Optimal News Bias with External Information[*]

Raquel Lorenzo Vidal[†]   Paul O. Richter[‡]

November 19, 2024

Abstract

Individuals are often exposed to information they did not actively seek, such as news shared by others, raising the question of how such information environments shape personal information choices. This paper studies how expectations of external information influence agents' choices of news bias. Extending a standard model of Bayesian learning from biased sources to account for the anticipation of additional information, we show that expected information critically impacts news bias choices. We characterize the optimal learning strategy depending on the decision maker's prior belief and the structure of the additional information, offering a novel explanation for why people often consume like-minded media news while also engaging with opposing ones. Applying this to social contexts, we find that highly uncertain agents tend to coordinate on the same news bias, whereas relatively certain individuals may opt for opposing ones. We also shed light on how to foster information acquisition among agents with more extreme beliefs.

Keywords: Learning, Expectations, Media-bias, Communication

JEL Codes: D83, D84, D71, L82, O33

1   Introduction

Concerns that the Internet may intensify ideological segregation have been greatly empha-sized (Sunstein 2001; Bishop 2009). In light of this, many papers have explored potential mechanisms behind eco-chambers, characterized by individuals segregating with like-minded peers, who influence their beliefs in a non rational manner (Levy and Razin 2019). How-ever, empirical evidence on the prevalence of online echo-chambers is mixed. While social media often promotes exposure to like-minded information, extreme echo-chamber patterns are rare, and few people have heavily skewed information diets (Gentzkow and Shapiro 2011, Bakshy et al. 2015, Flaxman et al. 2016, Boxell et al. 2017, Dubois and Blank 2018, Nyhan et al. 2023). Both online and offline, people are regularly exposed to information shared by others (Wojcieszak and Mutz 2009, Minozzi et al. 2020). This raises the question of how people's information environments affect their choice of news bias.

Most papers studying information choices either assume agents either have full control over the information that they process or no control at all. In this paper, we study the effect of expecting information beyond one's control on agents' choices of news bias. By adding the expectation of additional information to a standard Bayesian learning model, we show that agents choices of news bias are highly dependent on anticipated external information. Our model captures well-documented patterns in news consumption, including the tendency to consume both like-minded and opposing media or to align choices of news bias with peers.

In our framework, a Bayesian decision-maker (DM) selects an action, left or right, to match an unknown true state, such as voting for the correct policy. Prior to making this decision, the DM can acquire information from one of two news sources (signals), each biased towards a particular action –i.e., left- or right-biased. After choosing a source but before acting, the DM also receives additional information from an independent exogenous source, of known structure, such as information shared by others.[1]

In this setting, the DM's problem can be framed as assessing the value of further learning from each news source at each of the beliefs she expects the exogenous source to take her. At more uncertain beliefs further information is more valuable, as choosing the wrong action is considered more likely. Hence, the DM's preferences for news at more frequent and uncertain beliefs tend to be more important in determining her choice of source.

First, we show that the mere expectation of receiving additional information significantly shapes optimal information choices. As long as neither source is objectively better than the other (i.e., yields higher expected utility across all priors), there always exists an independent anticipated signal that can alter the DM's choice between sources, regardless of her prior.

---

[1]By independent sources we refer to independent signals, conditional on the state.

This holds broadly, for sources with finite messages in finite-state, finite-action settings. For news consumption, this implies that expecting information, such as incidental political insights from peers, can be sufficient to affect media choices. It also means that beliefs alone cannot fully predict agents' learning choices when additional information is expected.

Second, our model rationalizes both *own-biased learning* (learning from sources biased towards the state one deems more likely) and *opposite-biased learning* (choosing sources biased towards the state one deems less likely). This is consistent with empirical findings showing that although individuals frequently consume like-minded news, many also seek sources with opposing views (e.g., see Gentzkow and Shapiro 2011). The optimality of one learning heuristic over the other depends on both prior beliefs and the nature of anticipated additional information. In particular, the DM's optimal strategy is characterized by intervals of priors, each corresponding to a different source choice, where both the source selected and interval thresholds depend critically on the expected additional information.

A relatively uncertain DM expecting additional information of similar quality to her news sources will optimally choose a source with the same bias as the additional information. Namely, if she anticipates receiving left-biased information from peers, she chooses left-biased news too, and similarly for right-biased information.[2] This offers a rationale for certain levels of ideological homophily observed in social media friend networks (Bakshy et al. 2015, Barberá et al. 2015), suggesting it may stem not only from a preference for like-minded individuals but also from optimal learning choices. Friends with sufficiently uncertain political views, even if leaning towards different directions, may choose media with the same bias (either left-biased or right-biased) just for learning optimally, which in turn makes it more likely that people within the same social circle end up sharing similar views.

A DM who is extremely sure about the state is indifferent between news sources, since none can change her action; while if she is very sure but can still learn something valuable, she chooses own-biased learning. The prior where she shifts from indifference to strictly preferring own-biased news varies depending on the type of expected additional information. In particular, an extreme agent might start seeking information as the anticipated information becomes of better quality or more credible at sending information that contradicts her prior. These insights can be used to understand how to incentivize information acquisition. For instance, to promote information seeking by COVID vaccine skeptics, a policy maker might want to prioritize improving the perceived accuracy of informational campaigns at identifying when the vaccine does not work as opposed to when it does. Or, a media outlet trying to

---

[2]Instead, when the additional information Blackwell dominates the sources available to the DM (better quality), the DM is indifferent between sources; and if it is Blackwell dominated (worse quality), the DM's incentives resemble a situation without additional information and she prefers own-biased learning.

penetrate the market of extreme agents who are not consuming any news, may want to offer free content, since the expectation of receiving it can incentivize some to acquire information.

Under mild conditions on the additional information, a moderately certain DM finds it optimal to choose opposite-biased learning. For instance, a moderately left-leaning voter might, expecting left-biased information from her social group, still choose right-biased news. Studying how the threshold that separates the region of very certain priors at which the DM prefers own-biased learning from where she prefers opposite-biased learning, we show tha there is a range of priors where the DM finds it optimal to mismatch the bias of the exogenous source. For instance, expecting right-biased information, a DM might choose left-biased news, and vice versa. We apply this result to understand the optimal learning choices of two teams, who, together, need to decide whether to adopt a new technology. Even if both teams believe that the technology has potential to improve the company's outcomes, they may behave as *advocates* of opposite alternatives (Dewatripont and Tirole 1999).

Finally, using these patterns as its main building blocks, we fully characterize the DM's optimal learning strategy. We also discuss several applications of our results, which are relevant in other learning environments where agents only have partial control over the information they receive. Similar to voters who choose news sources yet cannot fully filter information shared by their social networks, doctors need to select diagnostic tests while expecting to observe patient symptoms in the future, and CEOs design their market research anticipating additional insights from competitors' market outcomes.

1.1   Illustrative Example

Ann must decide whether to vote for the left or right policy tomorrow. Voting for the correct policy gives her a payoff of 1 util, and voting incorrectly gives 0. She believes the left policy is correct with probability 0.6, making her *left-biased*.

Today, Ann can choose to consult one of two news sources, each reporting either "left" or "right". One source is left-biased and the other is right-biased. If the correct policy is left, the left-biased source reports "left" with 80% probability and "right" only with 50% probability if the correct policy is right. Similarly, the right-biased source reports "right" with 80% probability right is the correct policy and "left" with 50% probability when left is. Thus, each source matches the correct policy more often when it aligns with its bias: the left-biased source reports more accurately for the left policy, and the right-biased for the right policy. However, a biased source contradicting its typical tendency (e.g., the left-biased source says "right") is stronger evidence about the correct policy.

Ann understands how the two news sources generate their messages. Therefore, she

chooses the source that maximizes the value of her vote tomorrow, given her prior about the correct policy and the way sources report in each state. If she made the choice in isolation, the left-biased source would be her optimal choice. We call this *own-biased learning*.

However, Ann does not learn in isolation. She knows that her friend Bob will share news with her before the election and expects Bob to share information from the right-biased source. Assuming that, conditional on the correct policy, Bob's report is independent of the report Ann would receive, what would be Ann's optimal choice of news bias?[3]

If Bob reports "left", Ann becomes quite confident (79%) that the correct policy is left, and further information that she can choose will not change her vote. But if he reports "right", her beliefs shift towards the right policy being correct (52%), though she remains uncertain enough that further information is valuable, since it can affect her vote. While at the first (left-biased) interim belief she is indifferent between any news source, at the latter (right-biased) interim belief she would prefer the right-biased source. All in all, Ann finds it optimal to select the right-biased source today. Since her prior is left-biased, this means that she chooses *opposite-biased learning*.

This simple example showcases how the expectation of future information can alter optimal learning decisions. It breaks with the standard prediction from the literature on Bayesian learning, which suggests that agents will generally prefer sources biased toward the state they initially consider more likely (own-biased learning).

1.2   Literature Review

This paper contributes to the literature on optimal Bayesian learning from biased sources. Papers showing the optimality of choosing an own-biased source for a Bayesian agent are Calvert 1985, Meyer 1991, Suen 2004, Burke 2008, Gentzkow and Shapiro 2006 and Mullainathan and Shleifer 2005. Our predictions are consistent with their finding when no additional information is expected. But our results show that with additional information opposite-biased learning can also be optimal. In contrast, Oliveros and Várdy (2015) find that voters who can abstain may prefer unbiased sources. Instead, we focus on two actions.

Recent work using dynamic models demonstrates that it can be optimal for a Bayesian agent to multi-home or learn from opposite-biased sources (Che and Mierendorff 2019, Nikandrova and Pancs 2018, Mayskaya 2024). In these models, an agent faces an optimal stopping problem with costly information acquisition. There, opposite-biased learning is chosen due to a trade-off between accuracy and delay. By offering stronger evidence of the state deemed more likely, opposite-biased learning can reduce future information costs for highly uncertain

---

[3]If reading the same type of news implies getting the same report, Ann's optimal source choice is trivial.

agents, who require a higher rise confidence to stop acquiring information. In these models the prior belief fully determines the optimal learning strategy. In our model additional information plays a critical role in determining the DM's information choice at any prior and in explaining the optimality of opposite-biased learning.

Gossner et al. (2021) and Liang et al. (2022) study the effect of exogenously manipulating attention of a DM who dynamically learns about the value of different items. Our model is different in several ways: (i) their DM chooses among signals about multiple independent states, while ours chooses among signals about a binary state, (ii) they consider an optimal stopping problem, while our problem has a given stopping time, (iii) they study the effect of manipulating attention at a given point in time on consequent learning choices, while we study how expected manipulations in the future affect current learning choices.

Recent developments in the literature on the value of signals, building on Blackwell et al. (1951), are conceptually related. Börgers et al. (2013) identify conditions under which certain signals increase or decrease each other's value, regardless of the decision problem. Brooks et al. (2023) develop a signal ordering robust to the presence of additional signals and decision problem. Our analysis focuses on the choice between two signals of comparable informativeness –choices depend on the DM's prior and the structure of additional information– , while they study signal dominance across all decision problems. Therefore, their analysis restricts to correlated signals while we focus on independent ones (conditional on the state).

The information design literature also recently considers how the value of one signal structure changes in the presence of another independent signal. Laclau et al. (2017); Kolotilin et al. (2017) and Dworczak and Pavan (2022) study the problem of a persuader who is uncertain about the receiver's private information. Bergemann et al. (2018) consider a data seller offering a menu of signals to screen buyers based on their beliefs. While in these models a DM chooses information –i.e., a signal for the receiver or a menu for buyers– taking into account posterior distributions, our work differs in the DM's goal: selecting a state-dependent action rather than persuading of a state-independent action or screening agents.

2   Learning Model

A Bayesian decision-maker (DM) must choose an action, $A^x \in \{A^L, A^R\}$, trying to match an unknown state $\theta \in \{L, R\}$. Her payoff is 1 if the chosen action matches the state ($x = \theta$) and 0 otherwise. Her prior is denoted by $p_0 \in (0, 1)$, which refers to the probability of state being R. Before taking action, the DM obtains information from two sources: one chosen, σ_c, and other given exogenously, σ_e. The first represents the information she can actively select and the second any additional information she expects to receive beyond her control.

Sources of information. A source σ^x can send two possible messages, l or r, and is characterized by two parameters: the probability of sending message l when the state is L, $\pi^x(l|L)$, and the probability of sending message r when the state is R, $\pi^x(r|R)$. Formally, a news source is a binary signal structure. This is illustrated in Table 1 below.

