Neutral News as a Bridge to Broader Opinion Exposure

Tereza Burýšková[1], CERGE-EI[2]

Abstract

This paper examines how individuals with limited attention allocate their information consumption across neutral and biased news sources. I develop a model in which agents rationally choose a portfolio of outlets to learn about an uncertain state of the world. While confirmation bias emerges as the benchmark outcome when only biased sources are available, the presence of a sufficiently precise neutral source fundamentally alters optimal information acquisition. Neutral sources do not merely moderate beliefs: they restructure the informational value of biased outlets. In particular, agents with weak prior beliefs may optimally combine neutral sources with sources opposing their prior views. This portfolio exploits an asymmetric complementarity between reliable neutral information and the possibility of immediate uncertainty resolution provided by opposing sources, leading to sharp posterior beliefs without requiring any intrinsic preference for disagreement.

JEL classification: D83, L82

Keywords: neutral media, information acquisition, media balance, confirmation bias

[1]née Hofrichterová, Email: tereza.buryskova@cerge-ei.cz

[2]CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politickych veznu 7, 111 21 Prague, Czech Republic.

1 Introduction

In the contemporary information environment, biased and unbiased news sources coexist and jointly shape how individuals acquire information. A prominent and well-documented pattern in this context is confirmation bias: individuals tend to seek out, prioritize, and recall information that aligns with their pre-existing beliefs (Lord et al., 1979; Rabin and Schrag, 1999; Gentzkow and Shapiro, 2006). As a result, information consumption often reinforces initial opinions and limits exposure to opposing views. At the same time, recent empirical evidence suggests that this pattern is not universal. In some settings, individuals consume a broader portfolio of news sources and engage with content that challenges their prior beliefs (Bursztyn et al., 2022; Levy, 2021). Existing explanations often emphasize emotional or psychological mechanisms, such as the affective response to opposing opinions.

This paper explores whether such behavior can arise purely from rational information acquisition when neutral (yet imperfect) media are available. I place neutral news sources at the center of the analysis. The key result is that neutral media fundamentally reshape the informational role of biased sources. When a sufficiently precise neutral source is available, it can be optimal for agents with weak prior beliefs to combine neutral information with information from sources that oppose their prior views. In this sense, neutral sources act as structural bridges to opposing perspectives: they transform exposure to disagreement from a psychological anomaly into a rational outcome of information design.

I develop a stylized static model in which individuals with limited attention choose a small portfolio of news sources to learn about a binary state of the world. News sources differ only in their error structures. A neutral source provides imperfect but state-independent information. Biased sources, by contrast, may fully reveal one state of the world while potentially misreporting the other. Given an agent's prior belief, one biased source is aligned with the prior, while the other opposes it. Agents commit to their information portfolio before observing any signals. This setting is motivated, for instance, by buying subscriptions or following certain profiles on social media, where there is (at least a mental) switching cost once the decision is made.

The central mechanism of the paper emerges when agents can consult multiple sources. In this case (and provided that the neutral source is sufficiently reliable), some agents optimally combine a neutral source with a source opposing their prior beliefs. This behavior departs from the standard notion of confirmation bias, but it does not rely on any behavioral assumptions. Instead, it is driven by a complementarity between neutral and opposing sources: the opposing source offers a chance of immediate resolution of uncertainty, while the neutral source provides reliable inference when such resolution does not occur.

For illustration, consider an agent deciding whether to support a complex policy such as the

European Green Deal. The actual effects of the policy are uncertain at the time of decision-making, but media outlets offer cues. Suppose the agent believes the policy is likely beneficial but is not fully confident. She can read only a limited number of news articles due to time constraints. Her options include Green Future Magazine (a source biased in favor of the policy), Our Coal Weekly (opposed to it), and the Daily News (neutral but still imperfect).

Reading from Green Future Magazine would likely confirm the agent's views and boost her confidence that the Green Deal is beneficial. If the magazine unexpectedly claimed the policy was harmful, the surprise would be so strong that it could prompt her to fully revise her belief. In both cases, the agent ends up quite confident about her opinion. In contrast, reading Our Coal Weekly and encountering a negative report leaves her uncertain — she cannot discern whether the policy is truly flawed or whether the report simply reflects the source's bias. The neutral Daily News is never entirely persuasive due to occasional random errors, but if those errors are sufficiently rare, it may still be the best choice.

The main result of this paper is that combining Daily News with Our Coal Weekly can be optimal for some agents, as it rarely leaves them in doubt. An agent who believes the Green Deal is slightly more likely to be good than bad assigns a meaningful probability to Our Coal Weekly publishing a surprisingly positive article. If that happens, the agent becomes confident the policy is beneficial regardless of what the Daily News reports. If Our Coal Weekly claims the policy is harmful, the agent can still cross-check the claim with the more reliable and unbiased Daily News. In either case, the agent ends up with a clearer view of the policy's effects.

The key contribution of this paper is to provide a focused mechanism for describing how neutral and biased news sources interact to shape optimal information portfolios. Neutral sources not only become optimal choices for agents with weak priors, but can also induce exposure to opposing sources that agents would otherwise avoid. The analysis clarifies how both confirmation bias and selective exposure to opposing views can arise from rational behavior in a stylized information environment.

The rest of this paper is organized as follows: Section 2 reviews the most relevant literature; Section 3 sets up the theoretical model; Section 4 presents the results and discusses their implications; Section 5 concludes.

2 Related Literature and Contribution

My research contributes to the literature on biased information, which lies in the intersection of psychology and economics. The benchmark result in this field is the confirmation bias (the tendency to read sources aligned with priors). Nickerson (1998) summarizes the earlier literature on this topic. Subsequent contributions from economists have offered various explanations for the phenomenon, including signal misinterpretation (Rabin and Schrag, 1999), reputation concerns (Gentzkow and Shapiro, 2006), preference for consistency (Yariv, 2005), bounded memory (Wilson, 2014), the dimensionality of information (Andreoni and Mylovanov, 2012), and different models of optimal information acquisition (Hu et al., 2024; Jann and Schottmüller, 2023; Allon et al., 2021; Montanari and Nunnari, 2023; Novák et al., 2024)[1]. Among these, the last class of models aligns most closely with the framework I develop in this paper. The primary innovation of my model lies in describing the key role of neutral media as disruptors of confirmation bias.

