\mu) = v'(\mu) = -f(p(\mu, c)) \cdot \frac{\partial}{\partial \mu} p(\mu, c),$$

and so,

$$h'(\mu) = -f'(p(\mu, c)) \left(\frac{\partial}{\partial \mu} p(\mu, c) \right)^2 - f(p(\mu, c)) \cdot \frac{\partial^2}{\partial \mu^2} p(\mu, c).$$

Therefore, the sign of $h'(\mu)$ is the same as the sign of

$$-\frac{f'(p(\mu, c))}{f(p(\mu, c))} - \frac{\frac{\partial^2}{\partial \mu^2} p(\mu, c)}{\left(\frac{\partial}{\partial \mu} p(\mu, c) \right)^2},$$

where $f(p(\mu, c)) > 0$ by assumption and $\partial p(\mu, c)/\partial \mu > 0$ follows $\gamma > 0$. Using (16) and substituting $\gamma = \frac{1-c}{c}\frac{1-p_s}{p_s}$, we get

$$-\frac{f'(p(\mu, c))}{f(p(\mu, c))} - \frac{\frac{\partial^2}{\partial \mu^2} p(\mu, c)}{\left(\frac{\partial}{\partial \mu} p(\mu, c) \right)^2} = -\frac{\partial}{\partial p} \log f(p(\mu, c)) - 2\frac{\gamma - 1}{\gamma}(1 + (\gamma - 1)\mu). \tag{17}$$

Substituting the value of γ into equation (16) gives: $p(\mu, c) = \frac{1-\mu}{1+(\gamma-1)\mu}$. Solving for μ,

$$\mu = \frac{1 - p(\mu, c)}{1 + (\gamma - 1)p(\mu, c)}. \tag{18}$$

Substituting for μ in equation (17), we get

$$\begin{aligned} -\frac{f'(p(\mu, c))}{f(p(\mu, c))} - \frac{\frac{\partial^2}{\partial \mu^2} p(\mu, c)}{\left(\frac{\partial}{\partial \mu} p(\mu, c) \right)^2} &= -\frac{\partial}{\partial p} \log f(p(\mu, c)) - 2\frac{\gamma - 1}{1 + (\gamma - 1)p(\mu, c)} \\ &= -\frac{\partial}{\partial p} \log f(p(\mu, c)) + g\left(p(\mu, c)\right), \end{aligned}$$

where $g(p) \equiv -2\frac{\gamma-1}{1+(\gamma-1)p}$. Note that $g(p)$ is increasing in p, is convex in p if $\gamma \le 1$, and is concave in p if $\gamma \ge 1$. Therefore,

$$\begin{aligned} \min_{p \in [0,1]} g'(p) &= \begin{cases} g'(0) & \text{if } \gamma \le 1, \\ g'(1) & \text{if } \gamma \ge 1 \end{cases} \\ &= \begin{cases} 2(\gamma - 1)^2 & \text{if } \gamma \le 1, \\ 2\frac{(\gamma-1)^2}{\gamma^2} & \text{if } \gamma \ge 1 \end{cases} \\ &= 2(\gamma - 1)^2 \min\left\{1, \frac{1}{\gamma^2}\right\}. \end{aligned} \tag{19}$$

If condition (6) holds, then

$$\frac{\partial^2}{\partial p^2} \log f(p) < \min_{p \in [0,1]} g'(p),$$

which implies

$$\frac{\partial^2}{\partial p^2}\log f(p) < g'(p) \quad \forall p \in [0,1]. \tag{20}$$

Our claim is that, under condition (6), $h'(\mu)$ satisfies the strict single-crossing-from-above property. To see this, take any two $\mu, \tilde{\mu}$ with $\tilde{\mu} < \mu$ and $h'(\mu) \geq 0$. Because $p(\mu, c)$ is strictly decreasing in μ, $p(\mu, c) < p(\tilde{\mu}, c)$. Since $h'(\mu) \geq 0$, $\frac{\partial}{\partial p}\log f(p(\mu,c)) - g\left(p(\mu,c)\right) \leq 0$. Then,

$$\begin{aligned}
&\frac{\partial}{\partial p}\log f(p(\tilde{\mu},c)) - g\left(p(\tilde{\mu},c)\right) \\
&\quad = \frac{\partial}{\partial p}\log f(p(\mu,c)) - g\left(p(\mu,c)\right) + \int_{p(\mu,c)}^{p(\tilde{\mu},c)} \underbrace{\left(\frac{\partial^2}{\partial p^2}\log f(p) - g'(p)\right)}_{<0 \text{ by } (20)} dp \\
&\quad < \frac{\partial}{\partial p}\log f(p(\mu,c)) - g\left(p(\mu,c)\right) \leq 0.
\end{aligned}$$

Therefore, $h'(\tilde{\mu}) > 0$. The result follows. □

B Proofs for Section 3

Proof of Proposition 3. Consider a receiver r who receives message m, given sender's strategy (σ^0, σ^1). By Bayes' Rule, his posteriors are given by:

$$\Pr_r(\theta = 1 | m = 1) = \frac{p_r \sigma^1}{p_r \sigma^1 + (1-p_r)\sigma^0}, \tag{21}$$

$$\Pr_r(\theta = 1 | m = 0) = \frac{p_r(1-\sigma^1)}{p_r(1-\sigma^1) + (1-p_r)(1-\sigma^0)}. \tag{22}$$

Note that both values are increasing in p_r. Moreover, since $\sigma^1 \geq \sigma^0$,

$$\Pr_r(\theta = 1|m = 1) \geq p_r \geq \Pr_r(\theta = 1|m = 0).$$

That is, good news updates the prior upwards, and bad news updates it downwards. Given the utility function of the receiver, his action is:[12]

$$a_r(m) = \begin{cases} 1 & \text{if} \quad \Pr_r(\theta = 1|m) \geq c_r, \\ 0 & \text{otherwise.} \end{cases} \tag{23}$$

Given the observations so far, sender's strategy (σ^0, σ^1) partitions the receivers into the following three groups:

1. A receiver r is a never-supporter if and only if

$$c_r > \Pr_r(\theta = 1|m = 1).$$

[12]The receiver is indifferent between the two actions when his posterior equals c_r. Consequently, there is an indeterminacy in this case. Since f is continuously differentiable and bounded over its support, such receivers have measure zero and are therefore inconsequential for the analysis.

Substituting (21), this is equivalent to

$$c_r > \frac{p_r \sigma^1}{p_r \sigma^1 + (1-p_r)\sigma^0}.$$

2. A receiver r is a complier if and only if

$$\Pr_r(\theta = 1 | m = 1) \geq c_r > \Pr_r(\theta = 1 | m = 0).$$

Substituting (21) and (22), this is equivalent to

$$\frac{p_r \sigma^1}{p_r \sigma^1 + (1-p_r)\sigma^0} \geq c_r > \frac{p_r(1-\sigma^1)}{p_r(1-\sigma^1) + (1-p_r)(1-\sigma^0)}.$$

3. A receiver r is an always-supporter if and only if

$$c_r \leq \Pr_r(\theta = 1 | m = 0).$$

Substituting (22), this is equivalent to:

$$c_r \leq \frac{p_r(1-\sigma^1)}{p_r(1-\sigma^1) + (1-p_r)(1-\sigma^0)}.$$

Rearranging the inequalities yields the result. □

Proof of Proposition 4. As discussed in the proof of Theorem 1, when h is single-peaked,

$$\{\mu \in [0,1] : V(\mu) = v(\mu)\} = \{0\} \cup [\hat{\mu}, 1].$$

If $p_s < \hat{\mu}$, the optimal policy generates two posteriors: $\mu \in \{0, \hat{\mu}\}$. This is achieved by two messages, with one perfectly revealing the bad state. When $p_s \geq \hat{\mu}$, the optimal policy is not revealing any information. This can be achieved by two messages, $m \in \{0,1\}$, and an information structure where $\Pr(m = 1 | \theta = 0) = \Pr(m = 1 | \theta = 1) = 1$. Message $m = 0$ will occur with probability zero, and the posterior beliefs following $m = 0$ will be free in a Perfect Bayesian Equilibrium. One can assign the posterior $\Pr_r(\theta = 0 | m = 0) = 1$ assign $m = 0$ as the message that perfectly reveals the bad state. This proves the first part of Proposition 4.