State/Message	l	r
$\theta = L$	$\pi^x(l|L)$	$1 - \pi^x(l|L)$
$\theta = R$	$1 - \pi^x(r|R)$	$\pi^x(r|R)$

Table 1: Signal structure of a source of information σ^x

Without loss of generality we assume l is the message that updates the DM's prior towards believing state L is more likely, i.e., $\pi^x(l|L) \geq 1 - \pi^x(r|R)$. A source is uninformative if $\pi^x(l|L) = 1 - \pi^x(r|R)$, fully revealing about state R if $\pi^x(l|L) = 1$ and fully revealing about state L if $\pi^x(r|R) = 1$. To avoid uninteresting limit cases, we restrict attention to sources with $\pi^x(l|L) > 1 - \pi^x(r|R)$ and $\pi^x(l|L) < 1$, $\pi^x(r|R) < 1$.

Since we are interested in studying choices among biased sources, we formalize this concept within our framework. The binary structure offers a clear interpretation of bias, as explained below. This is in line with other approaches to model biasedness (Che and Mierendorff 2019), high-bar or low-bar experiments (Gans 2023) or information skewness (Masatlioglu et al. 2023), which also focus on binary signals.

Definition 1 *For any source σ^x, characterized by $\pi^x(l|L)$, and $\pi^x(r|R)$:*

  i) *The source is right-biased if $\pi^x(l|L) < \pi^x(r|R)$,*

 ii) *The source is left-biased if $\pi^x(l|L) > \pi^x(r|R)$,*

iii) *The source is unbiased if $\pi^x(l|L) = \pi^x(r|R)$.*

A biased source is more likely to match the true state when it aligns with its bias; while an unbiased source matches the state with equal probability, regardless of the true state. From the perspective of a DM with a neutral prior, a right-biased source is expected to send a right message more often than a left message, and similarly, a left-biased source is more likely to send a left message than a right message. Furthermore, she would update her prior further upon receiving a left message from the right-biased source or a right message from the left-biased source than when receiving a message that agrees with the source's bias.

To isolate the effect of bias from that of informativeness, we consider a choice set containing two symmetric signals. The following definition formalizes the concept of symmetry.[4]

Definition 2 *Two sources σ^x and σ^y are symmetric if $\pi^x(l|L) = \pi^y(r|R)$ and $\pi^x(r|R) = \pi^y(l|L)$.*

When two sources are symmetric, the probability that their messages match the state is the same, conditional on the state either aligning or misaligning with their bias. However, their accuracy might differ between aligned and misaligned states. For two symmetric biased sources, when the both states are equally likely, the left-biased source sends message l (r) with the same probability as the right-biased source sends message r (l). If two symmetric sources are unbiased, they are effectively the same source.

We assume that signals are independent conditional on the state. Therefore, if the DM's choice of source depends on σ_e, it is not due to correlation. While in reality messages from different sources may exhibit correlation beyond the state, we abstract from this to isolate the effect of other features of the sources' structure, such as bias and informativeness.

Timing, information and beliefs. First, Nature draws a state θ. Then, the DM chooses among two symmetric biased sources, where σ^R denotes the right-biased source and σ^L the left-biased. After choosing her source, $\sigma_c \in \{\sigma^R, \sigma^L\}$, the DM observes two (independent) messages: one from the chosen source σ_c, and one from the exogenous source, σ_e. Finally, after updating her beliefs with the information received from each source, the DM chooses an action A^x and her payoff is realized.

| Nature draws θ | DM chooses σ_c | DM sees σ_e and σ_c's messages | DM chooses A^x |

Figure 1: Timing of the learning problem

When selecting σ_c, the DM knows the σ_e's structure –i.e., the probability of each message given the state–, but not the actual message. She also knows she will observe this message after selecting σ_c but before choosing A^x. Our analysis is robust to other timing assumptions as long as these two aspects are maintained. If the DM observed σ_e's message before selecting a source, she would simply update her belief and choose as if learning in isolation. Conversely, if σ_e's message arrived after taking action, it would be irrelevant for decision-making and thus

---

[4]Note the difference with respect to Masatlioglu et al. (2023)'s definition of symmetry, which refers to unbiasedness in our language.

not affect her source choice. In our framework, however, since the DM receives the exogenous information after selecting a source and before choosing an action, her informational choice depends on how well σ_c complements the expected information from σ_e.[5]

We denote the updated belief of the DM after observing message s from source σ^x alone by $p(s^x)$ and the updated belief after observing message s from source σ^x and message m from source σ^y by $p(s^x, m^y)$. Both of them refer to the probability that the state is R.

2.1 Preliminaries

Before anakyzing the DM's optimal learning strategy in Section 3, we discuss how to study the DM's problem. One approach is to treat the chosen information source and the exogenous source as a bundle, which together form one aggregate source with four possible messages: (ll, lr, rl, rr). In that case, the DM chooses between two such bundles, each containing a fixed exogenous component, i.e. (σ^L, σ_e) or (σ^R, σ_e). However, this complicates the analysis of how different features of the exogenous source impact the optimal choice of source.

To address this, we decompose the DM's problem in two steps. First, we calculate all possible posteriors resulting from the DM observing messages solely from the exogenous source. We refer to these posteriors as *interim posteriors*. Second, we assess the difference in expected value between the left- and right-biased sources at each of those interim posteriors, without the presence of any additional information. This step captures the extent to which the DM prefers one biased source over the other in isolation, akin to analyses in the literature that exclude exogenous information.

Thus, the DM's problem can be viewed as computing a weighted average of the belief-dependent preferences over the left- and right-biased source across all possible interim posteriors. The weights correspond to the probabilities of each interim posterior, induced by the exogenous source. Observation 1 below formalizes this perspective.

Observation 1 *For any exogenous source σ_e,*

$$EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) \geq 0 \iff \sum_{s \in \{l,r\}} \mathbb{P}(s^e|p_0)\Big(EU(\sigma^L|p(s^e)) - EU(\sigma^R|p(s^e))\Big) \geq 0,$$

where $\mathbb{P}(s^e|p_0) = p_0\pi^e(s^e|R) + (1-p_0)\pi^e(s^e|L)$ is the probability that σ_e sends s, given prior p_0. $EU(\cdot, \cdot|p)$ is the expected utility, from an ex-ante perspective, of receiving information from a bundle of sources given prior p. $EU(\cdot|p)$ is the ex-ante expected utility of receiving information from a single source, without expecting further information, given prior p.

---

[5]Through the lens of Börgers et al. (2013)'s framework all the signals in our environment are substitutes. However, in our specific decision problem, a signal can lower or increase the marginal value of another.

This approach can be easily represented graphically in the binary case.[6] Figure 2 below illustrates the ex-ante expected utility of receiving information from two symmetrically biased sources. The red curve represents the expected utility of receiving a message from a right-biased source σ^R at different prior beliefs (x-axis), absent of an exogenous source. The blue curve corresponds to its symmetric left-biased source σ^L. The gray curve shows the expected utility in the absence of any information.


Figure 2: Expected utility of two symmetric signals, σ^L and σ^R

Both the red and blue curves consist of three linear segments. The two outer segments correspond to priors where the DM's action is unaffected by the source's message. At these priors, she chooses as if she received no information. For priors near 0, the DM is so confident that the state is L that she always chooses A^L, regardless of the source's message. Similarly, for priors near 1, the DM is so confidence that the state is R that she always chooses A^R. The middle segment corresponds to priors where the DM is uncertain about the best action to take and *follows the source's recommendation* by choosing A^L if the message is l and A^R if the message is r. Moreover, the gain in expected utility due to information is larger for more uncertain priors. This can be seen by comparing the blue and red curves with the gray one. Intuitively, information is more valuable when the DM is more unsure about what is the best action to take. All of this is true for any binary signal.

Since the sources are symmetric, the left-biased source is comparatively more accurate when the true state is L. Therefore, if the DM finds L more likely according to her prior, she gets weakly greater expected utility from choosing the left-biased source. Similarly, if the DM finds R more likely, the right-biased source provides higher expected utility. This goes along with existing literature showing that Bayesian agents prefer learning from sources

---

[6]Although we focus on binary signals and two states, this perspective of taking weighted averages over posteriors is easily generalized to multiple states and more complex signals.

biased toward the state they deem more likely (own-biased learning).

The violet curve in Figure 3 represents the difference in value between the left- and right-biased sources in Figure 2, without any additional information, for various beliefs of the DM. A similar graph would result from any pair of binary sources. In the absence of the exogenous source, the DM will optimally choose σ^L at priors where the curve is above 0 and σ^R at priors where it is below 0. Consistent with the principle of own-biased learning, the curve takes weakly positive values when the DM assigns higher probability to state L and weakly negative values when she favors state R. This graph, together with Observation 1, can serve as a useful tool for evaluating the expected difference in value between the left- and right-biased sources when information from an exogenous source is expected.

Since the DM is Bayesian, the expectation of her interim posteriors must equal her prior belief. In the case of a binary exogenous source, one interim belief will lie weakly to the left of the prior, while the other will lie weakly to the right. The probability of each interim belief is proportional to its relative distance from the prior. Consequently, the difference in expected value between the left- and right-biased sources, when expecting information from σ_e, corresponds to the height at which the DM's prior intersects the segment connecting the value of the violet curve at each of the two interim posteriors.


Figure 3: Example where σ_e changes DM's optimal choice of source

For example, consider a DM with a prior p_0 just below 0.5 and an exogenous source which results in interim posteriors $p(l^e)$ when it sends the left message and $p(r^e)$ when it sends the right message, as illustrated in Figure 3. To calculate the weighted average of the sources' relative values across these interim posteriors, one can draw a line connecting the curve at each posterior. The height of this line at the prior, represented by a blue dot, is the expected difference in value between the left- and right-biased sources, given the expectation of information from the exogenous source. In this case, the blue dot lies below 0, meaning that the right-biased source, σ^R, provides more value to the DM.

Although the DM considers state L more likely than R and would choose the left-biased

10

source without any additional information (own-biased learning), expecting exogenous information makes it optimal for the left-biased DM to choose the right-biased source (opposite-biased learning). Using our framework, it is easy to show that the expectation of information in the future can reverse the DM's optimal choice of source, even when this information is independent conditional on the state.[7] Unlike in standard models of optimal learning from biased sources, the prior alone no longer predicts the DM's optimal choice of source.

Proposition 1 below illustrates this point. It shows that, regardless if the DM's prior belief, expected independent information can always make one source optimal over another, as long as that source is optimal for at least one prior. This implies that exogenous information can change the DM's informational choice for any non-trivial choice between sources.[8]

Proposition 1 *Fix any two sources, σ^x and σ^y. If there exists a belief, p, s.t.*
$EU(\sigma^x|p) > EU(\sigma^y|p)$, *then,* $\forall p_0 \in (0,1)$, $\exists \sigma_e$ *s.t.* $EU(\sigma^x, \sigma_e|p_0) > EU(\sigma^y, \sigma_e|p_0)$

The proof, relegated to the appendix, can be extended to signals with non-binary message spaces and more general action and state spaces (see note in the appendix).

As explained in Observation 1, the DM understands that the exogenous source can lead to different interim posteriors, where further information (of her choice) may have different value. At very certain beliefs, additional information is less valuable because the DM assigns a low probability to choosing the wrong action. However, at more uncertain beliefs, further information becomes highly valuable. Therefore, the DM's preference for information at these uncertain beliefs plays a critical role in her choice of source.

This insight underpins the proof of Proposition 1. The most extreme example of an exogenous source that can reverse the DM's informational choice is one that either makes the DM fully certain of the state (where further information has no value) or places her at an interim posterior where she strictly prefers σ^x. Given such exogenous source, the DM prefers σ^x from an ex-ante perspective because, at full certainty, she is indifferent between sources, and otherwise, she strictly prefers σ^x. While this is a convenient extreme case, many other exogenous sources can similarly cause preference reversal.

This result highlights the importance of accounting for information beyond the DM's control when analyzing problems of information choice, as such exogenous information can significantly alter the DM's informational preferences. When the DM anticipates receiving exogenous information in the future, her prior alone does not fully determine her choice of information source. Instead, different features of the expected exogenous information

---

[7]When correlation between σ_e and the sources in the DM's choice set is allowed, finding σ_e that reverses the preferences is even easier. For instance, making σ_e fully correlated with the originally optimal source.

[8]If a source is less valuable for all priors, no prior will ever find it optimal. Therefore, the choice is trivial.

affect the DM's optimal learning strategy in distinct ways. Nevertheless, the prior remains a critical determinant of the DM's choice of source. Crucially, both the DM's prior and the structure of the exogenous information jointly dictate her optimal selection of information.

The next example illustrates how the DM's prior and the structure of the exogenous source interact to determine her learning choices. Many insights from this example will extend to the full characterization of the DM's optimal strategy, as reflected in Section 3.

2.2   Illustrative Example – Continuation

Consider again Ann choosing between two symmetric news sources: a left-biased, σ^L and a right-biased source σ^R, where $\pi^L(l|L) = \pi^R(r|R) = 0.8$ and $\pi^L(r|R) = \pi^R(l|L) = 0.5$. Rather than focusing on her choice for a given prior, we now analyze how her learning strategy varies with her prior.