Several studies already describe behaviors departing from the general preference for aligned information. Charness et al. (2021) derive that this pattern may occur when sources are biased by omission — in case the state of the world is unfavorable for them, they do not send any signal. Jann and Schottmüller (2024) demonstrate the consumption of an opposing source when deriving utility also from learning the opinions of other agents. Che and Mierendorff (2019) show that such behavior can arise once a dynamic trade-off between learning the information early and learning it more precisely is introduced. The basic features of their model are similar to mine; however, the key innovation of my paper is that I model neutral sources and show that they induce demand for opposing sources, even without assuming dynamic trade-offs. Moreover, I study the case of simultaneous information source choice and full commitment, in which agents cannot alter their choices after observing the signals. Therefore, my model provides a suitable description of situations where the agent does not dynamically re-optimize her choices and instead chooses a limited amount of news to consume, sticking to her choice for all future topics she becomes informed about. This could, for instance, describe purchasing subscriptions or choosing which accounts to follow on social media. Once the choice is made, there are some monetary or, at least, mental costs associated with rebalancing the source portfolio, which could keep the agent committed to her initial choice.

This paper could also help to interpret empirical evidence on general news consumption patterns. Ambuehl and Li (2018) find that when facing biased information sources, agents underrespond to a subjective increase in the informativeness of the source and overvalue sources providing certainty. Calford and Chakraborty (2023) run an experiment with costly biased signals and report that, when the costs are equal, agents choose a diverse portfolio of signals. In a setting similar to mine, Montanari and Nunnari (2023) experimentally show that even though it would be optimal for a Bayesian agent to always acquire an aligned signal, agents sometimes tend to choose the opposing signal.

[1] Similar considerations have been discussed also in the context of optimal innovation strategy of a firm (Zhong, 2022; Gans, 2023).

3 Model set-up

Let there be two states of the world $\omega \in \{a, b\}$ ex-ante equally likely to occur. The agent (news consumer) has a prior belief represented by the probability she assigns to the state of the world being a. I denote this prior belief by q. The objective of the agent is to match her action $t \in \{a, b\}$ to the true state of the world, in which case she gets a payoff of one; and a payoff of zero otherwise. That implies a utility function

$$U(t) = \begin{cases} 1, & \text{if } t = \omega \\ 0, & \text{if } t \neq \omega \end{cases} \tag{1}$$

Before taking the action, the agent can consume news articles from three different news sources available, denoted as $M \in (\alpha, \nu, \beta)$. These sources send signals $s \in \{A,B\}$ (representing short news articles). Each source, however, is imprecise and sometimes sends signals that are inconsistent with the actual state of the world. The probability of the mistake depends on the true state of the world and on the source. There is a source that is biased towards a, a neutral source, and a source biased towards b, all of them independent conditional on ω. Table 1 defines the mistake probabilities for different information sources. I assume that the biases of the biased sources are of equal magnitudes, described by $e \in (0,1)$, whereas the bias of the neutral source $e_\nu \in (0,0.5)$ may differ. Furthermore, I assume full commitment: once the agent has made her portfolio selection, she cannot alter it, and all signals are revealed simultaneously.

Table 1: The news sources and their biases

news sources		
α biased towards a	ν neutral	β biased towards b
$P(A\|b) = e$	$P(A\|b) = e_\nu$	$P(A\|b) = 0$
$P(B\|a) = 0$	$P(B\|a) = e_\nu$	$P(B\|a) = e$

In this paper, I use the terms *opposing* and *aligned* sources to capture the relationship between an agent's prior belief q and the biased sources in my model. For $q > 0.5$ (the agent believing that a is more likely), the source biased towards a is the *aligned* source, whereas the source biased towards b is the *opposing* source.

The agent is exogenously limited to n signals she can consume. Her task is to choose an unordered n-tuple W from elements of the set $M \in (\alpha, \nu, \beta)$ to maximize her expected payoff.

The problem of the agent is

$$\max_{(t, W \in \{\alpha,\nu,\beta\}^n)} \mathbf{E}\left[U(t) | q, W\right], \tag{2}$$

where $U(t)$ and q, n, e, e_ν are defined above.

This problem prescribes that once the realization of all signals S is observed, the optimal action of the agent is

(3)
$$t = \begin{cases} a, & \text{if } r_{WqS} > 0.5 \\ b, & \text{if } r_{WqS} < 0.5 \\ a \text{ or } b, & \text{if } r_{WqS} = 0.5 \end{cases},$$

where $r_{WqS} = \mathbf{E}\left[\mathbf{I}_{\omega=a}|q,S,W\right]$, denotes the posterior belief of the agent induced by the signal received.

The expected payoff, as depicted in Figure 1, is given by:

(4)
$$\mathbf{E}\left[U(t)|r_{WqS}\right] = \max\left(r_{WqS}, 1 - r_{WqS}\right)$$

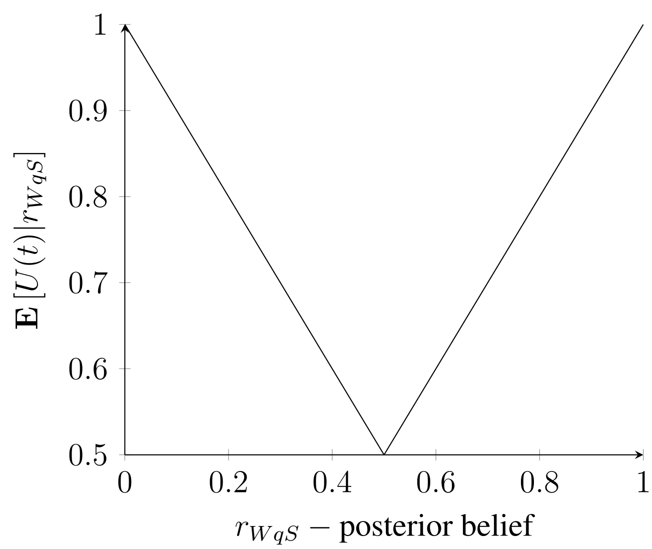

Figure 1: Expected utility as a function of the posterior belief

4 Results

4.1 Demand for One Signal

I begin with a setting where the neutral source is not available, and the agent can consume only a single signal from biased sources $\{\alpha, \beta\}$. This benchmark will serve as a reference point for all subsequent results.

Proposition 1 (Optimal Choice of One Signal, No Neutral Source)**.** *If the neutral source ν is not available, the aligned source always at least weakly dominates the opposing source.*

Proof. See Appendix A.1. □

I explain the intuition for an agent with a belief ($q > 0.5$). This agent expects source α to send signal A, and is quite confident about that because she believes the state of the world is a, which aligns with the fact that this source is biased towards a. In that case, her belief shifts towards one strengthening the prior of the agent. Even if the source α sends signal B, the agent is well off because the source reveals that the state of the world must be b. Therefore, all signal realizations are favorable for the agent.