Because the optimal policy involves the bad news perfectly revealing the bad state, $\sigma^1 = 1$ under the optimal policy. Substituting into (7), $\overline{p}(c) = 1$ for all c, i.e., there are no always-supporters. Moreover,

$$\underline{p}(c) = \frac{c\sigma^0}{c\sigma^0 + (1-c)} \leq c.$$

Therefore, for any r with $p_r \geq c_r$, we have: $p_r \geq \underline{p}(c_r)$. By Proposition 3, then, every ex-ante supporter is a complier. □

C Proofs for Section 4

We begin by introducing some notation and preliminary results for the remaining proofs.

Lemma 1. *The value function $v(\mu)$ satisfies*

$$\lim_{\mu\to 0}\mu v'(\mu) = \lim_{\mu\to 1}(1-\mu)v'(\mu) = 0.$$

Proof. First, note that

$$\begin{aligned} v'(\mu) &= \int_0^1 f(p, c(\mu,p)).\frac{\partial}{\partial\mu}c(\mu,p)dp \\ &= \int_0^1 f(p,c(\mu,p))\frac{p(1-p)p_s(1-p_s)}{(p_s(1-p)+\mu(p-p_s))^2}dp, \end{aligned}$$

But since f is bounded, there exist some $C > 0$ such that

$$\begin{aligned} |v'(\mu)| &\le C\int_0^1 \frac{p(1-p)p_s(1-p_s)}{(p_s(1-p)+\mu(p-p_s))^2}dp \\ &= C\frac{p_s(1-p_s)}{(p_s-\mu)^3}\left[2(\mu-p_s)-(\mu(1-p_s)+p_s(1-\mu))\log\left(\frac{\mu(1-p_s)}{p_s(1-\mu)}\right)\right]. \end{aligned}$$

On the other hand,

$$\begin{aligned} &\lim_{\mu\to 0}\mu\cdot\frac{p_s(1-p_s)}{(p_s-\mu)^3}\left[2(\mu-p_s)-(\mu(1-p_s)+p_s(1-\mu))\log\left(\frac{\mu(1-p_s)}{p_s(1-\mu)}\right)\right] \\ &\quad = \lim_{\mu\to 0}\frac{-(1-p_s)}{p_s}\mu\log(\mu) = 0, \end{aligned}$$

and

$$\begin{aligned} &\lim_{\mu\to 1}(1-\mu)\cdot\frac{p_s(1-p_s)}{(p_s-\mu)^3}\left[2(\mu-p_s)-(\mu(1-p_s)+p_s(1-\mu))\log\left(\frac{\mu(1-p_s)}{p_s(1-\mu)}\right)\right] \\ &\quad = \lim_{\mu\to 1}\frac{-p_s}{1-p_s}(1-\mu)\log(1-\mu) = 0. \end{aligned}$$

Therefore,

$$\lim_{\mu\to 0}\mu v'(\mu) = \lim_{\mu\to 1}(1-\mu)v'(\mu) = 0.$$

This completes the proof of the Lemma. □

Consider a single-peaked virtual density $h(\mu)$. As discussed in the proof of Theorem 1, $\{\mu \in [0,1] : V(\mu) = v(\mu)\} = \{0\} \cup [\hat\mu, 1]$ for some $\hat\mu \in [0,1]$. Note that:

- $v'(\mu)\mu < v(\mu)$ for all $\mu \in (0,1)$ if and only if $\hat\mu = 0$.
- $v'(\mu)\mu > v(\mu)$ for all $\mu \in (0,1)$ if and only if $\hat\mu = 1$.
- When $\hat\mu \in (0,1)$, it satisfies:

$$v'(\hat\mu)\hat\mu = v(\hat\mu). \tag{24}$$

Let

$$y(\mu) \equiv v'(\mu)\mu - v(\mu) = h(\mu)\mu - \int_0^\mu h(\tilde{\mu})\tilde{\mu}, \qquad \forall \mu \in [0,1]. \tag{25}$$

Then, $\hat{\mu} \in (0,1)$ is characterized by the equation: $y(\hat{\mu}) = 0$. We start with three remarks.

Remark 1. $\lim_{\mu \to 0} y(\mu) = 0$. This follows Lemma 1 and the fact that $v(0) = 0$.

Remark 2. *$y(\mu)$ is continuous in μ over* $(0,1)$. This is because f is continuous over its support.

Remark 3. *$y(\mu)$ is first strictly increasing and then strictly decreasing.* This is because $y'(\mu) = v''(\mu)\mu + v'(\mu) - v'(\mu) = v''(\mu)\mu = h'(\mu)\mu$. Since $h'(\mu)$ satisfies strict single crossing from above, so does $y'(\mu)$, and the remark follows.

Define the set

$$\mathcal{U}_y \equiv \{\mu \in [0,1]: \quad y(\mu) \geq 0\}.$$

Based on Remarks 1, 2 and 3, we conclude that $\mathcal{U}_y$ has the following form:

$$\mathcal{U}_y = [0, \hat{\mu}].$$

Our approach through the rest of this section is built on showing that $y(\mu)$ changes in a predictable manner (and so does $\hat{\mu}$).

Proof of Theorem 2. Let:

$$y_1(\mu) = h_1(\mu)\mu - \int_0^\mu h_1(\tilde{\mu})\tilde{\mu},$$
$$y_2(\mu) = h_2(\mu)\mu - \int_0^\mu h_2(\tilde{\mu})\tilde{\mu}.$$

Since both virtual densities are single-peaked:

$$\mathcal{U}_{y_1} = \{\mu \in [0,1]: \quad y_1(\mu) \geq 0\} = [0, \hat{\mu}_1],$$
$$\mathcal{U}_{y_2} = \{\mu \in [0,1]: \quad y_2(\mu) \geq 0\} = [0, \hat{\mu}_2].$$

Take any $\mu \in (0,1]$, and suppose $y_1(\mu) \geq 0$. Then,

$$h_1(\mu)\mu - \int_0^\mu h_1(\tilde{\mu})\tilde{\mu} \geq 0,$$

which implies

$$\frac{h_1(\mu)\mu}{\int_0^\mu h_1(\tilde{\mu})} \geq 1.$$

On the other hand, since h_2 is larger than h_1 in the reversed hazard rate order,

$$\frac{h_2(\mu)\mu}{\int_0^\mu h_2(\tilde{\mu})} \geq \frac{h_1(\mu)\mu}{\int_0^\mu h_1(\tilde{\mu})} \geq 1,$$

which implies

$$h_2(\mu)\mu - \int_0^\mu h_2(\tilde{\mu})\tilde{\mu} \geq 0.$$

Therefore, $y_2(\mu) \geq 0$. We conclude that

$$\{\mu \in [0,1] : y_1(\mu) \geq 0\} \subseteq \{\mu \in [0,1] : y_2(\mu) \geq 0\},$$

and therefore, $[0, \hat{\mu}_1] \subseteq [0, \hat{\mu}_2]$ and $\hat{\mu}_1 \leq \hat{\mu}_2$. To conclude the proof, consider three cases:

1. If $p_s \geq \hat{\mu}_2$, the optimal policy does not reveal any information in either case, and we pick $\sigma_1^0 = \sigma_2^0 = 1$.