First, consistent with previous literature and the analysis above, when Ann does not expect any additional information, she weakly prefers own-biased learning. As shown in Figure 4, she is indifferent between sources when her prior assigns more than $\frac{5}{7}$ probability to one policy. That is because for such priors no report from any of the available news sources would alter her vote. For priors within the intermediate range, Ann prefers the source biased towards the policy she deems more likely to be correct (own-biased learning).

$$—EU(\sigma^L|p_0) = EU(\sigma^R|p_0)$$
$$—EU(\sigma^L|p_0) > EU(\sigma^R|p_0)$$
$$—EU(\sigma^L|p_0) < EU(\sigma^R|p_0)$$

p_0 ├──────────────┼────────────┼──────────────┼──────────────────────┤

0 $\frac{2}{7}$ $\frac{1}{2}$ $\frac{5}{7}$ 1

Figure 4: Optimal choice of signal in isolation

p_0 ├────────┼──────────┼────┼────────┼──────────┼────────┤

0 $\frac{1}{5}$ $\frac{5}{13}$ $\frac{1}{2}$ $\frac{35}{61}$ $\frac{5}{7}$ $\frac{25}{29}$ 1

Figure 5: Optimal source choice, expecting an additional message from a right-biased source

Next, suppose Ann expects Bob to share news with her. In line with Proposition 1, this may affect the news she reads. Her optimal strategy now depends more intricately on

12

Figure 6: Optimal source choice, expecting an additional message from a left-biased source

her prior, partitioning it into six adjacent, non-symmetric intervals. The structure of Bob's information also plays a critical role in shaping her informational decision. Figures 5 and 6 illustrate her optimal learning choices under the expectation that Bob will share right-biased or left-biased news, respectively.

Recall from Observation 1 that Ann's problem can be broken down as follows. She first considers the interim posterior beliefs induced by each possible report from Bob. She then evaluates the value of consulting each news source at those interim posteriors. Finally, she weights these values by the probability of reaching each posterior. We will see how this type of reasoning will be helpful to understand Ann's behavior at each region of priors.

For very certain priors, Ann is indifferent between sources because no combination of reports (from her and Bob's source) can change her vote. In a slightly less certain range, Bob's information alone cannot affect what policy she considers more likely to be correct, but certain combinations of messages from both sources can. In this range, additional learning (beyond Bob's information) is valuable, and own-biased learning is optimal from the ex ante perspective, since her interim beliefs are still biased towards the same policy as her prior.

When Ann becomes more uncertain, her problem is more difficult. At these priors, Bob's report alone may shift her belief about the correct policy, making her preferences for news sources differ across interim posteriors. Her optimal choice thus depends on two factors: i) how uncertain she is at each interim posterior, which determines the value of further learning, and ii) the probability she assigns to each posterior. Our earlier example illustrates this: Ann's prior was left-biased ($p_0 = 0.4$), and, expecting right-biased news from Bob, she finds it optimal to choose the right-biased source herself, engaging in opposite-biased learning.

The last remark from this example is that certain features of the information that Ann expects to receive significantly influence her optimal learning strategy. An example of this is the bias of Bob's news. Comparing Figures 5 and 6, both cases exhibit two extreme regions of indifference, two adjacent regions of own-biased learning, and two intermediate regions of opposite-biased learning. However, the specific thresholds of these intervals shift depending on Bob's bias and this leads to somewhat interesting predictions. For example, around the neutral prior ($p_0 = 0.5$), Ann prefers the right-biased source if Bob provides right-biased news but switches to the left-biased source if Bob provides left-biased news. At moderately certain priors (e.g., $p_0 = \frac{1}{3}$ or $p_0 = \frac{2}{3}$) the pattern reverses: she prefers the right-biased source

13

if Bob provides left-biased news and the left-biased source if Bob provides right-biased news. Furthermore, compared to the case with no additional information, the range of priors where Ann strictly prefers a source is larger, especially when her prior aligns with Bob's news bias.

This exemplifies how the expectation of future information reshapes optimal learning decisions. Many features of Ann's optimal learning strategy extend to other settings where agents anticipate receiving additional information beyond their control. The following sections explain how these insights generalize and the logic behind the different regions.

3   The Optimal Learning Strategy

Before introducing the full characterization of the DM's optimal strategy, we first analyze the problem of a DM with different levels of certainty about the state of the world (i.e., priors closer to or further from the neutral prior). We identify key regions of priors where the DM's optimal choice of source follows consistent heuristics. These regions serve as building blocks for the complete charaterization of the optimal learning strategy, outlined in Section 4.

3.1   Learning Choices of an Uncertain DM

We first analyze the problem of a DM who is relatively uncertain about the state. When the DM's prior is close to full uncertainty three main heuristics describe her optimal choice of information source.

Matching the bias of the exogenous source – If the exogenous source is biased and moderately informative, the DM selects a source with the same bias. A biased exogenous source can either make an uncertain DM highly confident in the state opposing its bias (rarely) or only slightly confident in the state aligning with it (frequently). Learning is less valuable when the DM is very certain about the state, but more valuable when she remains uncertain. Since the DM often reaches the interim posterior where she is still uncertain and biased towards the same state as the exogenous source, this uncertain posterior dominates in shaping her ex ante optimal choice. As the DM prefers own-biased learning from the perspective of her interim posteriors, her ex ante optimal strategy is to select the source matching the exogenous source's bias.[9]

Figure 3 exemplifies this heuristic. A relatively uncertain DM with $p_0 = 0.48$ expects a right-biased exogenous source to either rarely suggest state being L ($p(l^e) = 0.2$) or, often suggest state being R ($p(r^e) = 0.6$). Because at the interim posterior biased towards R

---

[9]If the exogenous source is unbiased, by the same logic, a fully uncertain DM ($p_0 = 0.5$) is exactly indifferent between sources, while a DM with nearby priors selects her own-biased source.

the difference in value between sources is greater, and the DM expects to reach it more frequently, she selects the right-biased source.

Indifference for highly informative exogenous sources – When the exogenous source is significantly more informative than the sources available to the DM, the DM's action is fully determined by the exogenous source's message, making further learning payoff irrelevant. Figure 7 showcases this scenario, where both interim posteriors are so certain that additional information provides no value, leaving the DM indifferent between sources.


Figure 7: Example where σ_e makes DM indifferent

Own-biased learning for weak exogenous sources – If the exogenous source is significantly less informative than the available sources, the DM always follows the recommendation of her chosen source, making the exogenous information payoff irrelevant. In this case, the DM's problem is equivalent to having no exogenous information, resulting in own-biased learning. Figure 8 exemplifies this. At both interim posteriors, regardless of the selected source, the DM would choose the action aligning with her source's message. This means that the exogenous source's message does not affect her payoff relevant action in any contingency. More generally, as long as the interim posteriors remain within the central linear segment, the DM's action is unaffected by σ_e, leading her to choose own-biased learning.


Figure 8: Example where σ_e does not affect the DM's choice of source

15

The following proposition summarizes the conditions under which each heuristic applies.

Proposition 2 *Fix any σ_e and symmetric σ^L and σ^R. There exists $\epsilon > 0$ s.t. $\forall p_0 \in N_\epsilon \left(\frac{1}{2}\right)$,*

  i) *If neither σ_e Blackwell Dominates σ^L and σ^R nor σ^L and σ^R Blackwell Dominate σ_e, the DM strictly prefers the source with the same bias as σ_e.*

  ii) *If σ_e Blackwell Dominates σ^L and σ^R, the DM is indifferent between sources.*

  iii) *If σ^L and σ^R Blackwell Dominate σ_e, the DM strictly prefers own-biased learning.*

Figure 9 illustrates this result. The x-axis corresponds to the probability that the exogenous source is accurate in state L, and the y-axis corresponds to the probability that it is accurate in R. Due to our assumption that *messages mean what they say* –i.e., l leads the DM to put a higher probability in state L–, without loss of generality we can restrict attention to upper triangle above the diagonal, which covers all the space of binary signals. Sources in the diagonal are fully uninformative, while the source at $(1,1)$ is fully informative. Intuitively, as we approach $(1,1)$ the exogenous source becomes more informative. Right-biased sources lie above the 45-degree line, while left-biased sources fall below it.
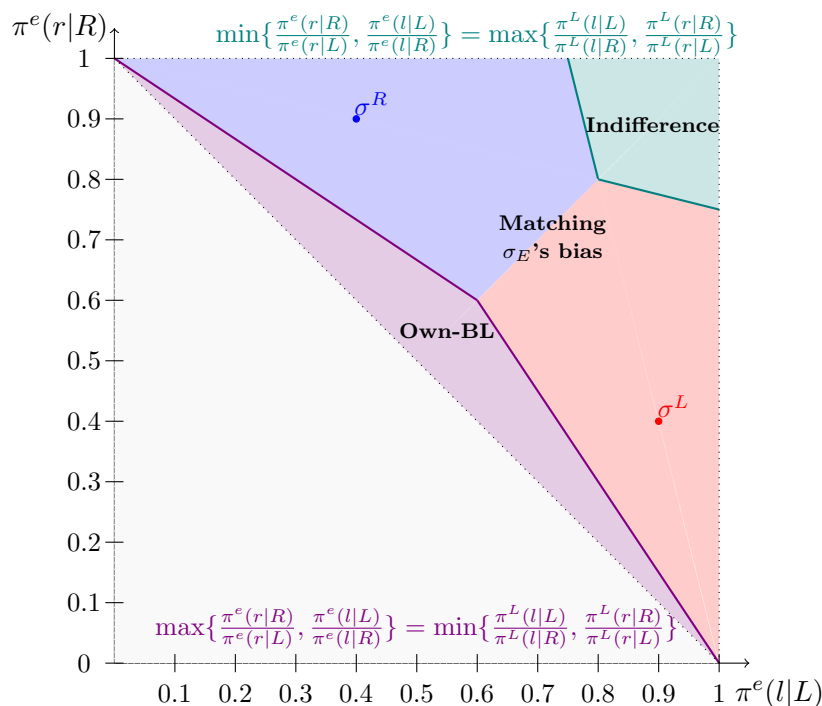

Figure 9: Illustration of Proposition 2

From Figure 9 one can see that Proposition 2 covers the entire space of sources. If the exogenous source is Blackwell more informative than the sources in the DM's choice set

(green region), further learning becomes irrelevant, leading to indifference. If the sources in the DM's choice set are Blackwell more informative than the exogenous source (violet region), the exogenous source has no impact on the DM's incentives, and she optimally follows own-biased learning. Otherwise, when the exogenous source is *comparably informative* (blue and red regions), the DM prefers the source matching the exogenous source's bias.

The figure also illustrates how Blackwell dominance translates to this framework. The following definition, which formalizes the notion of *informativeness of a message*, will be useful to understand this.

Definition 3 *The informativeness of a source σ^x's left message is given by $\frac{\pi^x(l|L)}{\pi^x(l|R)}$ and the informativeness of its right message is $\frac{\pi^x(r|R)}{\pi^x(r|L)}$.*

These likelihood ratios reflect how much the DM updates her beliefs: a higher $\frac{\pi^x(l|L)}{\pi^x(l|R)}$ increases the belief update towards state L upon observing l, and similarly a higher $\frac{\pi^x(r|R)}{\pi^x(r|L)}$ increases the belief update towards state R upon observing r. The ratios always exceed 1.

The following lemma restates Blackwell dominance in this framework: a source Blackwell Dominates the other if and only it is more informative across all messages.

Lemma 1 *A source σ^x Blackwell Dominates σ^y i.f.f.*

$$\min\left\{\frac{\pi^x(r|R)}{\pi^x(r|L)}, \frac{\pi^x(l|L)}{\pi^x(l|R)}\right\} > \max\left\{\frac{\pi^y(r|R)}{\pi^y(r|L)}, \frac{\pi^y(l|L)}{\pi^y(l|R)}\right\}.$$

3.1.1 Learning with peers: eco-chambers

Consider two agents who are unsure about whether to vote for the left policy, L, or the right policy, R. Each agent first chooses whether to consult a left-biased (σ^L) or right-biased (σ^R) news outlet. From this, they obtain a recommendation of what to vote (a message). After, they share these recommendations and each chooses what to vote. Their goal is to vote for the correct policy. In this game, the exogenous source is endogenized to be someone else's chosen source.

Following Proposition 2, when the agents are sufficiently uncertain, the game has two Nash equilibria: one where both of them choose the right-biased source, and another where both choose the left-biased source. Even if one agent slightly favors L and the other slightly favors R, in equilibrium, both will choose to consult the same type of biased news.

The literature in eco-chambers often explains that people within the same group tend to consume the same type of news because they belong to groups of like-minded people, which influences their learning in a non-rational way (Levy and Razin (2019)). Here, we offer an alternative explanation: even rational, non-partisan individuals may optimally choose to consume the same type of biased news with the sole goal of learning which policy is best.