Conversely, source β does not offer sufficient certainty. If it sends signal b, the agent is left unsure whether it reflects the true state of the world or whether it is just a bias of the source. Of course, there is still a chance for full revelation with this source; However, the agent does not consider this situation very likely, especially if the error e of the source is high. For illustration of the posterior beliefs and expected payoffs, see Figure 2.

A similar result has already been discussed in the literature (Montanari and Nunnari, 2023; Charness et al., 2021). Below, I extend their baseline result by including the neutral source.

Proposition 2 (Optimal Choice of One Signal, Neutral Source Available)**.** *If the neutral source ν is available, the optimal choice depends on the prior q and the relative precision e_ν/e:*

- *If $e_\nu/e > 0.5$, the aligned source is always optimal.*
- *If $e_\nu/e < 0.5$, the neutral source ν is optimal for agents with $q \in \left(\frac{e_\nu}{e}, 1 - \frac{e_\nu}{e}\right)$; otherwise, the aligned source is optimal.*

Proof. See Appendix A.1. □

This proposition suggests a potential way to mitigate confirmation bias — introducing neutral media. If these sources are of sufficient quality, they have the power to push the agent's beliefs towards certainty, thereby avoiding the situation where the agent would be left in doubt. This implies their dominance over the aligned source for agents with weak beliefs.

Figure 2 demonstrates the consideration of the agent. The black line shows the convex mapping between the agent's posterior belief and its payoff. The agent starts with a prior belief associated with the payoff denoted by the black dot. Then, she selects one of the available signals. The signal shifts her belief according to the rules described above. Here, we obtain two possible belief shifts corresponding to two possible signal realizations for each source, which are denoted by red, blue, and green dots. Clearly, receiving signal A from any source shifts the belief to the right, and receiving signal B shifts the belief to the left.

Moreover, given that our agent is Bayesian, there is a simple way to illustrate her expected payoff for each source. The lines connecting the two possible realizations of the signal show all the convex linear combinations of these two payoffs. Furthermore, its intersection with the vertical line drawn at the prior belief of the agent (dashed) determines the overall expected payoff for that source and agent with the given prior belief. This yields from the fact that the expected posterior belief is equal to the prior belief of the agent. Therefore, it is easy to illustrate the intuition behind optimal source choice by comparing the vertical coordinates of intersections with the dashed line denoting the prior belief.

Panel (a) of Figure 2 illustrates the case of an agent with a strong prior belief. For her, the highest payoff is achieved by the aligned source, whereas the other two sources could result in weaker beliefs and lower expected payoffs.

Conversely, an agent with a weak prior belief benefits the most from choosing source ν, as demonstrated in Panel (b). Here, the aligned source is unable to convince the agent that the state of the world is a as effectively as the neutral source is.

I express the optimality conditions as a function of e and e_ν showing that agents never opt for their opposing signal. Instead, depending on the strength of their prior belief and the quality of the sources available, they might seek out a neutral source. The optimal signal choice is completely determined by q (the prior belief) and the ratio of e_ν and e (capturing the relative bias of the neutral source). For $\frac{e_\nu}{e} > 0.5$, it is never optimal to choose the neutral source. However, below this threshold, there is always a band of agents $q \in \left(\frac{e_\nu}{e}, 1 - \frac{e_\nu}{e}\right)$ who find it optimal to choose the neutral source. Figure 3 illustrates the optimal signal choices as a function of e and e_ν.

This result suggests that, absent neutral media and assuming the agent pays only minimal attention, confirmation bias arises even from the simple objective of guessing the true state of the world, given that agents already have preexisting beliefs. However, when high-quality neutral sources are introduced, some agents with weak prior beliefs may prefer them to biased sources. Nevertheless, this benchmark result takes an interesting twist when I allow the agents to choose multiple signals.

[]

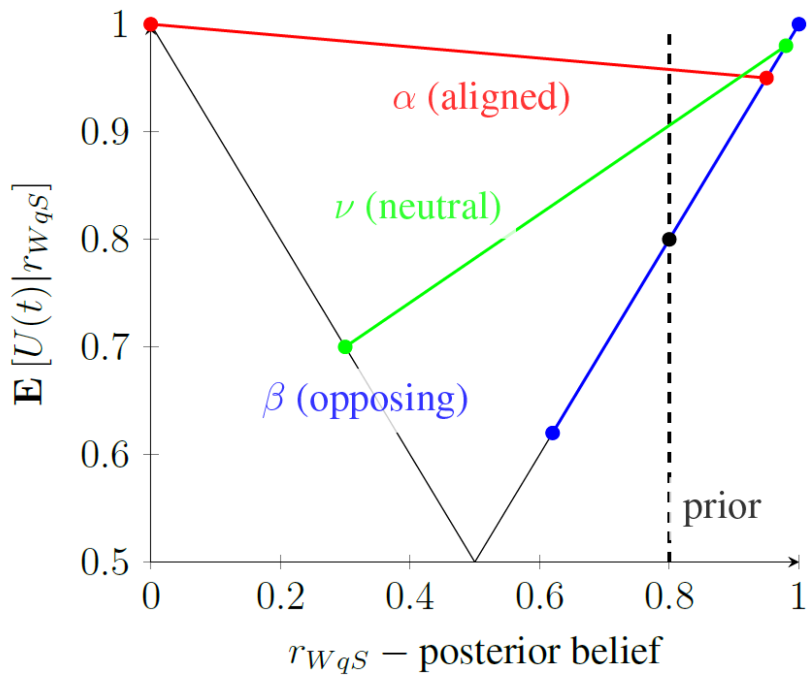

(a) strong prior ($q = 0.8$), optimal choice: α

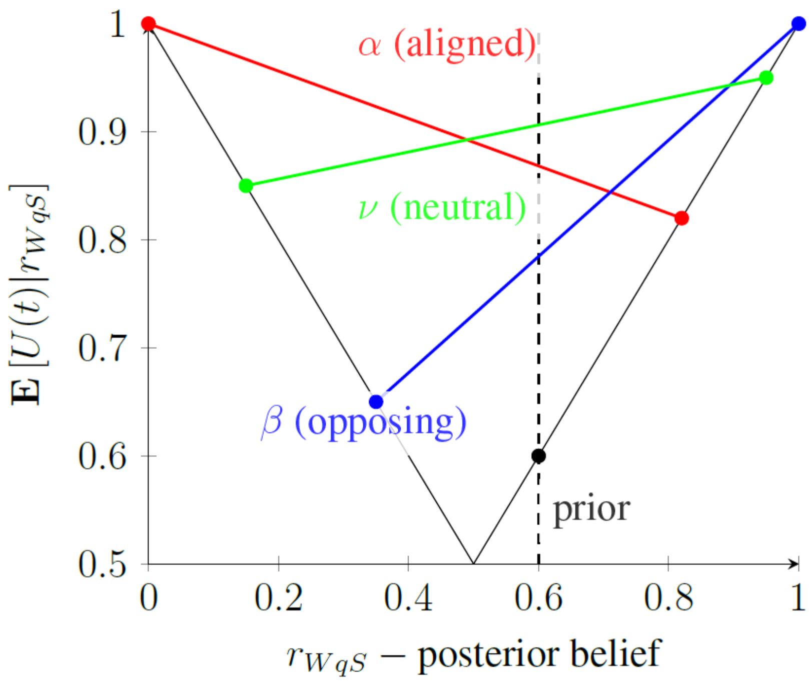

(b) weak prior ($q = 0.6$), optimal choice: ν

Figure 2: Posterior beliefs and expected payoffs for the case of one signal.