2. If $\hat{\mu}_2 > p_s \geq \hat{\mu}_1$, the optimal policy under $h_1(\mu)$ does not reveal any information. In this case, we pick $\sigma_1^0 = 1$ and $\sigma_2^0 < 1$.

3. If $p_s > \hat{\mu}_2$, the optimal policies σ_1^0 and σ_2^0 satisfy:

$$\frac{p_s}{p_s + (1-p_s)\sigma_1^0} = \hat{\mu}_1, \qquad \frac{p_s}{p_s + (1-p_s)\sigma_2^0} = \hat{\mu}_2.$$

   Then, $\hat{\mu}_2 \geq \hat{\mu}_1$ implies $\sigma_2^0 \leq \sigma_1^0$.

In any case, $\sigma_2^0 \leq \sigma_1^0$, and the result follows. □

Proof of Corollary 2. Suppose $p_r = p_s$ for all r, and let f denote the density of costs. As discussed in the proof of Proposition 1, $h'(\mu) = f'(\mu)$ in this case. Therefore, when f_1 is smaller than f_2 in the reversed hazard rate order, h_1 is smaller than h_2 in the reversed hazard rate order as well. The result follows from Theorem 2. □

Proof of Corollary 3. Suppose $c_r = c$ for all r, and let f denote the density of priors. As discussed in the proof of Proposition 2,

$$h(\mu) = -f(p(\mu,c)) \cdot \frac{\partial}{\partial \mu} p(\mu, c),$$

where

$$p(\mu,c) = \frac{1-\mu}{1+(\gamma-1)\mu}, \qquad \gamma = \frac{1-c}{c}\frac{1-p_s}{p_s}. \tag{26}$$

Therefore,

$$\frac{\partial}{\partial \mu} p(\mu,c) = -\frac{\gamma}{(1+(\gamma-1)\mu)^2}. \tag{27}$$

Solving (26) for μ, we get

$$\mu = \frac{1-p(\mu,c)}{1+(\gamma-1)p(\mu,c)}. \tag{28}$$

Substituting (28) into (27) gives

$$\frac{\partial}{\partial \mu} p(\mu, c) = -\frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma}. \tag{29}$$

Therefore,

$$h(\mu) = f(p(\mu,c)) \cdot \frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma}. \tag{30}$$

Finally, note that

$$\int_0^{\mu} h(\tilde{\mu}) d\tilde{\mu} = v(\mu) = \int_{p(\mu,c)}^{1} f(\tilde{p}) d\tilde{p}, \tag{31}$$

where the second equality is (15). Now, consider two densities of costs, f_1 and f_2. By equations (30) and (31), for any $\mu > 0$, the virtual densities satisfy,

$$\frac{h_1(\mu)}{\int_0^{\mu} h_1(\tilde{\mu}) d\tilde{\mu}} = \frac{f_1(p(\mu,c))}{\int_{p(\mu,c)}^{1} f_1(\tilde{p}) d\tilde{p}} \cdot \frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma},$$

and

$$\frac{h_2(\mu)}{\int_0^{\mu} h_2(\tilde{\mu}) d\tilde{\mu}} = \frac{f_2(p(\mu,c))}{\int_{p(\mu,c)}^{1} f_2(\tilde{p}) d\tilde{p}} \cdot \frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma}.$$

When f_1 is larger than f_2 in the hazard rate order,

$$\frac{f_1(p(\mu,c))}{\int_{p(\mu,c)}^{1} f_1(\tilde{p}) d\tilde{p}} \leq \frac{f_2(p(\mu,c))}{\int_{p(\mu,c)}^{1} f_2(\tilde{p}) d\tilde{p}}$$

for all μ. Then,

$$\frac{h_1(\mu)}{\int_0^{\mu} h_1(\tilde{\mu}) d\tilde{\mu}} \leq \frac{h_2(\mu)}{\int_0^{\mu} h_2(\tilde{\mu}) d\tilde{\mu}}$$

for all μ. Thus, h_1 is smaller than h_2 in the reversed hazard rate order. The result follows from Theorem 2. □

Proof of Theorem 3. We begin with two remarks.

Remark 4. *If $h_1(\mu)$ is a single-peaked distribution, then any distribution with density*

$$h_2(\mu) = \frac{(h_1(\mu))^{\alpha}}{\kappa} \qquad \text{for all } \mu \in [0,1], \text{ where } \alpha \geq 1, \kappa > 0$$

is single-peaked. To see this, suppose $h_1(\mu)$ is single-peaked. Then, $h_1'(\mu)$ satisfies the strict single-crossing-from-above condition:

If $h_1'(\mu) \geq 0$ for some $\mu \in [0,1]$, then $h_1'(\tilde{\mu}) > 0$ for all $\tilde{\mu} < \mu$.

Note that

$$h_2'(\mu) = \alpha \frac{(h_1(\mu))^{\alpha-1}}{\kappa} h_1'(\mu),$$

which implies that the sign of $h_2'(\mu)$ is the same as the sign of $h_1'(\mu)$. The remark follows.

Remark 5. *If $h(\mu)$ is a single-peaked distribution, then $h(\mu) > 0$ for all $\mu \in (0,1)$.* This is a simple consequence of the fact that, for any single-peaked distribution, there exists some $\tilde{\mu}$ such that $h'(\mu) > 0$ for all $\mu \in [0, \tilde{\mu})$ and $h'(\mu) < 0$ for all $\mu \in (\tilde{\mu}, 1]$.

Now, take a single-peaked distribution $h(\mu)$. Consider a family of distributions $\{h_\alpha\}_{\alpha \geq 1}$ characterized by

$$h_\alpha(\mu) = \frac{(h(\mu))^\alpha}{\kappa(\alpha)}, \qquad \text{for all } \mu \in [0,1], \alpha \geq 1,$$

where $\kappa(\alpha)$ is the normalization constant given by

$$\kappa(\alpha) \equiv \int_0^1 (h(t))^\alpha dt.$$

The corresponding cdf's are given by:

$$H_\alpha(\mu) \equiv \int_0^\mu h_\alpha(x) dx = \frac{\int_0^\mu (h(x))^\alpha dt}{\kappa(\alpha)}.$$

By Remark 4, any distribution in this family is single-peaked. Take any such distribution h_α, and let

$$y_\alpha(\mu) \equiv h_\alpha(\mu)\mu - H_\alpha(\mu).$$

Then, by the argument in the proof of Theorem 2, the set $\mathcal{U}_{y_\alpha} \equiv \{\mu \in [0,1] : y_\alpha(\mu) \geq 0\}$ has the following form:

$$\mathcal{U}_{y_\alpha} = [0, \hat{\mu}_\alpha].$$

The proof goes through by showing that $\hat{\mu}_\alpha$ is decreasing in α. We start with two important remarks.

Remark 6. $y_\alpha'(\hat{\mu}_\alpha) < 0$. This follows from the fact that $\mathcal{U}_{y_\alpha} = [0, \hat{\mu}_\alpha]$. Then, $y_\alpha(\mu)$ crosses zero from above at $\hat{\mu}_\alpha$. Since $y_\alpha(\mu)$ is differentiable, the remark follows.