17

3.2 Learning Choices of a Highly Certain DM

Next, we consider the problem of a DM who is very sure about the state. When the DM's prior is close to full certainty she will weakly prefer own-biased learning, with indifference between no information and any source for extremely certain priors and a strict preference for own-biased information at slightly less certain ones. Interestingly, as the exogenous source becomes more informative, a wider range of these priors strictly prefer to acquire information. The intuition is as follows.

Indifference for extremely certain priors – When the DM's prior is extremely certain, no combination of messages from the chosen and exogenous sources can affect her payoff-relevant action. Her action is fully determined by her prior, making her learning choice irrelevant. As a result, the DM is indifferent between information the two information sources, and also between those and no information at all.

Own-biased learning for very certain priors – Consider a DM so certain about the state that information from the exogenous source alone cannot change her payoff-relevant action. However, she is still uncertain enough that a combination of messages from both the chosen and exogenous sources can alter her decision. In this case, the DM can still obtain valuable information from the chosen source. Since her belief remains own-biased after any message from the exogenous source, and she prefers own-biased learning at both interim posteriors, it is optimal for her to choose own-biased learning from an ex ante perspective.

There are always sufficiently certain priors where the DM is indifferent between obtaining information from any source and no information. This is due to the assumption that the exogenous source is not fully informative about any state. As these regions of indifference are adjacent to regions where the DM strictly prefers own-biased learning over no information, the narrowing of the indifference region means the DM is willing to pay for information over a larger set of certain priors. This narrowing occurs when the exogenous source becomes more informative in the Blackwell sense. As the informativeness of the exogenous source increases, even for more certain priors, combinations of messages from both sources can influence the DM's action. Thus, better exogenous information incentivizes information acquisition. This can be seen in our illustrative example: comparing Figures 4 and 5 the intervals of very certain priors at which Ann is indifferent between sources (in gray) are smaller when she expects to receive additional information from Bob compared to when she expects no information from Bob (learning in isolation).

More specifically, the set of extremely certain priors near $p_0 = 0$ (biased towards state L) shrinks as the informativeness of the exogenous source's right message increases. Similarly, the set of extremely certain priors near $p_0 = 1$ (biased towards state R) shrinks as the informativeness of the exogenous source's left message increases. Hence, a highly certain

DM who is indifferent between sources or no information will pay a positive amount for information if the exogenous source's *opposite-biased message* –one that recommends an action against her bias– is sufficiently informative. This is because only opposite-biased messages can alter the DM's action, and thus only the informativeness of the exogenous source's opposite-biased message influences her willingness to pay for information.

The following proposition formalizes this result.

Proposition 3 *Fix any σ_e and symmetric σ^L and σ^R. There exist cutoffs $0 < \underline{p}_I < \underline{p}_O < \frac{1}{2} < \bar{p}_O < \bar{p}_I < 1$ s.t.*

    i) If $p_0 \in (\underline{p}_I, \underline{p}_O) \cup (\bar{p}_O, \bar{p}_I)$, the DM strictly prefers own-biased learning.

    ii) If $p_0 \in (0, \underline{p}_I) \cup (\bar{p}_I, 1)$, the DM is indifferent between any source and no information.

As σ_e's informativeness increases (in the Blackwell sense), $\underline{p}_I$ decreases and $\bar{p}_I$ increases. Additionally, $\bar{p}_I$ strictly increases with the informativeness of σ_e's left message, while $\underline{p}_I$ strictly decreases with the informativeness of σ_e's right message.

The concepts of bias can also be be interpreted through the informativeness of messages. As Lemma 2 shows, a left-biased source is more informative when sending message r than l, while a right-biased source is more informative when sending l than r. Also note that for any pair of symmetric sources, σ^L and σ^R, the left-biased source has the same informativeness for message r (l) as the right-biased source does for message l (r).

Lemma 2 *A signal σ^x is left-biased (right-biased) i.f.f. $\frac{\pi^x(l|L)}{\pi^x(l|R)} < \frac{\pi^x(r|R)}{\pi^x(r|L)}$ $\left(\frac{\pi^x(l|L)}{\pi^x(l|R)} > \frac{\pi^x(r|R)}{\pi^x(r|L)} \right)$.*

This combined with Proposition 3 implies that for a given pair of symmetric biased exogenous sources, there exists a set of extremely certain priors for which the DM strictly prefers own-biased learning when the exogenous source is own-biased, while being indifferent when the exogenous source is opposite-biased. Our illustrative example reflects this. Looking at Figure 5, when Ann expects Bob to share right-biased news, the interval of very certain left-biased priors where she is indifferent between any source and no information is larger than the analogous interval of right-biased priors. On the contrary, when she expects Bob to share left-biased news (Figure 6), the interval of left-biased priors becomes the smaller one.

This is stated in the following proposition.

Proposition 4 *Fix any symmetric σ^L and σ^R, and a pair of symmetric biased exogenous sources σ^{e_L} (left-biased) and σ^{e_R} (right-biased). Restrict $\sigma_e \in \{\sigma^{e_L}, \sigma^{e_R}\}$. There exist thresholds $0 < \underline{p}_A < \bar{p}_A < \frac{1}{2}$ s.t.*

    i) If $p_0 \in [\underline{p}_A, \bar{p}_A]$, the DM strictly prefers σ^L over no information $\iff \sigma^e = \sigma^{e_L}$.

    ii) If $p_0 \in [1 - \bar{p}_A, 1 - \underline{p}_A]$, the DM strictly prefers σ^R over no information $\iff \sigma^e = \sigma^{e_R}$.

3.2.1 Incentivizing information acquisition: extreme agents

Policymakers may want to incentivize information acquisition by extreme agents. This can be useful to mitigate political polarization (Levy and Razin 2019). Informing skepticals can also play an important role in dealing with public health crisis, such as climate change or COVID-19 (Angelucci et al. 2021, Andre et al. 2024).

Consider an agent deciding whether to get vaccinated, L, or not, R. She can choose whether to consult a pro-vaccine (σ^L) or anti-vaccine (σ^R) source of information for a negligible positive cost c, or, opt for no information at zero cost. Additionally, the agent expects to receive exogenous information beyond her control, σ_e.

A planner can perturb the exogenous information by modifying its accuracy. This can be interpreted as being able to partially affect the information that the agent expects to receive beyond her control, for instance, by organizing information campaigns. The perturbed exogenous information $\tilde{\sigma}_e$ has the structure $\tilde{\pi}^e(l|L) = \pi^e(l|L) + \delta_l$ and $\tilde{\pi}^e(r|R) = \pi^e(r|R) + \delta_r$, where the planner chooses δ_l and δ_r. The planner's goal is to maximize the probability that the agent pays c for information, while minimizing the total perturbation, $\delta_l + \delta_r$. For instance, because collecting and communicating information is costly.

First, knowing the agent's prior, the planner announces the altered information structure $\tilde{\sigma}_e$. The agent then decides whether to acquire her own information or not. If she does, she selects either σ^L or σ^R. After, she observes the messages from $\tilde{\sigma}_e$ and the chosen source (if any), updates her prior and makes her decision.

Suppose the agent is so sure about the best alternative being not getting vaccinated that in absence of the planner's intervention she would choose not to inform herself. In that case, the planner wants to increase the informativeness of σ_e's message contradicting her bias. That is, the planner wants to boost the informativeness of the pro-vaccine message. But, since perturbations are costly, only until the point at which the agent becomes indifferent between paying for anti-vaccine information (the own-biased source) and opting for no information at all. Denote this level of informativeness by I_L^*.[10]

The pro-vaccine message informativeness grows slower with the source's accuracy when it is best to get vaccinated, $\pi^e(l|L)$, than when not, $\pi^e(r|R)$. Therefore, the planner's optimal choice is to set $\delta_r = \frac{\pi^e(l|R)I_L^* - \pi^e(l|L)}{I_L^*}$ while keeping $\delta_l = 0$. When designing informational campaigns in this context, rather than generating the expectation that they are very accurate when the vaccine works, the government might prefer to focus on improving the expected accuracy when it does not. This would more efficiently enhance credibility of pro-vaccine messages and incentivize skeptics to engage in independent information acquisition.

---

[10]As $c \to 0$, this makes the agent's prior exactly equal to $\bar{p}_I$ in Proposition 3.

Similarly, a media outlet that wants to penetrate the market of extreme agents, may want to increase the informativeness available information (increasing both δ_r and δ_l) since its expectation can incentivize those agents to inform themselves.

3.3 Learning Choices of a Moderately Certain DM

For priors that lie between highly certain and relatively uncertain, the DM may use alternative heuristics to choose her source of information. When the exogenous source is sufficiently accurate (i.e., $\pi^e(l|L)$ and $\pi^e(r|R)$ are sufficiently large), there is a region of moderately certain priors on both sides of the neutral prior where the DM strictly prefers opposite-biased learning –selecting the source that biased towards the state she considers less likely. This region of is adjacent to the one of very certain priors where the DM prefers own-biased learning. As the exogenous source becomes more informative, an even larger range of very certain priors prefers opposite-biased learning over own-biased learning. The intuition is as follows.

Opposite-biased learning – A sufficiently accurate exogenous source can lead a moderately certain DM to either become so convinced about her own bias that no further information can affect her action, or to be slightly confident in the state opposing her bias. In the latter case, additional information is valuable because it can improve the accuracy of her action. Consequently, the DM's preference for information at her most uncertain interim posterior drives her optimal choice of source. Since this posterior is biased towards the state opposite to her prior, it is optimal for her to choose opposite-biased learning.

Figure 10 illustrates this: a right-biased DM, expecting right-biased information, prefers to choose a left-biased source. The right-biased DM anticipates that a right-biased source will either make her very confident that the state is R or only moderately confident that it is L. In the first scenario, further learning is not helpful, as she is already so certain that her action will not change. In the second scenario, additional information can be useful and learning from the left-biased source is more valuable than from the right-biased one. Thus, she strictly prefers the left-biased source, since it is the most valuable when she is left more uncertain by the exogenous information.

As the exogenous source becomes less informative, it is less powerful in moving the DM's priors. Therefore, if the informativeness drops enough, its messages alone cannot affect the state that the DM considers more likely. At this point, both interim posteriors become own-biased, and following the discussion in Section 3.2, own-biased learning becomes optimal. Therefore, as the exogenous source becomes less informative, the priors at which the DM follows the opposite-biased learning heuristic must become more uncertain.

21

Figure 10: Example where σ_e makes DM choose opposite-biased learning

Since the DM always remains own-biased after receiving a message that aligns with the state she deems more likely –an *own-biased message*–, whether she finds it optimal to choose the own- or opposite-biased source crucially depends on whether an opposite-biased message of the exogenous source can affect her bias. In particular, as the informativeness of the exogenous source's right message increases, left-biased priors need to be less uncertain to follow the opposite-biased learning heuristic. Similarly, as the exogenous source's left message informativeness increases, right-biased priors also need to be less uncertain to follow the heuristic.

The following proposition formalizes this and provides an interpretable sufficient condition for the opposite-biased learning heuristic to arise.[11]

Proposition 5 *Fix any σ_e and symmetric σ^L and σ^R. If $\frac{\pi^e(r|R)}{\pi^e(r|L)}\frac{\pi^e(l|L)}{\pi^e(l|R)} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$, then, there exist cutoffs $0 < \underline{p}_O < \underline{p}_C < \frac{1}{2} < \bar{p}_C < \bar{p}_O < 1$ s.t.*

i) *If $p_0 \in (\underline{p}_O, \underline{p}_C) \cup (\bar{p}_C, \bar{p}_O)$, the DM strictly prefers opposite-biased learning.*

ii) *If $p_0 \in (0, \underline{p}_O) \cup (\bar{p}_O, 1)$, the DM weakly prefers own-biased learning.*

As σ_e's (Blackwell) informativeness increases, $\underline{p}_O$ decreases and $\bar{p}_O$ increases. Additionally, $\bar{p}_O$ strictly increases with σ_e's left message informativeness and $\underline{p}_O$ strictly decreases with σ_e's right message informativeness.

The sufficient condition in Proposition 5 requires that a measure of joint informativeness of the left and right messages from the exogenous source exceeds a certain threshold.[12] This threshold corresponds to the highest informativeness of messages from the sources in the DM's choice set. Since the informativeness of a message is greater than 1, any exogenous

---

[11] The full characterization offers the exact necessary and sufficient condition.

[12] This joint measure is explained more in detail in Section 4, since it will be relevant for the full characterization.

source with the same structure as σ^L or σ^R would satisfy this condition, as well as any message that Blackwell dominates them. Figure 11 illustrates this sufficient condition.