Note: For both figures $e_\nu = 0.1$ and $e = 0.4$. The prior of the agent is expressed with a black dot and a dashed black line. The solid black line plots the expected utility as a function of posterior belief. The points demonstrate the potential posterior beliefs and the payoffs associated with them. The intersection of the lines with the dashed line shows the expected payoff for a given source.

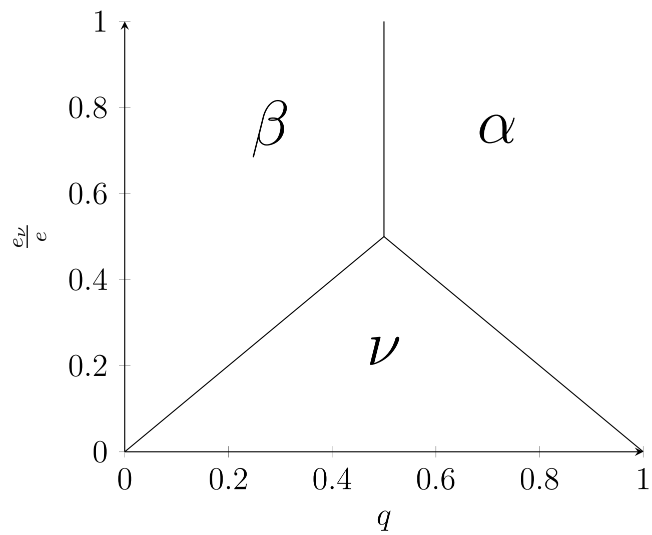

Figure 3: Optimal signal choice for different parameter combinations

4.2 Demand for Two Signals

Now, suppose the agent chooses a portfolio of two signals, possibly from the same source. I assume that the agent can distinguish which source sent which signal. Proposition 3 summarizes the conditions for optimality in this case, focusing only on agents with $q \in (0.5,1)$; the claim holds symmetrically for $q \in (0,0.5)$.

Proposition 3 (Optimal Choice of a Two-Source Portfolio)**.** *Fix signal accuracies e and e_ν. There exist thresholds $q_{\nu\nu} \in (0.5,1)$ and $q_{\beta\nu H} \in (0.5,1)$ such that, for an agent with prior belief $q \in (0.5,1)$, the optimal two-source portfolio satisfies the following properties:*

- *For sufficiently high prior beliefs, $q > \max\{q_{\nu\nu}, q_{\beta\nu H}\}$, the unique optimal portfolio is the aligned portfolio $\alpha\alpha$.*

- *For intermediate prior beliefs, $q \in (0.5, \max\{q_{\nu\nu}, q_{\beta\nu H}\})$, the optimal portfolio is either $\alpha\alpha$, $\nu\nu$, or $\beta\nu$ — double aligned, double neutral, or mix of neutral and opposing signal. Each of these dominates over a nontrivial range of parameter space (q,e,e_ν).*

Proof. See Appendix A.2. □

Proposition 3 shows that allowing the agent to consult two sources substantially enlarges the set of optimal information portfolios. While the intuition behind the double aligned ($\alpha\alpha$) and double neutral ($\nu\nu$) portfolios closely parallels the one-signal case, the new element is the potential optimality of the mixed portfolio $\beta\nu$, which combines an opposing and a neutral source. This portfolio becomes optimal for agents with sufficiently weak prior beliefs when the neutral source is sufficiently precise.

4.2.1 Baseline Observations

Some portfolios are never optimal. The portfolio $\alpha\beta$, which combines aligned and opposing biased sources, rarely improves upon the agent's prior belief and does not reliably induce sharp posterior beliefs. Similarly, the double-opposing portfolio $\beta\beta$ performs poorly for agents with $q > 0.5$, as it offers limited opportunities for belief improvement compared to portfolios involving either the aligned or neutral source. For clarity, these portfolios are omitted from the graphical illustrations.

Figure 4 demonstrates the choice between source portfolios using similar plots as in the case of one signal. Since some of the source portfolios offer three different payoff-relevant signal realizations, it is no longer possible to illustrate the expected payoff as a simple intersection of two lines. However, the logic remains the same: the expected payoff of a given source is the weighted average of the three potential payoffs, as shown at the vertical dashed line representing the prior belief.

Panel (a) demonstrates that for agents with a strong prior belief, no other source can overrule the one-sided aligned portfolio $\alpha\alpha$ as it always induces a sharp posterior belief.

Nevertheless, as the prior belief weakens, portfolios that capture the neutral source become more appealing. For some agents with moderate belief, source $\nu\nu$ might become the most convenient, as shown in panel (b). This choice is always associated with some probability of receiving two different signals and being left with the prior belief, but if it is of sufficient quality, this option happens only with a little probability, and at the same time, the beliefs when the signals are aligned are sufficiently sharp for this portfolio to still be the optimal choice.

Complementarity Between Neutral and Opposing Sources

For agents with sufficiently weak prior beliefs, the mixed portfolio $\beta\nu$ becomes optimal - Figure 4, Panel (c). The dominance of this portfolio follows a simple intuition. Either by sending signal A, source β reveals the true state of the world completely. If not (and the agent gets B from β), she can benefit from the signal acquired from the neutral source.. The posterior belief will, in all cases, be quite close to 0 or 1, making this combination an optimal portfolio choice.

An important implication of this mechanism is that the portfolio $\beta\nu$ can strictly dominate the aligned–neutral portfolio $\alpha\nu$. Under $\alpha\nu$, receiving a favorable signal from the aligned source combined with an unfavorable neutral signal leaves the agent with a posterior close to 0.5, substantially reducing expected payoffs. No analogous low-payoff realization arises under $\beta\nu$, since conflicting signals still shift beliefs in the direction favored by the more precise neutral source.