Remark 7. *If $\hat{\mu}_\alpha \in (0,1)$, then $y_\alpha(\hat{\mu}_\alpha) = 0$, or equivalently,*

$$h_\alpha(\hat{\mu}_\alpha)\hat{\mu}_\alpha = \int_0^{\hat{\mu}_\alpha} h_\alpha(x) dx. \tag{32}$$

By Remark 7, $\hat{\mu}_\alpha \in (0,1)$ satisfies

$$y_\alpha(\hat{\mu}_\alpha) = 0.$$

Implicitly differentiate with respect to α to get

$$\frac{\partial}{\partial\alpha} y_\alpha(\hat{\mu}_\alpha) + y'_\alpha(\hat{\mu}_\alpha)\frac{\partial\hat{\mu}_\alpha}{\partial\alpha} = 0.$$

By Remark 6, $y'_\alpha(\hat{\mu}_\alpha) < 0$. Then, $\frac{\partial\hat{\mu}_\alpha}{\partial\alpha} \le 0$ if and only if

$$\left.\frac{\partial}{\partial\alpha} y_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \le 0.$$

Note that, for any $\mu \in [0,1]$,

$$\frac{\partial}{\partial\alpha} y_\alpha(\mu) \le 0 \iff \frac{\partial}{\partial\alpha} h_\alpha(\mu)\mu \le \int_0^\mu \frac{\partial}{\partial\alpha} h_\alpha(x)dx.$$

Recall that $h_\alpha(\mu) = \frac{(h(\mu))^\alpha}{\kappa(\alpha)}$. Therefore, for any $x \in (0,1)$, $\frac{\partial}{\partial\alpha} h_\alpha(x) = h_\alpha(x)\left(\log h(x) - \frac{\kappa'(\alpha)}{\kappa(\alpha)}\right)$, and so,

$$\left.\frac{\partial}{\partial\alpha} y_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \le 0 \iff h_\alpha(\hat{\mu}_\alpha)\log h(\hat{\mu}_\alpha)\hat{\mu}_\alpha \le \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log h(x)dx.$$

Using (32) to substitute $h_\alpha(\hat{\mu}_\alpha)\hat{\mu}_\alpha = \int_0^{\hat{\mu}_\alpha} h_\alpha(x)dx$ on the left hand side of the above inequality, we have

$$\begin{aligned}
\left.\frac{\partial}{\partial\alpha} y_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \le 0 &\iff \log h(\hat{\mu}_\alpha)\int_0^{\hat{\mu}_\alpha} h_\alpha(x)dx \le \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log h(x)dx \\
&\iff \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log h(\hat{\mu}_\alpha)dx \le \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log h(x)dx \\
&\iff \alpha\int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log h(\hat{\mu}_\alpha)dx \le \alpha\int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log h(x)dx \\
&\iff \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log(h(\hat{\mu}_\alpha))^\alpha dx \le \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log(h(x))^\alpha dx \\
&\iff \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log\left(\frac{(h(\hat{\mu}_\alpha))^\alpha}{(h(x))^\alpha}\right)dx \le 0 \\
&\iff \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log\left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)}\right)dx \le 0.
\end{aligned}$$

For any real number $z > 0$, $\log(z) \le z - 1$, with a strict inequality for any $z \ne 1$. Therefore,

$$\int_0^{\hat{\mu}_\alpha} h_\alpha(x)\log\left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)}\right)dx \le \int_0^{\hat{\mu}_\alpha} h_\alpha(x)\left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)} - 1\right)dx.$$

Therefore, $\left.\frac{\partial}{\partial\alpha} y_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \le 0$ as long as

$$\begin{aligned}
\int_0^{\hat{\mu}_\alpha} h_\alpha(x)\left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)} - 1\right)dx \le 0 &\iff \int_0^{\hat{\mu}_\alpha} \left(h_\alpha(\hat{\mu}_\alpha) - h_\alpha(x)\right)dx \le 0 \\
&\iff \int_0^{\hat{\mu}_\alpha} h_\alpha(\hat{\mu}_\alpha)dx \le \int_0^{\hat{\mu}_\alpha} h_\alpha(x)dx
\end{aligned}$$

$$\iff h_\alpha(\hat{\mu}_\alpha)\hat{\mu}_\alpha \leq \int_0^{\hat{\mu}_\alpha} h_\alpha(x)dx,$$

which is guaranteed by (32). We conclude that for any $\hat{\mu}_\alpha \in (0,1)$, $\frac{\partial \hat{\mu}_\alpha}{\partial \alpha} \leq 0$.

Since h_1 and h_2 are within the family we considered (with h_1 corresponding to $\alpha = 1$ and h_2 corresponding to some $\alpha \geq 1$), $\hat{\mu}_2 \leq \hat{\mu}_1$. Repeating the same argument in the proof of Theorem 2, we conclude that $\sigma_2^0 \geq \sigma_1^0$. □

D Proofs for Section 5

Note that single-dippedness of the virtual density is equivalent to the following *strict single-crossing-from-below property* for $h'(\mu)$:

If $h'(\mu) \geq 0$ for some $\mu \in [0,1]$, then $h'(\tilde{\mu}) > 0$ for all $\tilde{\mu} > \mu$.

Proof of Proposition 5. If $h'(\mu)$ satisfies the strict single-crossing-from-below condition, by definition, so does $v''(\mu)$. Therefore, whenever $v(\mu)$ is convex at μ, it is strictly convex at any $\hat{\mu} \geq \mu$. This means that $v(\mu)$ is first strictly concave and then strictly convex. Therefore, the set where the concave closure of $v(\mu)$ coincides with $v(\mu)$ has the following form:

$$\{\mu \in [0,1] : V(\mu) = v(\mu)\} = [0, \hat{\mu}] \cup \{1\}.$$

When $p_s < \hat{\mu}$, the optimal policy is not revealing any information. This can be achieved by two messages, $m \in \{0,1\}$, and an information structure where $\Pr(m = 1|\theta = 0) = \Pr(m = 1|\theta = 1) = 0$. Message $m = 1$ will occur with probability zero, and the posterior beliefs following $m = 1$ will be free in a Perfect Bayesian Equilibrium. One can assign posteriors $\Pr_r(\theta = 1|m = 1) = 1$ to make $m = 1$ as the message that perfectly reveals the good state.

When $p_s \geq \hat{\mu}$, by Corollary 2 of Kamenica and Gentzkow (2011), the optimal policy generates two posteriors: $\mu \in \{\hat{\mu}, 1\}$. This is achieved by two messages, $m \in \{0,1\}$, where message $m = 1$ perfectly reveals the good state.

Because the optimal policy involves the good news perfectly revealing the good state, $\sigma^0 = 0$ in the optimal policy. Substituting into (7), $\underline{p}(c) = 0$ for all c, i.e., there are no never-supporters. Moreover,

$$\overline{p}(c) = \frac{c}{c + (1-c)(1-\sigma^1)} \geq c.$$

Therefore, for any r with $p_r < c_r$, we have: $p_r < \overline{p}(c_r)$. By Proposition 3, then, every ex-ante opponent is a complier. □

Proof of Proposition 6. The proof of Proposition 6 is identical to the proof of Proposition 1. □

Proof of Proposition 7. The proof follows identical steps to that of Proposition 2 until equation (19). The rest of the argument is provided below.