Figure 11: Illustration of sufficient condition in Proposition 5

Using Proposition 5 we also show that for a given pair of symmetric biased exogenous sources, there exists a set of moderately certain priors where the DM strictly prefers to mismatch the bias of the exogenous source. The reasoning is as follows.

Mismatching the bias of the exogenous source – For moderately certain priors, a DM expects an own-biased exogenous source to either make her very certain about her own bias or only slightly certain that the state is opposite to her bias. Conversely, an opposite-biased exogenous source either makes her even more confident in her bias, or slightly uncertain but still biased towards the same state as her prior. Since her most uncertain interim posterior has different biases in each case, her interim preferences differ at the point where information is most valuable. As a result, from an ex ante perspective, she prefers different sources as well. In particular, when expecting an own-biased exogenous source she strictly prefers to choose an opposite-biased source herself; while when expecting an opposite-biased source she strictly prefers to choose an own-biased source.

This result is summarized in the following proposition.

Proposition 6 *Fix symmetric σ^L and σ^R and another symmetric pair σ^{e_L} (left-biased) and σ^{e_R} (right-biased) s.t. $\forall e \in \{e_L, e_R\}$ $\frac{\pi^e(r|R)}{\pi^e(r|L)}\frac{\pi^e(l|L)}{\pi^e(l|R)} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$. Restrict $\sigma_e \in \{\sigma^{e_L}, \sigma^{e_R}\}$. There exist thresholds $0 < \underline{p}_M < \bar{p}_M < \frac{1}{2}$ s.t. if $p_0 \in [\underline{p}_M, \bar{p}_M] \cup [1 - \bar{p}_M, 1 - \underline{p}_M]$,*

i) *The DM strictly prefers σ^L over σ^R $\iff$ $\sigma^e = \sigma^{e_R}$.*

ii) *The DM strictly prefers σ^R over σ^L $\iff$ $\sigma^e = \sigma^{e_L}$.*

3.3.1 Learning with peers: advocates

Suppose a company decides to explore a disruptive technology, R, which could significantly boost net revenue. The alternative to adopting R is to leave things as they are L. There are two teams in the company (e.g. operations and finance). Both agree that the new technology seems promising compared to the status quo (moderately certain about R), but have the chance to perform independent research before deciding. This research consists of biased tests.[13] Each team is rewarded if it supports the most profitable option.

First, each team chooses which type of test to conduct, either one that tends to favor R, but occasionally gives a clear signal that the new technology is not worth pursuing (σ_R) or one that tends to favor the status quo, but sometimes gives a strong signal that R is worth adopting (σ_L).[14] After conducting their tests, the teams share their test designs and outcomes, and then each team votes on the best course of action. Eventually, the profitability of R becomes known (e.g., through market outcomes of competitors), and the teams receive their payoffs based on whether they backed the right option.

For a range of moderately certain priors about R being a better choice, the two teams coordinate by selecting opposing tests, despite their tests being independent conditional on the best technology. Specifically, one conducts the test that favors the existing technology, while the other conducts the one that favors the new one. Although the two teams share the same objective and prior belief, they behave as *advocates* (in the sense of Dewatripont and Tirole (1999)) for opposite alternatives.

4 Characterization

In the previous section, we introduced the main heuristics that guide the DM's optimal learning choices when additional information is expected. These heuristics are the building blocks for the full characterization of her optimal strategy, which we present in this section. To do so, we first define a measure of a source's aggregate informativeness, which will help in understanding different regions of the characterization.

---

[13]See Gans (2023) for a discussion on firms using biased tests or experiments to evaluate disruptive technologies.

[14]The two tests are a pair of symmetric biased signals. Note that since both teams have access to the same two types of sources, the sufficient condition in Proposition 5 holds.

Definition 4 *The index of aggregate informativeness of a source σ^x is the product of its messages informativeness, namely,* $\frac{\pi^x(l|L)}{\pi^x(l|R)}\frac{\pi^x(r|R)}{\pi^x(r|L)}$.

This index measures the joint informativeness of a source's left and right message, establishing a complete ordering for binary signals (unlike Blackwell's partial ordering). Moreover, Blackwell implies this order. In the sense that if one source Blackwell dominates another, it will also have higher aggregate informativeness. Figure 12 plots isocurves of sources, each corresponding to a given level of aggregate informativeness. Movements towards $(1,1)$ indicate higher values of the index. Increases in aggregate informativeness can result from a higher probability of correct messages (e.g., moving from point C to B) or greater "biasedness," that is, a larger difference in the probability of sending the correct message across states (e.g., moving from B to A). The intuition behind the second is that a drop in probability of a correct message in a given state can be over-compensated by the rise in informativeness of a specific message.


Figure 12: Isocurves for $\frac{\pi^x(l|L)}{\pi^x(l|R)}\frac{\pi^x(r|R)}{\pi^x(r|L)} = 27$; $\frac{\pi^x(l|L)}{\pi^x(l|R)}\frac{\pi^x(r|R)}{\pi^x(r|L)} = 3.86$; $\frac{\pi^x(l|L)}{\pi^x(l|R)}\frac{\pi^x(r|R)}{\pi^x(r|L)} = 1.5$

The next three propositions outline the structure of the DM's optimal strategy when the exogenous source is (i) comparably informative to the sources in the DM's choice set (Proposition 7), (ii) more informative (Proposition 8), or (iii) much less informative (Proposition 9), based on the index of aggregate informativeness. Section 4.4 addresses the special case when the exogenous source is slightly less informative. The specific threshold values that determine the optimal strategy are detailed in the Appendix.

4.1 Comparably informative exogenous source

Proposition 7 describes the structure of the DM's optimal learning strategy when the exogenous source is neither too informative nor too uninformative compared to those in the DM's choice set. Not too informative means that the exogenous source's message alone cannot fully determine the DM's optimal action. The lower bound on the index of aggregate informativeness ensures both that no single source in the DM's set can fully determine the DM's action, and, that the opposite-biased learning heuristic arises for some of the DM's priors.

Proposition 7 *When* $\frac{\pi^L(r|R)}{\pi^L(r|L)} \frac{\pi^R(l|L)}{\pi^R(l|R)} \geq \frac{\pi^e(r|R)}{\pi^e(r|L)} \frac{\pi^e(l|L)}{\pi^e(l|R)} > \max \left\{ \frac{\pi^L(r|R)}{\pi^L(r|L)}, \frac{\pi^R(r|R)}{\pi^R(r|L)} \frac{\pi^L(l|L)}{\pi^L(l|R)} \right\}$,

i) *If* $p_0 \in (0, p_1) \cup (p_5, 1)$, *the DM is indifferent between sources*

ii) *If* $p_0 \in (p_1, p_2) \cup (p_3, p_4)$, *the DM chooses* σ_L

iii) *If* $p_0 \in (p_2, p_3) \cup (p_4, p_5)$, *the DM chooses* σ_R

where $0 < p_1 < p_2 < p_3 < p_4 < p_5 < 1$.

To see how this relates to the heuristics introduced in Section 3, consider an unbiased exogenous source that meets these conditions. When the DM is almost certain of the state, she is *indifferent* between sources, as no message will change her action. For intermediate levels of certainty, where the exogenous source alone cannot determine her action but a combination of messages can, she chooses *own-biased learning*. Finally, if she is uncertain or moderately certain of the state, she chooses *opposite-biased learning*, since the exogenous source's message can alter her bias, leaving her most uncertain when it does.

If the exogenous source is biased, these regions shift, creating asymmetries that lead to patterns like matching or mismatching the bias of the exogenous information, or valuing information only when it is own-biased, as explained along Section 3. The natural case where $\sigma_e \in \{\sigma^L, \sigma^R\}$ always fits within this proposition.[15] Therefore, Figures 5 and 6 follow this structure, showing, for example, that a biased exogenous source may shift p_3 above or below $\frac{1}{2}$, causing an uncertain DM to chooses the source matching σ_e's bias.

4.2 Very informative exogenous source

When the exogenous source is very informative compared to the DM's choice set, meaning its message alone fully determines the DM's action for some priors, Proposition 8 applies.

---

[15]To see this, note that, by Lemma 2, for any pair of symmetric biased σ^L and σ^R, $\frac{\pi^L(r|R)}{\pi^L(r|L)} \frac{\pi^R(l|L)}{\pi^R(l|R)} \geq \frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^L(r|R)}{\pi^L(r|L)} > \max \left\{ \frac{\pi^L(r|R)}{\pi^L(r|L)}, \frac{\pi^R(r|R)}{\pi^R(r|L)} \frac{\pi^L(l|L)}{\pi^L(l|R)} \right\}$.

Proposition 8 *When* $\frac{\pi^e(r|R)}{\pi^e(r|L)} \frac{\pi^e(l|L)}{\pi^e(l|R)} > \frac{\pi^L(r|R)}{\pi^L(r|L)} \frac{\pi^R(l|L)}{\pi^R(l|R)}$,

  i) *If* $p_0 \in (0, p_1) \cup (p_3, p_4) \cup (p_6, 1)$, *the DM is indifferent between sources*

  ii) *If* $p_0 \in (p_1, p_2) \cup (p_4, p_5)$, *the DM chooses* σ_L

  iii) *If* $p_0 \in (p_2, p_3) \cup (p_5, p_6)$, *the DM chooses* σ_R

where $0 < p_1 < p_2 < p_3 < p_4 < p_5 < p_6 < 1$.

This structure is analogous to Proposition 7, with the addition of a central region of indifference. For an unbiased, very informative exogenous source, the DM is *indifferent* between sources both if she is extremely sure of the state or if she is relatively uncertain about it. In the first case she is so sure, that no messages can change her action; while in the latter the exogenous source's messages determine her action fully, regardless of any chosen source's message. In any case, no source in her choice set is valuable.[16] In between the two types of indifference regions, the DM chooses *own-biased learning* for very certain priors and *opposite-biased learning* for moderately certain ones. When the exogenous source is biased, the regions of priors are also shifted making the structure asymmetric.

4.3   Very uninformative exogenous source

When the exogenous source is much less informative than the sources in the DM's choice set, Proposition 9 describes the optimal strategy.

Proposition 9 *When* $\min \left\{ \frac{\pi^L(r|R)}{\pi^L(r|L)}, \frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^R(r|R)}{\pi^R(r|L)} \right\} > \frac{\pi^e(r|R)}{\pi^e(r|L)} \frac{\pi^e(l|L)}{\pi^e(l|R)}$,

  i) *If* $p_0 \in (0, p_1) \cup (p_3, 1)$, *the DM is indifferent between sources*

  ii) *If* $p_0 \in (p_1, p_2)$, *the DM chooses* σ_L

  iii) *If* $p_0 \in (p_2, p_3)$, *the DM chooses* σ_R

where $0 < p_1 < p_2 < p_3 < 1$.

---

[16]Note that while in the first case, no combination of messages from the chosen and exogenous source can affect the DM's action (thus, they do not affect her expected utility), in the second case, they can. However, in that case, the two types of sources (exogenous and chosen) act as substitutes. Thus, while any of the sources in the choice set would be valuable for an uncertain DM who does not expect any additional information; when expecting information from the exogenous source, the sources in the choice set cannot add any value.

In this case the DM's optimal strategy is similar to when there is no additional information, with *indifference* between sources at extreme priors and *own-biased learning* otherwise. However, if the exogenous source is sufficiently informative, the source switch that happens at $\frac{1}{2}$ under own-biased learning shifts following Proposition 2: relatively uncertain agents prefer to match the chosen source with the bias of the exogenous information.

4.4   Relatively uninformative exogenous source

Latly, the optimal strategy in the gap left out by Propositions 7, 8 and 9 depends on whether the difference in informativeness between the two messages from the sources in the DM's choice set is large enough.[17]   If $\frac{\pi^L(r|R)}{\pi^L(r|L)} > \frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^R(r|R)}{\pi^R(r|L)}$, the difference in informativeness between the two messages of the sources is relatively high, which can be interpreted as the sources being *sufficiently biased*. In this scenario, Propositions 7 and 9 extend to cover the gap (see the Appendix for details). Figure 13 illustrates this: the blue region corresponds to exogenous sources leading to the structure described in Proposition 9; the violet region to Proposition 7; and the orange to Proposition 8.


Figure 13: Scope of Propositions 7, 8 and 9 when $\frac{\pi^L(r|R)}{\pi^L(r|L)} > \frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^R(r|R)}{\pi^R(r|L)}$.

On the other hand, when the sources are *sufficiently unbiased* such that $\frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^R(r|R)}{\pi^R(r|L)} >$

[17]The gap is σ_e s.t. $\max\left\{\frac{\pi^L(r|R)}{\pi^L(r|L)}, \frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^R(r|R)}{\pi^R(r|L)}\right\} > \frac{\pi^e(l|L)}{\pi^e(l|R)} \frac{\pi^e(r|R)}{\pi^e(r|L)} > \min\left\{\frac{\pi^L(r|R)}{\pi^L(r|L)}, \frac{\pi^L(l|L)}{\pi^L(l|R)} \frac{\pi^R(r|R)}{\pi^R(r|L)}\right\}$.