This asymmetry is driven by the insurance role of the neutral source. When the opposing source β fails to reveal the true state, the neutral source provides a reliable corrective signal that still produces a sharp posterior. By contrast, in the aligned-neutral portfolio $\alpha\nu$, contradictory signals can trap the agent near an uninformative posterior close to 0.5, substantially reducing expected payoff. The neutral source, therefore, complements the opposing source but does not symmetrically complement the aligned source.

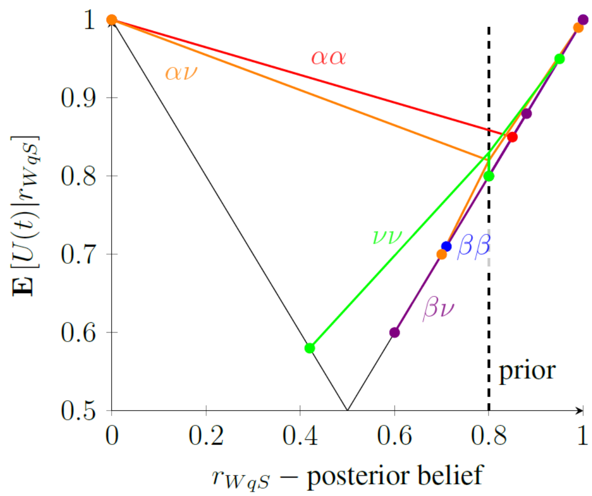

(a) strong prior ($q = 0.8$), optimal choice: $\alpha\alpha$

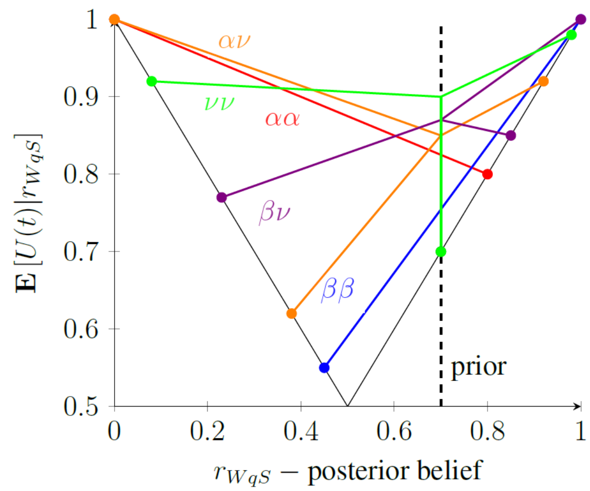

(b) medium prior ($q = 0.7$), optimal choice: $\nu\nu$

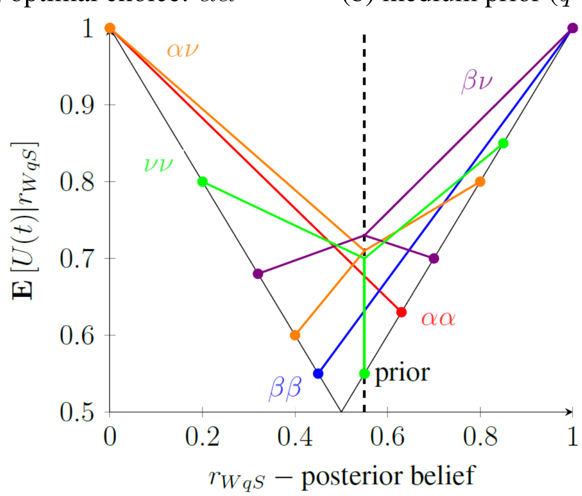

(c) weak prior ($q = 0.55$), optimal choice: $\beta\nu$

Figure 4: Posterior beliefs and expected payoffs for the case of two signals.

Note: For both figures $e_\nu = 0.3$ and $e = 0.8$. The prior of the agent is expressed with a black dot and a dashed black line. The solid black line plots the relation between belief and expected utility. Points show the potential belief shifts achieved by consuming the given source. The intersection of the colorful lines with the dashed line shows the expected payoff for a given source. For clarity, I omit the never optimal signal combination $\alpha\beta$.

Optimality Conditions

In the case of one signal, the optimal choice of sources was completely determined by the combination of q and $\frac{e_\nu}{e}$. For two signals, the optimal signal choice depends on the particular values of all three parameters e, e_ν, and q. I derive the exact optimality conditions in Appendix B. Furthermore, I illustrate this result in Figure 5, which shows the optimal source choices for several parameter configurations.

For low values of e, I obtain the classical confirmation bias result: it is always optimal to choose two signals from the aligned source. Nevertheless, once e reaches a certain threshold, and e_ν remains sufficiently small, for q around 0.5, it is optimal to choose the opposing source in addition to the neutral one. As e further increases, two signals from the neutral source become optimal for medium priors. However, some agents with weak beliefs are always left with a combination of neutral and opposing sources.