Recall that $g(p) \equiv -2\frac{\gamma-1}{1+(\gamma-1)p}$ and $g(p)$ is increasing in p, convex in p if $\gamma \le 1$, and concave in p if $\gamma \ge 1$. Therefore,

$$\max_{p\in[0,1]} g'(p) = \begin{cases} g'(1) & \text{if } \gamma \le 1 \\ g'(0) & \text{if } \gamma \ge 1 \end{cases} = \begin{cases} 2\frac{(\gamma-1)^2}{\gamma^2} & \text{if } \gamma \le 1 \\ 2(\gamma-1)^2 & \text{if } \gamma \ge 1 \end{cases} = 2(\gamma-1)^2 \max\left\{1, \frac{1}{\gamma^2}\right\}. \tag{33}$$

If condition (13) holds,

$$\frac{\partial^2}{\partial p^2} \log f(p) > \max_{p\in[0,1]} g'(p),$$

and so

$$\frac{\partial^2}{\partial p^2} \log f(p) > g'(p) \quad \forall p \in [0,1]. \tag{34}$$

Our claim is that, under condition (13), $h'(\mu)$ satisfies the strict single-crossing-from-below property. To see this, take any two $\mu, \tilde{\mu}$ with $\tilde{\mu} > \mu$ and $h'(\mu) \ge 0$. Because $p(\mu, c)$ is strictly decreasing in μ, $p(\tilde{\mu}, c) < p(\mu, c)$. Since $h'(\mu) \ge 0$, $\frac{\partial}{\partial p} \log f(p(\mu,c)) - g\left(p(\mu,c)\right) \le 0$. Then,

$$\begin{aligned} &\frac{\partial}{\partial p} \log f(p(\tilde{\mu},c)) - g\left(p(\tilde{\mu},c)\right) \\ &= \frac{\partial}{\partial p} \log f(p(\mu,c)) - g\left(p(\mu,c)\right) - \int_{p(\tilde{\mu},c)}^{p(\mu,c)} \underbrace{\left(\frac{\partial^2}{\partial p^2} \log f(p) - g'\left(p\right)\right)}_{>0 \text{ by } (34)} dp \\ &< \frac{\partial}{\partial p} \log f(p(\mu,c)) - g\left(p(\mu,c)\right) \le 0. \end{aligned}$$

Therefore, $h'(\tilde{\mu}) > 0$. The result follows. □

We continue by introducing some notation and preliminary results for the remaining proofs.

Consider a single-dipped virtual density $h(\mu)$. As discussed in the proof of Proposition 5, $\{\mu \in [0,1] : V(\mu) = v(\mu)\} = [0, \hat{\mu}] \cup \{1\}$ for some $\hat{\mu} \in [0,1]$. Note that:

- $v'(\mu)(1-\mu) > 1 - v(\mu)$ for all $\mu \in (0,1)$ if and only if $\hat{\mu} = 1$.
- $v'(\mu)(1-\mu) < 1 - v(\mu)$ for all $\mu \in (0,1)$ if and only if $\hat{\mu} = 0$.
- When $\hat{\mu} \in (0,1)$, it satisfies:

$$v'(\hat{\mu})(1-\hat{\mu}) = 1 - v(\hat{\mu}). \tag{35}$$

Let

$$z(\mu) \equiv v'(\mu)(1-\mu) - (1 - v(\mu)) = h(\mu)(1-\mu) - \int_{\mu}^{1} h(\tilde{\mu})\tilde{\mu}, \qquad \forall \mu \in [0,1]. \tag{36}$$

Then, $\hat{\mu} \in (0,1)$ is characterized by the equation: $z(\hat{\mu}) = 0$. We start with three remarks.

Remark 8. $\lim_{\mu\to 1} z(\mu) = 0$. This follows Lemma 1 and the fact that $1 - v(1) = 0$.

Remark 9. *$z(\mu)$ is continuous in μ over $(0,1)$.* This is because f is continuous over its support.

Remark 10. *$z(\mu)$ is first strictly decreasing and then increasing.* This is because $z'(\mu) = v''(\mu)(1-\mu) - v'(\mu) + v'(\mu) = v''(\mu)(1-\mu) = h'(\mu)(1-\mu)$. Since $h'(\mu)$ satisfies strict single crossing from below, so does $z'(\mu)$, and the remark follows.

Define the set

$$\mathcal{L}_z \equiv \{\mu \in [0,1] : \quad z(\mu) \leq 0\}.$$

Based on Remarks 8, 9 and 10, we conclude that $\mathcal{L}_z$ has the following form:

$$\mathcal{L}_z = [\hat{\mu}, 1].$$

Our approach through the rest of this section is built on showing that $z(\mu)$ changes in a predictable manner (and so does $\hat{\mu}$).

Proof of Theorem 4. Let:

$$z_1(\mu) = h_1(\mu)(1-\mu) - \int_{\mu}^{1} h_1(\tilde{\mu})\tilde{\mu},$$
$$z_2(\mu) = h_2(\mu)(1-\mu) - \int_{\mu}^{1} h_2(\tilde{\mu})\tilde{\mu}.$$

Since both virtual densities are single-dipped:

$$\mathcal{L}_{z_1} = \{\mu \in [0,1] : \quad z_1(\mu) \leq 0\} = [\hat{\mu}_1, 1],$$
$$\mathcal{L}_{z_2} = \{\mu \in [0,1] : \quad z_2(\mu) \leq 0\} = [\hat{\mu}_2, 1].$$

Take any $\mu \in [0,1)$, and suppose $z_1(\mu) \leq 0$. Then,

$$h_1(\mu)(1-\mu) - \int_{\mu}^{1} h_1(\tilde{\mu})\tilde{\mu} \leq 0,$$

which implies

$$\frac{h_1(\mu)(1-\mu)}{\int_{\mu}^{1} h_1(\tilde{\mu})} \leq 1.$$

On the other hand, since h_2 is larger than h_1 in the hazard rate order,

$$\frac{h_2(\mu)(1-\mu)}{\int_{\mu}^{1} h_2(\tilde{\mu})} \leq \frac{h_1(\mu)(1-\mu)}{\int_{\mu}^{1} h_1(\tilde{\mu})} \leq 1,$$

and so,

$$h_2(\mu)(1-\mu) - \int_{\mu}^{1} h_2(\tilde{\mu})\tilde{\mu} \leq 0.$$

Therefore, $z_2(\mu) \leq 0$. We conclude that

$$\{\mu \in [0,1] : z_1(\mu) \leq 0\} \subseteq \{\mu \in [0,1] : z_2(\mu) \leq 0\},$$

and therefore, $[\hat{\mu}_1, 1] \subseteq [\hat{\mu}_2, 1]$ and $\hat{\mu}_1 \geq \hat{\mu}_2$. To conclude the proof, consider three cases:

1. If $p_s \leq \hat{\mu}_2$, the optimal policy does not reveal any information in either case, and we pick $\sigma_1^1 = \sigma_2^1 = 0$.

2. If $\hat{\mu}_2 < p_s \leq \hat{\mu}_1$, the optimal policy under $h_1(\mu)$ does not reveal any information. In this case, we pick $\sigma_1^1 = 0$ and $\sigma_2^1 > 0$.

3. If $p_s < \hat{\mu}_2$, the optimal policies σ_1^1 and σ_2^1 satisfy:

$$\frac{p_s(1-\sigma_1^1)}{p_s(1-\sigma_1^1)+(1-p_s)} = \hat{\mu}_1, \qquad \frac{p_s(1-\sigma_2^1)}{p_s(1-\sigma_2^1)+(1-p_s)} = \hat{\mu}_2.$$

   Then, $\hat{\mu}_2 \leq \hat{\mu}_1$ implies $\sigma_2^1 \geq \sigma_1^1$.