$\frac{\pi^L(r|R)}{\pi^L(r|L)}$, a fourth type of optimal strategy may arise. Proposition 10 describes it and the condition under which it emerges.

Proposition 10 *When* $\frac{\pi^L(l|L)}{\pi^L(l|R)}\frac{\pi^R(r|R)}{\pi^R(r|L)} > \frac{\pi^e(r|R)}{\pi^e(r|L)}\frac{\pi^e(l|L)}{\pi^e(l|R)} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$ *and* $\frac{\pi^L(l|L)}{\pi^L(l|R)} > \max\left\{\frac{\pi^e(l|L)}{\pi^e(l|R)}, \frac{\pi^e(r|R)}{\pi^e(r|L)}\right\}$,

i) *If* $p_0 \in (0, p_1) \cup (p_6, 1)$, *the DM is indifferent between sources*

ii) *If* $p_0 \in (p_1, p_2) \cup \left(p_3, \frac{1}{2}\right) \cup [p_4, p_5]$, *the DM chooses* σ_L

iii) *If* $p_0 \in (p_2, p_3) \cup \left(\frac{1}{2}, p_4\right) \cup (p_5, p_6)$, *the DM chooses* σ_R

where $0 < p_1 < p_2 < p_3 < \frac{1}{2} < p_4 < p_5 < p_6 < 1$.

Proposition 10 applies when the exogenous source's aggregate informativeness falls within $\left[\frac{\pi^L(l|L)}{\pi^L(l|R)}\frac{\pi^R(r|R)}{\pi^R(r|L)}, \frac{\pi^L(r|R)}{\pi^L(r|L)}\right]$ and none of its messages are too informative. In this case, the structure is similar to that of Proposition 7, with the addition of a central region of own-biased learning for uncertain priors. This creates four regions of priors on each side of $\frac{1}{2}$.

A DM with a prior slightly below $\frac{1}{2}$ (who believes that state L is slightly more likely) chooses σ_L (*own-biased learning*). As she becomes more confident in state L, she will eventually choose σ_R (*opposite-biased learning*). However, with further certainty, she returns to σ_L (*own-biased learning*). Finally, when she is extremely certain of L, she becomes indifferent between sources. This structure works symmetrically for priors on the right side of $\frac{1}{2}$.

The key feature of the exogenous sources that lead to this type of strategy is that none of their messages are sufficiently informative to significantly influence the choice of source for an uncertain DM, but they are still informative enough to affect the decisions of a moderately certain or highly certain one. In the remaining cases, the optimal learning strategy follows the structure of Proposition 9. This is showcased in Figure 14.
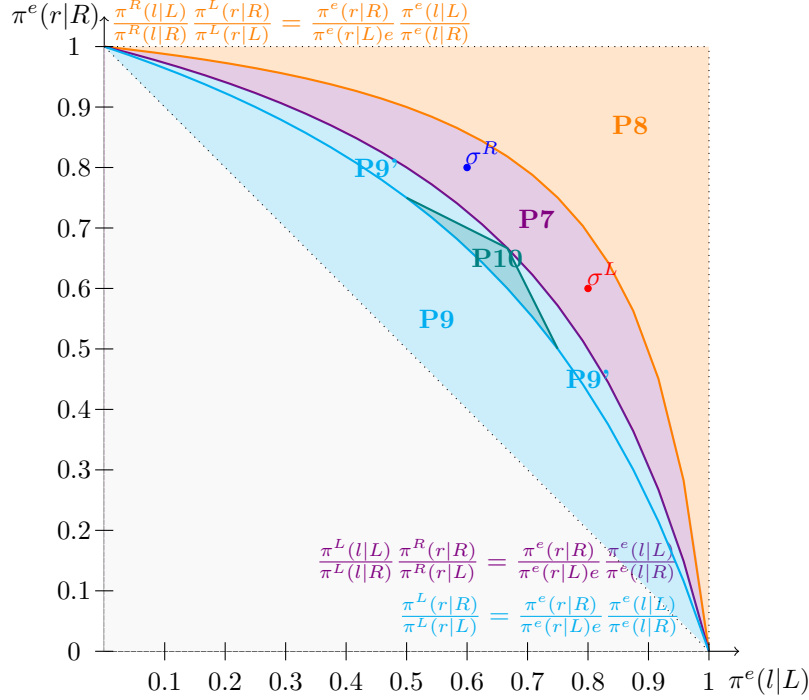
Figure 14: Scope of Propositions 7, 8 and 9 when $\frac{\pi^L(l|L)}{\pi^L(l|R)}\frac{\pi^R(r|R)}{\pi^R(r|L)} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$.

5 Discussion

This paper highlights the importance of expected information in shaping optimal information choices. By incorporating anticipated external information into a standard learning model, we identify a novel mechanism that explains why individuals may consume both like-minded and opposing news sources, aligning with empirical findings in the social media and news consumption literature.

Throughout the paper we also identify rich ways in which the optimal choice of news bias depends on the type of expected information. This leads to interesting predictions for applications to learning with peers and incentivizing information acquisition. Below, we discuss some additional insights about these issues.

5.1 Learning with peers

Consider, as in Sections 3.1.1 and 3.3.1, a scenario where two identical agents first choose one of two biased information sources, $\forall i \in \{1, 2\}, \sigma_i \in \{\sigma^L, \sigma^R\}$ then share their information and update their beliefs before making an individual decision (e.g., voting). Using our characterization of the optimal learning strategy, we determine how each agent responds to

the other's choice and identify the Nash Equilibria (NE) of this information choice game.

Proposition 11 *For any pair of symmetric biased sources, σ^L, σ^R, there exist thresholds $0 < p_1 < p_2 < p_3 < p_4 < p_5 < \frac{1}{2}$ such that:*

    i) If $p_0 \in (0, p_1) \cup (1 - p_1, 1)$, the set of NE are $\{\sigma^L, \sigma^R\}^2$.

    ii) If $p_0 \in (p_1, p_2) \cup (p_5, 1 - p_5) \cup (1 - p_2, 1 - p_1)$, the set of NE are (σ^L, σ^L) and (σ^R, σ^R).

    iii) If $p_0 \in (p_2, p_3) \cup (1 - p_5, 1 - p_4)$, the unique NE is (σ^L, σ^L).

    iv) If $p_0 \in (p_3, p_4) \cup (1 - p_4, 1 - p_3)$, the set of NE are (σ^L, σ^R) and (σ^R, σ^L).

    v) If $p_0 \in (p_4, p_5) \cup (1 - p_3, 1 - p_2)$, the unique NE is (σ^R, σ^R).

This result implies that when agents have low or relatively high certainty ($p_0 \in (p_1, p_2) \cup (p_5, 1 - p_5) \cup (1 - p_2, 1 - p_1)$), they tend to choose the same type of news bias in equilibrium, while moderately certain agents ($p_0 \in (p_3, p_4) \cup (1 - p_4, 1 - p_3)$) may select different biases.

For extremely certain agents ($p_0 \in (0, p_1) \cup (1 - p_1, 1)$), indifference between sources makes any pair of informational choices an equilibrium; however, as soon as there is some cost to acquire information those agents would choose not to inform themselves. A related result is true for agents with relatively high certainty ($p_0 \in (p_1, p_2) \cup (1 - p_2, 1 - p_1)$) who choose the same type of information in equilibrium: as soon as there is some cost to acquire information the two possible equilibrium informational choices are to either both choose the own-biased source, or, both choose no information. At the remaining priors agents have a dominant strategy to either choose σ^L or σ^R and, thus, the equilibrium is unique: both choose the same source but not due to the choice of their peer.

5.2 Incentivizing information acquisition

Now consider a planner who selects the exogenous information σ_e to maximize information acquisition among agents with heterogenous priors. As the receiver in Section 3.2.1, each agent decides whether to pay a small cost to acquire information from one of two biased sources, knowing that the planner will also provide additional information.

Given the distribution of priors, the planner publicly announces the structure of the additional information (e.g., the type of information campaign). Each agent then decides whether to acquire information at a negligible cost ($c \to 0$) and, if so, chooses between σ^L or σ^R. The agent then uses her chosen information (if any) and the planner's additional information to make a final decision (e.g., whether to get vaccinated or what to vote).

Even if additional information is costless for the planner, he faces a trade-off: enhancing the informativeness of σ_e incentivizes highly certain agents to acquire information but may deter uncertain agents from doing so. As the informativeness of σ_e increases (in the Blackwell sense), highly certain agents are more likely to acquire information, while highly uncertain agents may abstain if the additional source becomes too informative (Proposition 8). Once this happens, the range of central priors opting out expands as σ_e's informativeness increases.

From the characterization of the DM's optimal strategy, the planner will always prefer an exogenous source with a minimum level of aggregate informativeness $-\frac{\pi^e(r|R)}{\pi^e(r|L)}\frac{\pi^e(l|L)}{\pi^e(l|R)} \geq \frac{\pi^L(r|R)}{\pi^L(r|L)}\frac{\pi^R(l|L)}{\pi^R(l|R)}$. Below this level, increasing the informativeness of σ_e's messages encourages more agents to acquire information without discouraging others. Beyond this threshold, the planner's optimal choice depends on the distribution of priors: with a high proportion of uncertain agents less informative exogenous information may be preferred, while if many agents are highly certain, greater informativeness may be better, as it increases information acquisition by the extreme agents.

When most agents have the same bias, a similar trade-off arises at the message level. For example, if most are anti-vaccine, increasing the informativeness of σ_e's pro-vaccine message encourages those who are highly skeptical but could discourage those with only mild doubts.

References

Andre, P., Boneva, T., Chopra, F., and Falk, A. (2024). Misperceived social norms and willingness to act against climate change. *Review of Economics and Statistics*, pages 1–46.

Angelucci, C., Khandelwal, A. K., Prat, A., and Swanson, A. (2021). Knowledge acquisition in a high-stakes environment: Evidence from the covid-19 pandemic.

Bakshy, E., Messing, S., and Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132.

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science*, 26(10):1531–1542.

Bergemann, D., Bonatti, A., and Smolin, A. (2018). The design and price of information. *American economic review*, 108(1):1–48.

Bishop, B. (2009). *The big sort: Why the clustering of like-minded America is tearing us apart.* Houghton Mifflin Harcourt.

Blackwell, D. et al. (1951). Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 1, page 26.

Börgers, T., Hernando-Veciana, A., and Krähmer, D. (2013). When are signals complements or substitutes? *Journal of Economic Theory*, 148(1):165–195.

Boxell, L., Gentzkow, M., and Shapiro, J. M. (2017). Greater internet use is not associated with faster growth in political polarization among us demographic groups. *Proceedings of the National Academy of Sciences*, 114(40):10612–10617.

Brooks, B., Frankel, A., and Kamenica, E. (2023). Comparisons of signals. Technical report, Working paper.[2198].

Burke, J. (2008). Primetime spin: Media bias and belief confirming information. *Journal of Economics & Management Strategy*, 17(3):633–665.

Calvert, R. L. (1985). The value of biased information: A rational choice model of political advice. *The Journal of Politics*, 47(2):530–555.

Che, Y.-K. and Mierendorff, K. (2019). Optimal dynamic allocation of attention. *American Economic Review*, 109(8):2993–3029.

Dewatripont, M. and Tirole, J. (1999). Advocates. *Journal of political economy*, 107(1):1–39.

Dubois, E. and Blank, G. (2018). The echo chamber is overstated: the moderating effect of political interest and diverse media. *Information, communication & society*, 21(5):729–745.

Dworczak, P. and Pavan, A. (2022). Preparing for the worst but hoping for the best: Robust (bayesian) persuasion. *Econometrica*, 90(5):2017–2051.

Flaxman, S., Goel, S., and Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly*, 80(S1):298–320.

Gans, J. S. (2023). Experimental choice and disruptive technologies. *Management Science*.

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

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

Gossner, O., Steiner, J., and Stewart, C. (2021). Attention please! *Econometrica*, 89(4):1717–1751.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., and Li, M. (2017). Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964.

Laclau, M., Renou, L., et al. (2017). Public persuasion. *Working Paper*.

Levy, G. and Razin, R. (2019). Echo chambers and their effects on economic and political outcomes. *Annual Review of Economics*, 11(1):303–328.

Liang, A., Mu, X., and Syrgkanis, V. (2022). Dynamically aggregating diverse information. volume 90, pages 47–80.

Masatlioglu, Y., Orhun, A. Y., and Raymond, C. (2023). Intrinsic information preferences and skewness. *American Economic Review*.