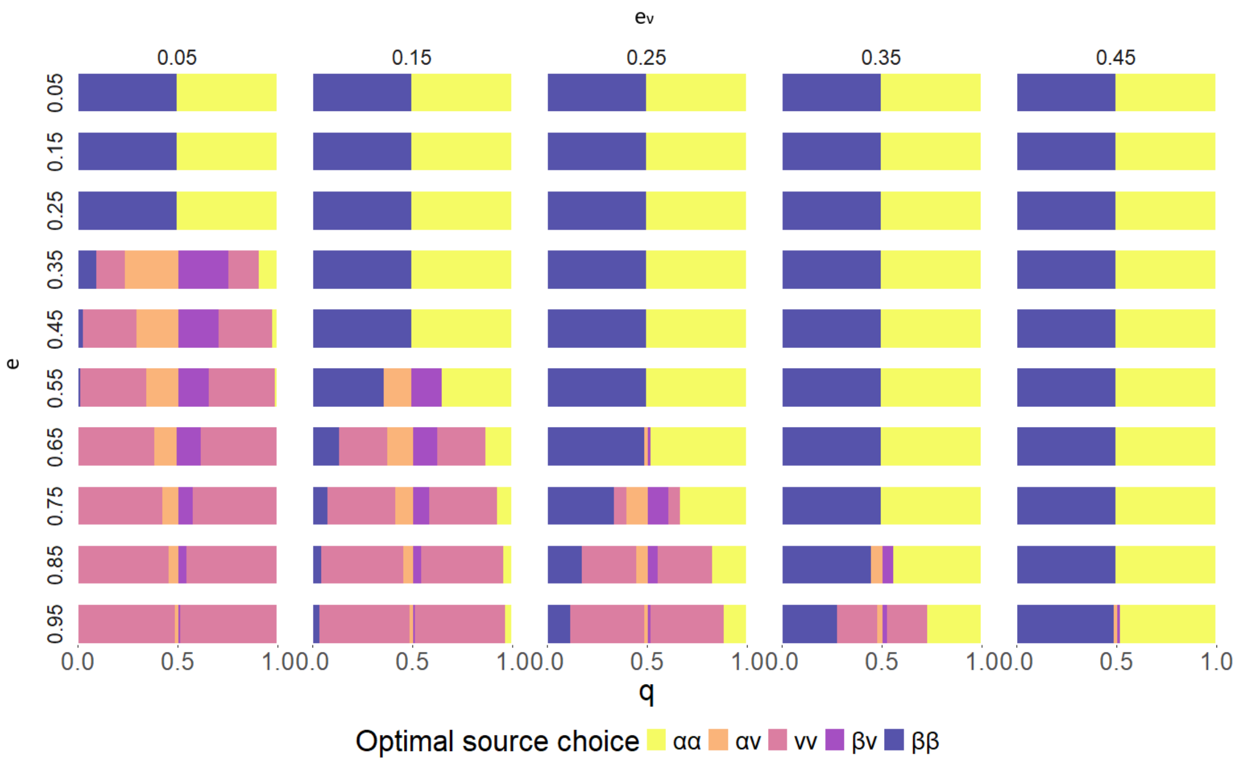

Figure 5: Optimal choice of two sources for different parameter combinations

Note: The rows represent different values of the biased sources error e, the columns represent different values of the neutral source error e_ν. The horizontal axis of each subplot denotes the prior belief of the agent.

The role of neutral sources as informational stabilizers and complements generalizes beyond two-signal portfolios.

4.3 Demand for n Signals

The role of neutral sources as informational stabilizers and complements generalizes beyond two-signal portfolios. Proposition 4 describes the optimal choice for a general number of signals. The problem of the agent would be to choose source portfolio represented by an ordered triplet $(k_\alpha, k_\nu, k_\beta)$, where $k_\alpha \geq 0, k_\nu \geq 0, k_\beta \geq 0$, and $k_\alpha + k_\nu + k_\beta = n$. That brings $\binom{3+n-1}{n}$ options for a given n.

Proposition 4 (Optimal Portfolios with n Signals)**.** *Let the agent be able to acquire $n > 2$ signals.*

1. *In the* ***absence of a neutral source****, the unique optimal portfolio consists of acquiring all n signals from the source aligned with the agent's prior belief.*

2. *When a* ***neutral source is available****, the optimal choice is either*

    - *all-aligned (for small e, high e_ν, strong q),*
    - *all-neutral (for small e_ν, weak q), or*
    - *one opposing, $n-1$ neutral.*

Proof. See Appendix A.3. □

The logic behind Proposition 4 closely parallels the two-signal case. When only biased sources are available, mixing aligned and opposing sources is never optimal: opposing signals either fail to improve beliefs or leave the agent with residual uncertainty, while aligned sources always weakly dominate. The earlier-derived inclination towards aligned sources is therefore a general result for any number of signals.

Introducing a neutral source still preserves the main intuition of the two-signal case. Neutral signals provide reliable but imperfect inference, while opposing signals offer a chance of immediate resolution of uncertainty. For agents with sufficiently weak prior beliefs, the all-neutral portfolio is the most informative, whereas for agents with strong beliefs still opt for the all-aligned portfolio.

Some agents may find one opposing and $n-1$ neutral signals optimal. By giving up one neutral signal, only a little of their informativeness is lost. A single opposing signal can (for some parameter values) offer a fair chance for learning the state completely. Again, neutral and opposing sources complement each other, and for agents with moderate priors, they become the optimal choice.

To illustrate the optimality regions, Figure 6 shows the optimum as a function of (e, e_ν, q) for three and five signals.

(a) Optimal choice of three signals ($n = 3$)

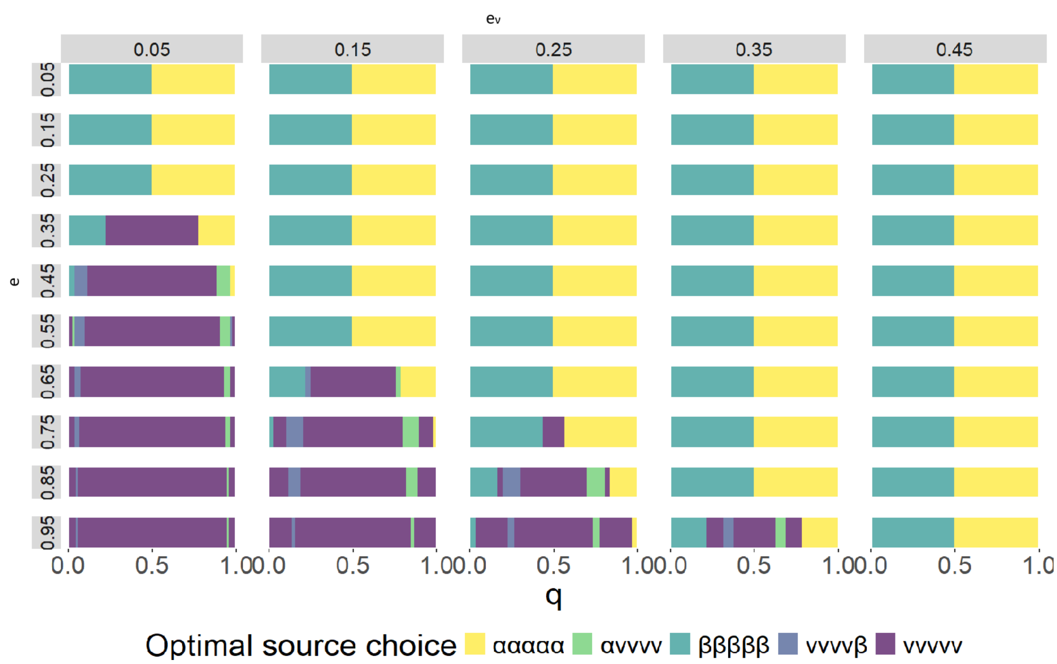

(b) Optimal choice of five signals ($n = 5$)

Figure 6: Optimal source choices for different values of q, e, and e_ν, and n.

5 Conclusion

This paper shows that neutral news sources play a structurally central role in shaping rational information acquisition. While confirmation bias remains the benchmark outcome when only biased sources are available, the introduction of sufficiently precise neutral media fundamentally alters optimal consumption patterns. Neutral sources do not simply moderate beliefs; they restructure the informational environment in which biased outlets operate.