In any case, $\sigma_2^1 \geq \sigma_1^1$, and the result follows. □

Proof of Corollary 4. Suppose $p_r = p_s$ for all r, and let f denote the density of costs. As discussed in the proof of Proposition 1, $h'(\mu) = f'(\mu)$ in this case. Therefore, when f_1 is smaller than f_2 in the hazard rate order, h_1 is smaller than h_2 in the hazard rate order as well. The result follows from Theorem 4. □

Proof of Corollary 5. Suppose $c_r = c$ for all r, and let f denote the density of priors. As discussed in the proof of Proposition 2,

$$h(\mu) = -f(p(\mu,c)) \cdot \frac{\partial}{\partial \mu} p(\mu,c),$$

where

$$p(\mu,c) = \frac{1-\mu}{1+(\gamma-1)\mu}, \qquad \gamma = \frac{1-c}{c}\frac{1-p_s}{p_s}. \tag{37}$$

Therefore,

$$\frac{\partial}{\partial \mu} p(\mu,c) = -\frac{\gamma}{(1+(\gamma-1)\mu)^2}. \tag{38}$$

Solving (37) for μ, we get

$$\mu = \frac{1-p(\mu,c)}{1+(\gamma-1)p(\mu,c)}. \tag{39}$$

Substituting (39) into (38) gives

$$\frac{\partial}{\partial \mu} p(\mu,c) = -\frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma}. \tag{40}$$

Therefore,

$$h(\mu) = f(p(\mu,c)) \cdot \frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma}. \tag{41}$$

Finally, note that

$$\int_{\mu}^{1} h(\tilde{\mu})d\tilde{\mu} = 1 - v(\mu) = \int_{0}^{p(\mu,c)} f(\tilde{p})d\tilde{p}, \tag{42}$$

where the second equality follows (15). Now, consider two densities of costs, f_1 and f_2. By equations (41) and (42), for any $\mu > 0$, the virtual densities satisfy,

$$\frac{h_1(\mu)}{\int_{\mu}^{1} h_1(\tilde{\mu})d\tilde{\mu}} = \frac{f_1(p(\mu,c))}{\int_0^{p(\mu,c)} f_1(\tilde{p})d\tilde{p}} \cdot \frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma},$$

and

$$\frac{h_2(\mu)}{\int_{\mu}^{1} h_2(\tilde{\mu})d\tilde{\mu}} = \frac{f_2(p(\mu,c))}{\int_0^{p(\mu,c)} f_2(\tilde{p})d\tilde{p}} \cdot \frac{(1+(\gamma-1)p(\mu,c))^2}{\gamma}.$$

When f_1 is larger than f_2 in the hazard rate order,

$$\frac{f_1(p(\mu,c))}{\int_0^{p(\mu,c)} f_1(\tilde{p})d\tilde{p}} \geq \frac{f_2(p(\mu,c))}{\int_0^{p(\mu,c)} f_2(\tilde{p})d\tilde{p}}$$

for all μ. Then,

$$\frac{h_1(\mu)}{\int_{\mu}^{1} h_1(\tilde{\mu})d\tilde{\mu}} \leq \frac{h_2(\mu)}{\int_{\mu}^{1} h_2(\tilde{\mu})d\tilde{\mu}}$$

for all μ. Thus, h_1 is larger than h_2 in the reversed hazard rate order. The result follows from Theorem 4. □

Proof of Theorem 5. We begin with two remarks.

Remark 11. *If $h_1(\mu)$ is a single-dipped distribution, then any distribution with density*

$$h_2(\mu) = \frac{(h_1(\mu))^{\alpha}}{\kappa} \qquad \textit{for all } \mu \in [0,1], \textit{ where } \alpha \geq 1, \kappa > 0$$

is single-dipped. To see this, suppose $h_1(\mu)$ is single-dipped. Then, $h_1'(\mu)$ satisfies the strict single-crossing-from-below property:

If $h_1'(\mu) \geq 0$ for some $\mu \in [0,1]$, then $h_1'(\tilde{\mu}) > 0$ for all $\tilde{\mu} > \mu$.

Note that

$$h_2'(\mu) = \alpha \frac{(h_1(\mu))^{\alpha-1}}{\kappa} h_1'(\mu),$$

which implies that the sign of $h_2'(\mu)$ is the same as the sign of $h_1'(\mu)$. The remark follows.

Remark 12. *If $h(\mu)$ is a single-dipped distribution, then $h(\mu) > 0$ for almost all μ.* This is a simple consequence of the fact that for any single-dipped distribution, there exists some $\tilde{\mu}$ such that $h'(\mu) < 0$ for all $\mu \in [0,\tilde{\mu})$ and $h'(\mu) > 0$ for all $\mu \in (\tilde{\mu},1]$. The only point μ at which $h(\mu)$ could be zero is $\tilde{\mu}$.

Now, take a single-dipped distribution $h(\mu)$. Consider a family of distributions $\{h_\alpha\}_{\alpha \geq 1}$ characterized by

$$h_\alpha(\mu) = \frac{(h(\mu))^\alpha}{\kappa(\alpha)}, \qquad \text{for all } \mu \in [0,1], \alpha \geq 1,$$

where $\kappa(\alpha)$ is the normalization constant given by

$$\kappa(\alpha) \equiv \int_0^1 (h(t))^\alpha dt.$$

The corresponding cdf's are given by:

$$H_\alpha(\mu) \equiv \int_0^\mu h_\alpha(x)dx = \frac{\int_0^\mu (h(x))^\alpha dt}{\kappa(\alpha)}.$$

By Remark 11, any distribution in this family is single-dipped. Take any such distribution h_α, and let

$$z_\alpha(\mu) \equiv h_\alpha(\mu)(1-\mu) - (1 - H_\alpha(\mu)).$$

Then, by the argument in the proof of Theorem 2, the set $\mathcal{L}_{z_\alpha} \equiv \{\mu \in [0,1] : z_\alpha(\mu) \leq 0\}$ has the following form:

$$\mathcal{L}_{z_\alpha} = [\hat{\mu}_\alpha, 1].$$

The proof goes through by showing that $\hat{\mu}_\alpha$ is decreasing in α. We start with two important remarks.

Remark 13. $z'_\alpha(\hat{\mu}_\alpha) < 0$. This follows from the fact that $\mathcal{L}_{z_\alpha} = [\hat{\mu}_\alpha, 1]$. Then, $z_\alpha(\mu)$ crosses zero from above at $\hat{\mu}_\alpha$. Since $z_\alpha(\mu)$ is differentiable, the remark follows.