Mayskaya, T. (2024). Following beliefs or excluding the worst? the role of unfindable state in learning. *European Economic Review*, 162:104653.

Meyer, M. A. (1991). Learning from coarse information: Biased contests and career profiles. *The Review of Economic Studies*, 58(1):15–41.

Minozzi, W., Song, H., Lazer, D. M., Neblo, M. A., and Ognyanova, K. (2020). The incidental pundit: Who talks politics with whom, and why? *American Journal of Political Science*, 64(1):135–151.

Mullainathan, S. and Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4):1031–1053.

Nikandrova, A. and Pancs, R. (2018). Dynamic project selection. *Theoretical Economics*, 13(1):115–143.

Nyhan, B., Settle, J., Thorson, E., Wojcieszak, M., Barberá, P., Chen, A. Y., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., et al. (2023). Like-minded sources on facebook are prevalent but not polarizing. *Nature*, 620(7972):137–144.

Oliveros, S. and Várdy, F. (2015). Demand for slant: How abstention shapes voters' choice of news media. *The Economic Journal*, 125(587):1327–1368.

Suen, W. (2004). The self-perpetuation of biased beliefs. *The Economic Journal*, 114(495):377–396.

Sunstein, C. R. (2001). *Republic. com*. Princeton university press.

Wojcieszak, M. E. and Mutz, D. C. (2009). Online groups and political discourse: Do online discussion spaces facilitate exposure to political disagreement? *Journal of communication*, 59(1):40–56.

Appendix A.

Proof of Proposition 1.   For any p_0 we construct σ_e and σ'_e that makes σ^x and σ^y optimal respectively.

By the assumed condition, $\exists p$ such that $EU(\sigma^x|p) > EU(\sigma^y|p)$

Case 1:

If $p < p_0$ construct a binary signal, σ_e, with posteriors p and 1 ($\pi(l|L) = 1$, solve for $\pi(r|R)$).

$$\sum_{s\in\{l,r\}} \mathbb{P}(s|p_0)\Big(EU(\sigma^x|p(s^e)) - EU(\sigma^y|p(s^e))\Big) =$$

$$\mathbb{P}(l|p_0)\underbrace{\Big(EU(\sigma^x|p) - EU(\sigma^y|p)\Big)}_{>0} + \mathbb{P}(r|p_0)\underbrace{\Big(EU(\sigma^x|1) - EU(\sigma^y|1)\Big)}_{0} > 0$$

By observation 1, this implies $EU(\sigma^x, \sigma_e|p_0) > EU(\sigma^y, \sigma_e|p_0)$

Case 2:

If $p > p_0$ construct a binary signal, σ_e, with posteriors p and 0 ($\pi(r|R) = 1$, solve for $\pi(l|L)$).

By same logic as in case 1, $EU(\sigma^x, \sigma_e|p_0) > EU(\sigma^y, \sigma_e|p_0)$

Case 3:

If $p = p_0$ construct a binary signal that is just noise (σ^e will not influence beliefs and at current belief σ^x is optimal).

Construction for σ'_e follows same procedure. $\blacksquare$

Note: this proof is generalizable to the finite state, general action setting. The sketch is the following. Construct an exogenous signal with one more signal realization than states. Signal realizations of the exogenous signal are perfectly informative about the state for all but one realization. For that realization, the signal induces the belief for which the "desired" signal is strictly preferred.

Proof of Lemma 1. σ^x Blackwell dominates $\sigma^y \iff \exists \sigma^z$ s.t. the distribution of posteriors from the bundle (σ^y, σ^z) is the same as the distribution of posteriors from σ^x.

Any such σ_z can be written as a binary signal with realizations l^z and r^z and it should break each of the posteriors from σ^y into σ^z's posteriors, that is,

$$p(l^y) = p(r^x)\gamma + p(l^x)(1 - \gamma), \quad \gamma = p(l^y)\pi^z(r|R) + (1 - p(l^y))\pi^z(r|L)$$

$$p(r^y) = p(r^x)\lambda + p(l^x)(1 - \lambda), \quad \lambda = p(r^y)\pi^z(r|R) + (1 - p(r^y))\pi^z(r|L)$$

solving for the structure of σ^z:

$$\pi^z(r|R) = \frac{\pi^x(l|L)(p_0\pi^x(r|R) + (1 - p_0)\pi^x(r|L))}{(\pi^x(l|L) - \pi^x(l|R))p_0}; \quad \pi^z(l|L) = \frac{\pi^x(r|R)(p_0\pi^x(l|R) + (1 - p_0)\pi^x(l|L))}{(\pi^x(l|L) - \pi^x(l|R))(1 - p_0)}.$$

For γ and λ are within the interval $[0, 1]$ i.f.f. $\frac{\pi^y(l|L)}{\pi^y(l|R)} < \frac{\pi^x(l|L)}{\pi^x(l|R)}$ and $\frac{\pi^y(r|R)}{\pi^y(r|L)} < \frac{\pi^x(r|R)}{\pi^x(r|L)}$. $\blacksquare$

This will be used in the next proofs. For any pair of symmetric biased sources σ^L, σ^R:

$$EU(\sigma^L|p) - EU(\sigma^R|p) = \begin{cases} 0 & p < \frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)} \\ \pi^L(r|R)p - \pi^L(r|L)(1-p) & p \in (\frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}, \frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)}) \\ (\pi^L(l|L) - \pi^L(r|R))(1-2p) & p \in (\frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)}, \frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)}) \\ \pi^L(r|L)p - \pi^L(r|R)(1-p) & p \in (\frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)}, \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)})) \\ 0 & \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)} < p \end{cases}$$

Proof of Proposition 2. If p_0 is sufficiently near $\frac{1}{2}$, $p_0 \in (\frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)}, \frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)})$. Thus, if σ_e is so weak that both $p(l^e)$ and $p(r^e)$ remain within $(\frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)}, \frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)})$, which is true for $p_0 = \frac{1}{2}$ i.f.f. $\frac{\pi^L(l|L)}{\pi^L(l|R)} > \max\left\{\frac{\pi^e(l|L)}{\pi^e(l|R)}, \frac{\pi^e(r|R)}{\pi^e(r|L)}\right\}$, then,

$$EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) =$$

$$\mathbb{P}(l^e|p_0)\underbrace{\left(EU(\sigma^L|p(l^e)) - EU(\sigma^R|p(l^e))\right)}_{(\pi^L(l|L)-\pi^L(r|R))(1-2p(l^e))} + \mathbb{P}(r^e|p_0)\underbrace{\left(EU(\sigma^L|p(r^e)) - EU(\sigma^R|p(r^e))\right)}_{(\pi^L(l|L)-\pi^L(r|R))(1-2p(r^e))}.$$

Plugging in for the probabilities and the interim posteriors,

$$EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) = (\pi^L(l|L) - \pi^L(r|R))(1 - 2p_0)$$

which means that own-biased learning is strictly optimal.

If, instead, σ_e is so strong that both $p(l^e) < \frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}$ and $p(r^e) > \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)}$, which is true for $p_0 = \frac{1}{2}$ i.f.f. $\min\left\{\frac{\pi^e(l|L)}{\pi^e(l|R)}, \frac{\pi^e(r|R)}{\pi^e(r|L)}\right\} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$, then, $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) = 0$, meaning that the DM is indifferent between sources.

The remaining cases are the following:

1. $p(l^e) \in (\frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}, \frac{1}{2})$ and $p(r^e) > \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)}$

2. $p(l^e) < \frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}$ and $p(r^e) \in (\frac{1}{2}, \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)})$

3. $p(l^e) \in (\frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)}, \frac{1}{2})$ and $p(r^e) \in (\frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)}, \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)})$

4. $p(l^e) \in (\frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}, \frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)})$ and $p(r^e) \in (\frac{1}{2}, \frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)})$

5. $p(l^e) \in (\frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}, \frac{\pi^R(r|L)}{\pi^R(r|L)+\pi^R(r|R)})$ and $p(r^e) \in (\frac{\pi^L(l|L)}{\pi^L(l|L)+\pi^L(l|R)}, \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)})$

When $p_0 = \frac{1}{2}$, $d(p(l^e), p_0) > d(p(r^e), p_0)$ i.f.f. σ_e is right-biased and $d(p(l^e), p_0) < d(p(r^e), p_0)$ i.f.f. σ_e is left-biased. Therefore, by symmetry, in cases 1 and 3 σ_e is left-biased while in 2 and 4 it is right-biased.

In 1 and 2 it is easy to see that at prior $\frac{1}{2}$ the source that matches σ_e's bias is strictly preferred, since at the most extreme interim posterior the DM is indifferent between sources, e.g., in case 1

$$EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) = \mathbb{P}(l^e|p_0)\underbrace{\left(EU(\sigma^L|p(l^e)) - EU(\sigma^R|p(l^e))\right)}_{>0} > 0.$$

Next, we show that in case 3 σ^L is strictly preferred. Suppose not, then

$$\mathbb{P}(l^e|p_0)\underbrace{\left(EU(\sigma^L|p(l^e)) - EU(\sigma^R|p(l^e))\right)}_{(\pi^L(l|L)-\pi^L(r|R))(1-2p(l^e))} + \mathbb{P}(r^e|p_0)\underbrace{\left(EU(\sigma^L|p(r^e)) - EU(\sigma^R|p(r^e))\right)}_{\pi^L(r|L)p(r^e)-\pi^L(r|R)(1-p(r^e))} \leq 0$$

$$\iff \frac{\pi^e(r|R)}{\pi^e(r|L)} \leq \frac{\pi^L(l|L)}{\pi^L(l|R)}.$$

36

However, this contradicts that $p(r^e) > \frac{\pi^L(l|L)}{\pi^L(l|L) + \pi^L(l|R)}$.

Similarly, in 4, σ^R is strictly preferred. Suppose not, then

$$\mathbb{P}(l^e|p_0) \underbrace{\Big(EU(\sigma^L|p(l^e)) - EU(\sigma^R|p(l^e))\Big)}_{\pi^L(r|R)p(l^e) - \pi^L(r|L)(1-p(l^e))} + \mathbb{P}(r^e|p_0) \underbrace{\Big(EU(\sigma^L|p(r^e)) - EU(\sigma^R|p(r^e))\Big)}_{(\pi^L(l|L) - \pi^L(r|R))(1-2p(r^e))} \geq 0$$

$$\iff \frac{\pi^L(l|L)}{\pi^L(l|R)} \geq \frac{\pi^e(l|L)}{\pi^e(l|R)},$$

which contradicts that $p(l^e) < \frac{\pi^R(r|L)}{\pi^R(r|L) + \pi^R(r|R)}$.

In the remaining case (5),

$$EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) =$$

$$\mathbb{P}(l^e|p_0) \underbrace{\Big(EU(\sigma^L|p(l^e)) - EU(\sigma^R|p(l^e))\Big)}_{\pi^L(r|R)p(l^e) - \pi^L(r|L)(1-p(l^e))} + \mathbb{P}(r^e|p_0) \underbrace{\Big(EU(\sigma^L|p(r^e)) - EU(\sigma^R|p(r^e))\Big)}_{\pi^L(r|L)p(r^e) - \pi^L(r|R))(1-p(r^e))}$$

which is positive i.f.f. $(\pi^e(l|L) - \pi^e(r|R))(\pi^L(r|R) - \pi^L(r|L)) > 0$ and negative i.f.f. $(\pi^e(l|L) - \pi^e(r|R))(\pi^L(r|R) - \pi^L(r|L)) < 0$. Because $\pi^L(r|R) - \pi^L(r|L) > 0$, σ^L is strictly preferred whenever σ_e is left-biased (i.e., $\pi^e(l|L) - \pi^e(r|R) > 0$) and σ^R when σ_e is right-biased.

Since all functions are continuous in p_0, this with Lemma 1 proves the statement. ∎

Proof of Proposition 3. Since $p(l^e) < p(r^e)$,

$$p(l^e), p(r^e) < \frac{\pi^L(r|L)}{\pi^L(r|L) + \pi^L(r|R)} \iff p_0 < \frac{\pi^L(r|L)\pi^e(r|L)}{\pi^L(r|L)\pi^e(r|L) + \pi^L(r|R)\pi^e(r|L)} =: \underline{p}_I;$$

$$p(l^e), p(r^e) > \frac{\pi^R(l|L)}{\pi^R(l|R) + \pi^R(l|L)} \iff p_0 > \frac{\pi^R(l|L)\pi^e(l|L)}{\pi^R(l|L)\pi^e(l|L) + \pi^R(l|R)\pi^e(l|R)} =: \bar{p}_I.$$

In both of these cases, $EU(\sigma^L|p) - EU(\sigma^R|p) = 0$ at both interim posteriors $p \in \{p(l^e), p(r^e)\}$. Thus, $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) = 0$ and the DM is indifferent at those priors.