In particular, agents with weak prior beliefs may optimally combine neutral sources with sources opposing their prior views. This portfolio exploits the complementarity between reliable neutral information and the possibility of immediate uncertainty resolution offered by opposing sources, generating sharp posterior beliefs without requiring any intrinsic preference for disagreement. This complementarity is asymmetric: neutral sources increase the informational value of opposing sources more than that of aligned sources, because they insure against the opposing source failing to reveal while preserving the possibility of immediate certainty. Exposure to diverse perspectives emerges endogenously from the structure of information, rather than from behavioral motives. These findings highlight a novel role of neutrality in the current information landscape, suggesting that neutral media can shape not only belief moderation but also the structure of engagement with biased outlets. This perspective may inform debates on media market design and policies aimed at fostering informed and balanced public discourse.

Acknowledgements

I am grateful to my supervisor, Filip Matějka, for invaluable guidance and constant support of this research idea. I would also like to thank Ole Jann, Jan Zápal, Yiman Sun, Konuray Mutluer, Avner Shaked, Gautam Rao, Louis Heizlsperger, and four anonymous referees for their helpful feedback and insightful questions. The project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 870245 (GEOCEP), from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreements No 678081 and No 101002898), and institutional support RVO 67985998 from the Czech Academy of Sciences.

References

Allon, G., K. Drakopoulos, and V. Manshadi (2021). Information inundation on platforms and implications. *Operations Research 69*(6), 1784–1804.

Ambuehl, S. and S. Li (2018). Belief updating and the demand for information. *Games and Economic Behavior 109*, 21–39.

Andreoni, J. and T. Mylovanov (2012). Diverging opinions. *American Economic Journal: Microeconomics 4*(1), 209–32.

Bursztyn, L., A. Rao, C. Roth, and D. Yanagizawa-Drott (2022). Opinions as facts. *The Review of Economic Studies 90*(4), 1832–1864.

Calford, E. M. and A. Chakraborty (2023). The value of and demand for diverse news sources.

Charness, G., R. Oprea, and S. Yuksel (2021). How do people choose between biased information sources? Evidence from a laboratory experiment. *Journal of the European Economic Association 19*(3), 1656–1691.

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

Gans, J. S. (2023). Experimental choice and disruptive technologies. *Management Science 69*(11), 7044–7058.

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

Hu, L., A. Li, and X. Tan (2024). A rational inattention theory of echo chamber.

Jann, O. and C. Schottmüller (2023). Why echo chambers are useful.

Jann, O. and C. Schottmüller (2024). Polarization in news consumption.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review 111*(3), 831–70.

Lord, C., L. Ross, and M. Lepper (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology 37*, 2098–2109.

Montanari, G. and S. Nunnari (2023). Audi alteram partem: An experiment on selective exposure to information.

Nickerson, R. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology 2*, 175–220.

Novák, V., A. Matveenko, and S. Ravaioli (2024). The status quo and belief polarization of inattentive agents: Theory and experiment. *American Economic Journal: Microeconomics 16*(4), 1–39.

Rabin, M. and J. L. Schrag (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics 114*(1), 37–82.

Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica 82*(6), 2257–2294.

Yariv, L. (2005). I'll see it when i believe it: A simple model of cognitive consistency.

Zhong, W. (2022). Optimal dynamic information acquisition. *Econometrica 90*(4), 1537–1582.

A Appendix

A.1 Proof of Propositions 1 and 2

Proof. I prove Propositions 1 and 2 (demand for one signal with/without neutral media) by directly computing the expected utilities for each source. First, I consider source α. An agent with belief q assigns probability $q + e(1-q)$ to this source sending signal A and leading the agent to form a belief of $\frac{q}{q+e(1-q)}$. This posterior belief translates into an expected payoff of $\max\left\{\frac{q}{q+e(1-q)}; 1-\frac{q}{q+e(1-q)}\right\}$. Otherwise, the source sends signal B, and the agent forms a belief of $q=0$ and gets a payoff of 1. Given that, the expected payoff for source α is

$$\text{(A.1)} \qquad \mathbf{E}\left[U|\alpha,q\right] = \left[q+e(1-q)\right]\cdot \max\left[\frac{q}{q+e(1-q)}, 1-\frac{q}{q+e(1-q)}\right] + (1-e)(1-q),$$

which can be simplified to

$$\text{(A.2)} \qquad \mathbf{E}\left[U|\alpha,q\right] = \max\left[q, e(1-q)\right] + (1-e)(1-q).$$

Using similar logic, we get

$$\text{(A.3)} \qquad \mathbf{E}\left[U|\beta,q\right] = \max\left[eq, (1-q)\right] + (1-e)q,$$

and

$$\text{(A.4)} \qquad \mathbf{E}\left[U|\nu,q\right] = \max\left[e_\nu q, (1-e_\nu)(1-q)\right] + \max\left[(1-e_\nu)q, e_\nu(1-q)\right].$$

By comparing these expected utilities, we get that

- source α dominates source β for all $q \in (0.5{,}1)$ (and vice versa for $q \in (0{,}0.5)$)
- for $e_\nu/e < 0.5$ and $q \in \left(\frac{e_\nu}{e}, 1-\frac{e_\nu}{e}\right)$, ν dominates both biased sources.

□

A.2 Proof of Proposition 3

Proof. Without loss of generality, I assume that $q \in (0.5{,}1)$. All the intuition applies symmetrically to $q \in (0{,}0.5)$.

Imposing this restriction on q allows me to simplify the expected utilities (Table A.1). These simplified expressions directly imply that portfolio $\alpha\beta$ is dominated by $\alpha\alpha$ and therefore is never the optimal choice.

Table A.1: Expected Utilities for $q > 0.5$

Portfolio	Expected Utility
$\alpha\alpha$	$1 - e^2(1-q)$
$\alpha\beta$	$1 - e(1-q)$
$\beta\beta$	$\begin{cases} 1 - e^2 q & \text{if } e < \sqrt{\frac{1-q}{q}} \\ q(e^2+1) - 1 & \text{otherwise} \end{cases}$
$\alpha\nu$	$\begin{cases} (1-e_\nu)q + \left[e(1-e_\nu) + (1-e)\right](1-q), & 0.5 \leq q < \frac{e(1-e_\nu)}{e(1-e_\nu)+e_\nu} \\ q + (1-e)(1-q), & q \geq \frac{e(1-e_\nu)}{e(1-e_\nu)+e_\nu} \end{cases}$
$\beta\nu$	$\begin{cases} (1-e)q + (1-q) & \text{if } q < \frac{e(1-e_\nu)}{e(1-e_\nu)+e_\nu} \\ (1-ee_\nu)q + (1-q)(1-e_\nu) & \text{if } \frac{e(1-e_\nu)}{e(1-e_\nu)+e_\nu} \leq q < \frac{ee_\nu}{ee_\nu+1-e_\nu} \\ q & \text{if } q \geq \frac{ee_\nu}{ee_\nu+1-e_\nu} \end{cases}$
$\nu\nu$	$\begin{cases} (1-e_\nu)^2 + 2(1-e_\nu)e_\nu q, & \text{if } 0.5 < q < \frac{(1-e_\nu)^2}{(1-e_\nu)^2+e_\nu^2} \\ q, & \text{if } 1 > q \geq \frac{(1-e_\nu)^2}{(1-e_\nu)^2+e_\nu^2} \end{cases}$