Remark 14. *If $\hat{\mu}_\alpha \in (0,1)$, then $z_\alpha(\hat{\mu}_\alpha) = 0$, or equivalently,*

$$h_\alpha(\hat{\mu}_\alpha)(1-\hat{\mu}_\alpha) = \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)dx. \tag{43}$$

By Remark 14, $\hat{\mu}_\alpha \in (0,1)$ satisfies

$$z_\alpha(\hat{\mu}_\alpha) = 0.$$

Implicitly differentiate with respect to α to get

$$\frac{\partial}{\partial \alpha} z_\alpha(\hat{\mu}_\alpha) + z'_\alpha(\hat{\mu}_\alpha)\frac{\partial \hat{\mu}_\alpha}{\partial \alpha} = 0.$$

By Remark 13, $z'_\alpha(\hat{\mu}_\alpha) < 0$. Then, $\frac{\partial \hat{\mu}_\alpha}{\partial \alpha} \leq 0$ if and only if

$$\left.\frac{\partial}{\partial \alpha} z_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \leq 0.$$

Note that, for any $\mu \in [0,1]$,

$$\frac{\partial}{\partial \alpha} z_\alpha(\mu) \leq 0 \iff \frac{\partial}{\partial \alpha} h_\alpha(\mu)(1-\mu) \leq \int_\mu^1 \frac{\partial}{\partial \alpha} h_\alpha(x) dx.$$

Recall that $h_\alpha(\mu) = \frac{(h(\mu))^\alpha}{\kappa(\alpha)}$. Therefore, for any x for which $h_\alpha(x) > 0$, we have $\frac{\partial}{\partial \alpha} h_\alpha(x) = h_\alpha(x)\left(\log h(x) - \frac{\kappa'(\alpha)}{\kappa(\alpha)}\right)$, and so,

$$\left.\frac{\partial}{\partial \alpha} z_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \leq 0 \iff h_\alpha(\hat{\mu}_\alpha)\log h(\hat{\mu}_\alpha)(1-\hat{\mu}_\alpha)) \leq \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log h(x) dx,$$

where we are using the fact that, by Remark 12, $h_\alpha(x) > 0$ almost everywhere. Using (43) to substitute $h_\alpha(\hat{\mu}_\alpha)(1-\hat{\mu}_\alpha) = \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)dx$ on the left hand side of the above inequality, we have

$$\begin{aligned}
\left.\frac{\partial}{\partial \alpha} z_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \leq 0 &\iff \log h(\hat{\mu}_\alpha) \int_{\hat{\mu}_\alpha}^1 h_\alpha(x) dx \leq \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log h(x) dx \\
&\iff \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log h(\hat{\mu}_\alpha) dx \leq \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log h(x) dx \\
&\iff \alpha \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log h(\hat{\mu}_\alpha) dx \leq \alpha \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log h(x) dx \\
&\iff \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log (h(\hat{\mu}_\alpha))^\alpha dx \leq \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log (h(x))^\alpha dx \\
&\iff \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log \left(\frac{(h(\hat{\mu}_\alpha))^\alpha}{(h(x))^\alpha}\right) dx \leq 0 \\
&\iff \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log \left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)}\right) dx \leq 0.
\end{aligned}$$

For any real number $z > 0$, $\log(z) \leq z-1$, with a strict inequality for any $z \neq 1$. Therefore,

$$\int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\log \left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)}\right) dx \leq \int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)} - 1\right) dx.$$

Therefore, $\left.\frac{\partial}{\partial \alpha} z_\alpha(\mu)\right|_{\mu=\hat{\mu}_\alpha} \leq 0$ as long as

$$\begin{aligned}
\int_{\hat{\mu}_\alpha}^1 h_\alpha(x)\left(\frac{h_\alpha(\hat{\mu}_\alpha)}{h_\alpha(x)} - 1\right) dx \leq 0 &\iff \int_{\hat{\mu}_\alpha}^1 \left(h_\alpha(\hat{\mu}_\alpha) - h_\alpha(x)\right) dx \leq 0 \\
&\iff \int_{\hat{\mu}_\alpha}^1 h_\alpha(\hat{\mu}_\alpha) dx \leq \int_0^{\hat{\mu}_\alpha} h_\alpha(x) dx \\
&\iff h_\alpha(\hat{\mu}_\alpha)(1-\hat{\mu}_\alpha) \leq \int_{\hat{\mu}_\alpha}^1 h_\alpha(x) dx,
\end{aligned}$$

which is guaranteed by (43). We conclude that for any $\hat{\mu}_\alpha \in (0,1)$, $\frac{\partial \hat{\mu}_\alpha}{\partial \alpha} \leq 0$.

Since h_1 and h_2 are within the family we considered (with h_1 corresponding to $\alpha = 1$ and h_2 corresponding to some $\alpha \geq 1$), $\hat{\mu}_2 \leq \hat{\mu}_1$. Repeating the same argument in the proof of Theorem 2, we conclude that $\sigma_2^1 \geq \sigma_1^1$. □

References

Abramowitz, Alan I., and Kyle L. Saunders. 2008. "Is Polarization a Myth?" *The Journal of Politics*, 70(2): 542–555.

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva. 2020. "The Polarization of Reality." *AEA Papers and Proceedings*, 110 324–28.

Alonso, Ricardo, and Odilon Câmara. 2016a. "Persuading Voters." *American Economic Review*, 106(11): 3590–3605.

Alonso, Ricardo, and Odilon Câmara. 2016b. "Bayesian Persuasion with Heterogeneous Priors." *Journal of Economic Theory*, 165 672–706.

Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder Jr. 2006. "Purple America." *Journal of Economic Perspectives*, 20(2): 97–118.

Arieli, Itai, and Yakov Babichenko. 2019. "Private Bayesian Persuasion." *Journal of Economic Theory*, 182 185–217.

Bagnoli, Mark, and Ted Bergstrom. 2005. "Log-Concave Probability and Its Applications." *Economic Theory*, 26(2): 445–469.

Bail, Christopher A., Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M.B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. "Exposure to Opposing Views on Social Media Can Increase Political Polarization." *Proceedings of the National Academy of Sciences*, 115(37): 9216–9221.

Bardhi, Arjada, and Yingni Guo. 2018. "Modes of Persuasion Toward Unanimous Consent." *Theoretical Economics*, 13(3): 1111–1149.

Baron, David P. 2006. "Persistent Media Bias." *Journal of Public Economics*, 90(1-2): 1–36.

Bergemann, Dirk, and Stephen Morris. 2019. "Information Design: A Unified Perspective." *Journal of Economic Literature*, 57(1): 44–95.

Bernhardt, Dan, Stefan Krasa, and Mattias Polborn. 2008. "Political Polarization and the Electoral Effects of Media Bias." *Journal of Public Economics*, 92(5-6): 1092–1104.

Besley, Timothy, and Andrea Prat. 2006. "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." *American Economic Review*, 96(3): 720–736.

Boleslavsky, Raphael, Mehdi Shadmehr, and Konstantin Sonin. 2021. "Media Freedom in the Shadow of a Coup." *Journal of the European Economic Association*, 19(3): 1782–1815.

Burke, Jeremy. 2008. "Primetime Spin: Media Bias and Belief Confirming Information." *Journal of Economics & Management Strategy*, 17(3): 633–665.

Campante, Filipe R., and Daniel A Hojman. 2013. "Media and Polarization: Evidence from the Introduction of Broadcast TV in the United States." *Journal of Public Economics*, 100 79–92.

Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang. 2019. "Pivotal Persuasion." *Journal of Economic Theory*, 180 178–202.

Chan, Jimmy, and Wing Suen. 2008. "A Spatial Theory of News Consumption and Electoral Competition." *The Review of Economic Studies*, 75(3): 699–728.

Chen, Heng, and Wing Suen. 2021. "Competition for Attention and News Quality."Technical report, Working Paper.

Chiang, Chun-Fang, and Brian Knight. 2011. "Media Bias and Influence: Evidence from Newspaper Endorsements." *The Review of Economic Studies*, 78(3): 795–820.

Chopra, Felix, Ingar Haaland, and Christopher Roth. 2021. "Do People Value More Informative News?" *Working Paper.*

Corneo, Giacomo. 2006. "Media Capture in a Democracy: The Role of Wealth Concentration." *Journal of Public Economics*, 90(1-2): 37–58.