Next, it is left to show that $\exists \epsilon > 0$ s.t. $\forall p_0 \in (\underline{p}_I, \underline{p}_I + \epsilon)$, $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) > 0$ and $\exists \delta > 0$ s.t. $\forall p_0 \in (\bar{p}_I - \epsilon, \bar{p}_I)$, $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) < 0$.

Let $p_0 = \underline{p}_I + \epsilon$. For any sufficiently small $\epsilon > 0$, $p(r^e) \in \left(\frac{\pi^L(r|L)}{\pi^L(r|L) + \pi^L(r|R)}, \frac{1}{2} \right)$ and $p(l^e) < \frac{\pi^L(r|L)}{\pi^L(r|L) + \pi^L(r|R)}$, which implies that $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) > 0$, since

$$\mathbb{P}(l^e|p_0) \underbrace{\Big(EU(\sigma^L|p(l^e)) - EU(\sigma^R|p(l^e))\Big)}_{0} + \mathbb{P}(r^e|p_0) \underbrace{\Big(EU(\sigma^L|p(r^e)) - EU(\sigma^R|p(r^e))\Big)}_{>0} > 0.$$

Similarly, for $p_0 = \bar{p}_I - \delta$ and any sufficiently small $\delta > 0$, $p(r^e) > \frac{\pi^R(l|L)}{\pi^R(l|R) + \pi^R(l|L)}$ and $p(l^e) \in \left(\frac{1}{2}, \frac{\pi^R(l|L)}{\pi^R(l|R) + \pi^R(l|L)} \right)$, which implies that $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) < 0$.

The comparative statics follow from $\underline{p}_I = \frac{1}{1 + \frac{\pi^L(r|R)\pi^e(r|R)}{\pi^L(r|L)\pi^e(r|L)}}$, $\bar{p}_I = \frac{\frac{\pi^R(l|L)\pi^e(l|L)}{\pi^R(l|R)\pi^e(l|R)}}{\frac{\pi^R(l|L)\pi^e(l|L)}{\pi^R(l|R)\pi^e(l|R)} + 1}$ and Lemma 1. ∎

Proof of Lemma 2. First, we want to show that σ^x being left-biased implies $\frac{\pi^x(l|L)}{\pi^x(l|R)} < \frac{\pi^x(r|R)}{\pi^x(r|L)}$.

Suppose not, then $\pi^x(l|L) - \pi^x(r|R) > 0$ while

$$\frac{\pi^x(l|L)}{\pi^x(l|R)} \geq \frac{\pi^x(r|R)}{\pi^x(r|L)} \iff \pi^x(l|L) - \pi^x(r|R) \geq (\pi^x(l|L) - \pi^x(r|R))(\pi^x(l|L) + \pi^x(r|R)).$$

But this implies that $1 \geq \pi^x(l|L) + \pi^x(r|R)$, which is a contradiction.

Next, we show that $\frac{\pi^x(l|L)}{\pi^x(l|R)} < \frac{\pi^x(r|R)}{\pi^x(r|L)}$ implies that σ^x is left-biased. Suppose not, then

$$\frac{\pi^x(l|L)}{\pi^x(l|R)} < \frac{\pi^x(r|R)}{\pi^x(r|L)} \iff \pi^x(l|L) - \pi^x(r|R) < (\pi^x(l|L) - \pi^x(r|R))(\pi^x(l|L) + \pi^x(r|R)),$$

while $\pi^x(l|L) - \pi^x(r|R) \leq 0$. This implies that either $0 < 0$ or $1 > \pi^x(l|L) + \pi^x(r|R)$, both reaching a contradiction. The proof for σ^x right-biased is symmetric. $\blacksquare$

Proof of Proposition 4. By σ^{eL}, σ^{eR} symmetric, $\frac{\pi^{eL}(r|R)}{\pi^{eL}(r|L)} = \frac{\pi^{eR}(l|L)}{\pi^{eR}(l|R)}$ and $\frac{\pi^{eL}(l|L)}{\pi^{eL}(l|R)} = \frac{\pi^{eR}(r|R)}{\pi^{eR}(r|L)}$. Therefore, using the proof of Proposition 3 and Lemma 2, for a given pair of σ^L, σ^R,

$$0 < \underline{p}_I(\sigma^{eL}) < \underline{p}_I(\sigma^{eR}) < \frac{1}{2} < \bar{p}_I(\sigma^{eL}) < \bar{p}_I(\sigma^{eR}) < 1.$$

Also by the proof of Proposition 3 and continuity, $\exists \epsilon \in (0, \underline{p}_I(\sigma^{eR}) - \underline{p}_I(\sigma^{eL}))$ s.t. $\forall p_0 \in (\underline{p}_I(\sigma^{eL}), \underline{p}_I(\sigma^{eL}) + \epsilon)$, $EU(\sigma^L, \sigma^{eL}|p_0) - EU(\sigma^R, \sigma^{eL}|p_0) > 0$ while $EU(\sigma^L, \sigma^{eR}|p_0) - EU(\sigma^R, \sigma^{eR}|p_0) = 0$. Moreover, $p_0 < \underline{p}_I(\sigma^{eR})$ also implies that $EU(\sigma^L, \sigma^{eR}|p_0) = EU(\sigma^R, \sigma^{eR}|p_0) = 0$.

Similarly, $\exists \delta \in (0, \bar{p}_I(\sigma^{eR}) - \bar{p}_I(\sigma^{eL}))$ s.t. $\forall p_0 \in (\bar{p}_I(\sigma^{eR}) - \delta, \bar{p}_I(\sigma^{eR}))$, $EU(\sigma^L, \sigma^{eR}|p_0) - EU(\sigma^R, \sigma^{eR}|p_0) < 0$ while $EU(\sigma^L, \sigma^{eL}|p_0) = EU(\sigma^R, \sigma^{eL}|p_0) = 0$.

Note that we can set $\underline{p}_A = \underline{p}_I(\sigma^{eL})$ since

$$1 - \underline{p}_I(\sigma^{eL}) = 1 - \frac{1}{1 + \frac{\pi^L(r|R)\pi^{eL}(r|R)}{\pi^L(r|L)\pi^{eL}(r|L)}} = \frac{\frac{\pi^R(l|L)\pi^{eR}(l|L)}{\pi^R(l|R)\pi^{eR}(l|R)}}{\frac{\pi^R(l|L)\pi^{eR}(l|L)}{\pi^R(l|R)\pi^{eR}(l|R)} + 1} = \bar{p}_I(\sigma^{eR}).$$

Finally, setting $\bar{p}_A = \underline{p}_I(\sigma^{eL}) + \gamma$ s.t. $\gamma > 0$ is sufficiently small to satisfy both the requirement for ϵ and δ (by continuity, it always exists) completes the proof. $\blacksquare$

Proof of Proposition 5. First, let $\underline{p}_O$ be the prior at which $p(r^e) = \frac{1}{2}$, that is, $\underline{p}_O = \frac{\pi^e(r|L)}{\pi^e(r|L) + \pi^e(r|R)}$ and $\bar{p}_O$ the prior at which $p(l^e) = \frac{1}{2}$, that is, $\bar{p}_O = \frac{\pi^e(l|L)}{\pi^e(l|L) + \pi^e(l|R)}$. Note that $0 < \underline{p}_O < \frac{1}{2} < \bar{p}_O < 1$. Since $\forall p_0 < \underline{p}_O$, $EU(\sigma^L|p) - EU(\sigma^R|p) \geq 0$ $\forall p \in \{p(l^e), p(r^e)\}$ and $\forall p_0 < \bar{p}_O$, $EU(\sigma^L|p) - EU(\sigma^R|p) \leq 0$ $\forall p \in \{p(l^e), p(r^e)\}$, at those priors, the DM weakly prefers own-biased learning (strictly in the intervals $(\underline{p}_I, \underline{p}_O)$ and $(\bar{p}_O, \bar{p}_I)$).

Next, we show that $\frac{\pi^e(l|L)}{\pi^e(l|R)} \frac{\pi^e(r|R)}{\pi^e(r|L)} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$ implies that for any sufficiently small $\epsilon > 0$ and $\delta > 0$, if $p_0 = \underline{p}_O + \epsilon$, $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) < 0$; and if $p_0 = \bar{p}_O - \delta$, $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) > 0$.

By continuity, since when $p_0 = \underline{p}_O$

$$\frac{\pi^e(l|L)}{\pi^e(l|R)} \frac{\pi^e(r|R)}{\pi^e(r|L)} > \frac{\pi^L(r|R)}{\pi^L(r|L)} \implies p(l^e) < \frac{\pi^L(r|L)}{\pi^L(r|L) + \pi^L(r|R)},$$

for a small enough $\epsilon > 0$, when $p_0 = \underline{p}_O + \epsilon$: (i) $p(r^e) \in (\frac{1}{2}, \frac{\pi^R(l|L)}{\pi^R(l|L) + \pi^R(l|R)})$, which implies that $EU(\sigma^L|p(r^e)) - EU(\sigma^R|p(r^e)) < 0$, and (ii) $p(l^e) < \frac{\pi^L(r|L)}{\pi^L(r|L) + \pi^L(r|R)}$, implying $EU(\sigma^L|p(l^e)) -$

$EU(\sigma^R|p(l^e)) = 0$. All in all, this implies that $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) < 0$.

Similarly, since when $p_0 = \bar{p}_O$

$$\frac{\pi^e(l|L)}{\pi^e(l|R)}\frac{\pi^e(r|R)}{\pi^e(r|L)} > \frac{\pi^L(r|R)}{\pi^L(r|L)} \implies p(r^e) > \frac{\pi^R(l|L)}{\pi^R(l|L) + \pi^R(l|R)},$$

for a sufficiently small $\delta > 0$, when $p_0 = \bar{p}_O - \delta$, $p(l^e) \in (\frac{\pi^L(r|L)}{\pi^L(r|L)+\pi^L(r|R)}, \frac{1}{2})$ and $p(r^e) > \frac{\pi^R(l|L)}{\pi^R(l|L)+\pi^R(l|R)}$, together leading to $EU(\sigma^L, \sigma_e|p_0) - EU(\sigma^R, \sigma_e|p_0) > 0$.

The comparative statics follow from rewritting $\underline{p}_O = \frac{1}{1+\frac{\pi^e(r|R)}{\pi^e(r|L)}}$, $\bar{p}_O = \frac{\frac{\pi^e(l|L)}{\pi^e(l|R)}}{\frac{\pi^e(l|L)}{\pi^e(l|R)}+1}$ and Lemma 1. ∎

Proof of Proposition 6. By $\sigma^{e_L}, \sigma^{e_R}$ symmetric, the proof of Proposition 5 and Lemma 2, for a given pair of σ^L, σ^R, $0 < \underline{p}_O(\sigma^{e_L}) < \underline{p}_O(\sigma^{e_R}) < \frac{1}{2} < \bar{p}_O(\sigma^{e_L}) < \bar{p}_O(\sigma^{e_R}) < 1$.

Also by the proof of Proposition 5 and continuity, when $\frac{\pi^e(l|L)}{\pi^e(l|R)}\frac{\pi^e(r|R)}{\pi^e(r|L)} > \frac{\pi^L(r|R)}{\pi^L(r|L)}$, $\exists \epsilon \in (0, \underline{p}_O(\sigma^{e_R}) - \underline{p}_O(\sigma^{e_L}))$ s.t. $\forall p_0 \in (\underline{p}_O(\sigma^{e_L}), \underline{p}_O(\sigma^{e_L})+\epsilon)$, $EU(\sigma^L, \sigma^{e_L}|p_0) - EU(\sigma^R, \sigma^{e_L}|p_0) < 0$ while $EU(\sigma^L, \sigma^{e_R}|p_0) - EU(\sigma^R, \sigma^{e_R}|p_0) > 0$. Similarly, $\exists \delta \in (0, \bar{p}_O(\sigma^{e_R}) - \bar{p}_O(\sigma^{e_L}))$ s.t. $\forall p_0 \in (\bar{p}_O(\sigma^{e_R}) - \delta, \bar{p}_O(\sigma^{e_R}))$, $EU(\sigma^L, \sigma^{e_R}|p_0) - EU(\sigma^R, \sigma^{e_R}|p_0) > 0$ while $EU(\sigma^L, \sigma^{e_L}|p_0) - EU(\sigma^R, \sigma^{e_L}|p_0) < 0$.

Setting $\underline{p}_M = \underline{p}_O(\sigma^{e_L})$, since $1 - \underline{p}_O(\sigma^{e_L}) = \bar{p}_O(\sigma^{e_R})$, and $\bar{p}_M = \underline{p}_O(\sigma^{e_L}) + \gamma$ s.t. $\gamma > 0$ is sufficiently small to satisfy both the requirement for ϵ and δ completes the proof. ∎

Proofs of Section 4 to be added.