DellaVigna, Stefano, and Ethan Kaplan. 2007. “The Fox News Effect: Media Bias and Voting.” *The Quarterly Journal of Economics*, 122(3): 1187–1234.

DiMaggio, Paul, John Evans, and Bethany Bryson. 1996. “Have Americans’ Social Attitudes Become More Polarized?” *American Journal of Sociology*, 102(3): 690–755.

Duggan, John, and Cesar Martinelli. 2011. “A Spatial Theory of Media Slant and Voter Choice.” *The Review of Economic Studies*, 78(2): 640–666.

Edmond, Chris. 2013. “Information Manipulation, Coordination, and Regime Change.” *Review of Economic Studies*, 80(4): 1422–1458.

Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. “Why Resource-Poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data.” *American Political Science Review*, 103(4): 645–668.

Fiorina, Morris P., and Samuel J. Abrams. 2008. “Political Polarization in the American Public.” *Annu. Rev. Polit. Sci.*, 11 563–588.

Flaxman, Seth, Sharad Goel, and Justin M. Rao. 2016. “Filter Bubbles, Echo Chambers, and Online News Consumption.” *Public Opinion Quarterly*, 80(S1): 298–320.

Fowler, Anthony, and Kisoo Kim. 2022. “An Information-Based Explanation for Partisan Media Sorting.” *Working Paper.*

Gehlbach, Scott, and Konstantin Sonin. 2014. “Government Control of the Media.” *Journal of Public Economics*, 118 163–171.

Gentzkow, Matthew. 2016. “Polarization in 2016.” *Toulouse Network of Information Technology White Paper.*

Gentzkow, Matthew, and Jesse M. Shapiro. 2006. “Media Bias and Reputation.” *Journal of Political Economy*, 114(2): 280–316.

Gentzkow, Matthew, and Jesse M. Shapiro. 2010. “What Drives Media Slant? Evidence from U.S. Daily Newspapers.” *Econometrica*, 78(1): 35–71.

Gentzkow, Matthew, and Jesse M. Shapiro. 2011. “Ideological Segregation Online and Offline.” *The Quarterly Journal of Economics*, 126(4): 1799–1839.

Gentzkow, Matthew, Jesse M. Shapiro, and Daniel F. Stone. 2015. “Media Bias in the Marketplace: Theory.” In *Handbook of Media Economics.* 1: Elsevier, 623–645.

Glaeser, Edward L., and Bryce A. Ward. 2006. “Myths and Realities of American Political Geography.” *Journal of Economic Perspectives*, 20(2): 119–144.

Groseclose, Tim, and Jeffrey Milyo. 2005. “A Measure of Media Bias.” *The Quarterly Journal of Economics*, 120(4): 1191–1237.

Guriev, Sergei, and Daniel Treisman. 2020. “A Theory of Informational Autocracy.” *Journal of Public Economics*, 186, p. 104158.

Herrera, Helios, and Ravideep Sethi. 2022. “Identity-Based Elections.” *Working Paper.*

Inostroza, Nicolas, and Alessandro Pavan. 2022. “Persuasion in Global Games with Application to Stress Testing.” *Working Paper.*

Iyengar, Shanto, and Kyu S. Hahn. 2009. “Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use.” *Journal of Communication*, 59(1): 19–39.

Kamenica, Emir, and Matthew Gentzkow. 2011. “Bayesian Persuasion.” *American Economic Review*, 101(6): 2590–2615.

Kerman, Toygar, P. Jean-Jacques Herings, and Dominik Karos. 2021. “Persuading Strategic Voters.” *Working Paper.*

Klein, Ezra. 2020. *Why We’re Polarized.* New York, NY: Simon and Schuster.

Kolotilin, Anton. 2015. “Experimental Design to Persuade.” *Games and Economic Behavior*, 90 215–226.

Kolotilin, Anton, Tymofiy Mylovanov, and Andriy Zapechelnyuk. 2021. "Censorship as Optimal Persuasion." *forthcoming in Theoretical Economics.*

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica*, 85(6): 1949–1964.

Kosterina, Svetlana. 2021. "Persuasion with Unknown Beliefs." *forthcoming in Theoretical Economics.*

Laclau, Marie, and Ludovic Renou. 2017. "Public Persuasion." *Working Paper.*

Larcinese, Valentino, Riccardo Puglisi, and James M. Snyder Jr. 2011. "Partisan Bias in Economic News: Evidence on the Agenda-Setting Behavior of US Newspapers." *Journal of Public Economics*, 95(9-10): 1178–1189.

Levendusky, Matthew S. 2013. "Why Do Partisan Media Polarize Viewers?" *American Journal of Political Science*, 57(3): 611–623.

Lin, Xiao, and Ce Liu. 2021. "Credible Persuasion." *Working Paper.*

Lott, John R., and Kevin A. Hassett. 2014. "Is Newspaper Coverage of Economic Events Politically Biased?" *Public Choice*, 160(1): 65–108.

Martin, Gregory J., and Ali Yurukoglu. 2017. "Bias in Cable News: Persuasion and Polarization." *American Economic Review*, 107(9): 2565–99.

Mathevet, Laurent, Jacopo Perego, and Ina Taneva. 2020. "On Information Design in Games." *Journal of Political Economy*, 128(4): 1370–1404.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches.* Cambridge, MA: MIT Press.

Milgrom, Paul, and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica* 157–180.

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review*, 95(4): 1031–1053.

Oliveros, Santiago, and Felix Várdy. 2015. "Demand for Slant: How Abstention Shapes Voters' Choice of News Media." *The Economic Journal*, 125(587): 1327–1368.

Perego, Jacopo, and Sevgi Yuksel. 2022. "Media Competition and Social Disagreement." *Econometrica*, 90(1): 223–265.

Petrova, Maria. 2008. "Inequality and Media Capture." *Journal of Public Economics*, 92(1-2): 183–212.

Prat, Andrea. 2015. "Media Capture and Media Power." In *Handbook of Media Economics.* eds. by Simon P. Anderson, Joel Waldfogel, and David Strömberg, 1 of Handbook of Media Economics: North-Holland, 669 – 686.

Prior, Markus. 2013. "Media and Political Polarization." *Annual Review of Political Science*, 16 101–127.

Puglisi, Riccardo, and James M. Snyder Jr. 2011. "Newspaper Coverage of Political Scandals." *The Journal of Politics*, 73(3): 931–950.

Shadmehr, Mehdi, and Dan Bernhardt. 2015. "State Censorship." *American Economic Journal: Microeconomics*, 7(2): 280–307.

Shaked, Moshe, and J George Shanthikumar. 2007. *Stochastic Orders.* New York, NY: Springer Science & Business Media.

Stroud, Natalie Jomini. 2010. "Polarization and Partisan Selective Exposure." *Journal of Communication*, 60(3): 556–576.

Sun, Junze, Arthur Schram, and Randolph Sloof. 2022. "Public Persuasion in Elections: Single-Crossing Property and the Optimality of Censorship." *Working Paper, Available at SSRN 4028840.*

Sunstein, Cass R. 2009. *Going to Extremes: How Like Minds Unite and Divide.* New York, NY: Oxford University Press.

Taneva, Ina. 2019. "Information Design." *American Economic Journal: Microeconomics*, 11(4): 151–85.

Wang, Yun. 2015. "Bayesian Persuasion with Multiple Receivers." *Working Paper.*