Moreover, $\alpha\alpha$ always dominates $\beta\beta$, because

(A.5) for $e \leq \sqrt{\dfrac{1-q}{q}}$: $\mathbf{E}\left[U|\alpha\alpha, q > 0.5\right] = 1 - e^2(1-q) \geq 1 - e^2 q = \mathbf{E}\left[U|\beta\beta, q > 0.5\right]$

(A.6) for $e > \sqrt{\dfrac{1-q}{q}}$: $\mathbf{E}\left[U|\alpha\alpha, q > 0.5\right] = qe^2 + 1 - e^2 \geq qe^2 + q - 1 = \mathbf{E}\left[U|\beta\beta, q > 0.5\right]$,

because by definition $q - 1 < 0 < 1 - e^2$.

Moreover, $\alpha\nu$ is also always dominated by either $\beta\nu$ or $\alpha\alpha$. I illustrate the intuition for $\beta\nu$ first. Given the symmetry of the problem, it has to be that at $q = 0.5$. Also, one could show that in some neighborhood of $q = 0.5$ it is that $\mathbf{E}\left[U|\alpha\nu\right] = (1 - ee_\nu)q + (1-q)(1-e_\nu)$, and $\mathbf{E}\left[U|\beta\nu\right] = (1 - ee_\nu)q + (1-q)(1-e_\nu)$. From there, it can be shown that the expected utility for $\alpha\nu$ decreases in q, whereas $\beta\nu$ increases in q. Thus, $\beta\nu$ dominates $\alpha\nu$ until these two curves intersect again, which happens only after q reaches $\frac{ee_\nu}{ee_\nu+1-e_\nu}$ and the utility curve for $\alpha\nu$ becomes increasing again. Then, utility curves for $\alpha\nu$ and $\beta\nu$ intersect at

(A.7) $$\hat{q} = \frac{e - e_\nu}{e - e_\nu + ee_\nu}.$$

By that, we have shown that for $0.5 < q < \hat{q}$, $\alpha\nu$ cannot be the optimal choice since it is dominated by $\beta\nu$.

When $1 > q > \hat{q}$, portfolio $\alpha\nu$ is dominated by $\alpha\alpha$. This could be easily shown by comparing the expected utilities

$$\mathbf{E}\left[U|\alpha\nu, q > \hat{q}\right] = q + (1-e)(1-q) > 1 - e^2(1-q) = \mathbf{E}\left[U|\beta\beta, q > \hat{q}\right], \tag{A.8}$$

as $0 < e < 1$.

By this, I have excluded $\alpha\nu$, $\beta\beta$, $\alpha\beta$ from the candidates for the optimal source portfolio for $q > 0.5$. Now, I establish conditions for each of the remaining portfolios, $\alpha\alpha$, $\nu\nu$, and $\beta\nu$, to be optimal. To do that, I use the following notation for the switching points

$$q_{\beta\nu L} = \frac{e(1-e_\nu)}{e(1-e_\nu) + e_\nu}, \tag{A.9}$$

$$q_{\beta\nu H} = \frac{ee_\nu}{ee_\nu + 1 - e_\nu}, \tag{A.10}$$

$$q_{\nu\nu} = \frac{(1-e_\nu)^2}{(1-e_\nu)^2 + e_\nu^2}. \tag{A.11}$$

It can be shown that $q_{\beta\nu L} \leq q_{\beta\nu H}$. However, $q_{\nu\nu}$ could lie above or between these values.

I proceed by showing that if $\beta\nu$ is optimal, it has to be that $q_{\beta\nu L} \leq 0.5 \leq q \leq q_{\beta\nu H}$. I prove this claim by contradiction. Assume that $1 > q > q_{\beta\nu H}$. Then, $\mathbf{E}\left[U|\beta\nu, q\right] = q$, which is always lower than $\mathbf{E}\left[U|\alpha\alpha, q\right] = 1 - e^2(1-q)$. Now, assume that $0.5 < q < q_{\beta\nu L}$. Then, $\mathbf{E}\left[U|\beta\nu, q\right] = (1-e)q + 1 - q$. This is greater than $\mathbf{E}\left[U|\alpha\alpha, q\right] = 1 - e^2(1-q)$ only if $q < e/(1+e)$. However, since $e \in (0,1)$, we get a contradiction with the condition $q > 0.5$. Therefore, $\beta\nu$ can only be optimal if $q_{\beta\nu L} \leq q \leq q_{\beta\nu H}$. Moreover, this implies, that when $q_{\beta\nu L} < 0.5$, $\beta\nu$ is not optimal for $q < q_{\beta\nu L}$. Similarly, it could be shown that in this case $\beta\nu$ is not optimal for $q \in (q_{\beta\nu L}, q_{\beta\nu H})$ either.

In a similar manner, $\nu\nu$ can only be optimal if $q \leq q_{\nu\nu}$, as for $q > q_{\nu\nu}$, it is dominated by $\alpha\alpha$.

Given that, we can finally describe the optimal portfolio choices for both the relative positions of the breaking points and all admissible parameter combinations.

- **Case A** $q_{\nu\nu} < q_{\beta\nu H}$
  - For $q \in (q_{\beta\nu H}, 1)$: $\alpha\alpha$ is optimal
  - For $q \in (q_{\nu\nu}, q_{\beta\nu H})$
    * if $q < \frac{e^2 + e_\nu}{e^2 + e_\nu - ee_\nu}$, $\alpha\alpha$ is optimal,

    * otherwise $\beta\nu$ is optimal
  - For $q \in (0.5, q_{\nu\nu})$
    * if $q < \frac{e^2+e_\nu}{e^2+e_\nu-ee_\nu}$, and at the same time $q < \frac{1-e^2+(1-e_\nu)^2}{2(1-e_\nu)e_\nu-e^2}$ $\alpha\alpha$ is optimal,
    * if $q \geq \frac{1-e^2-(1-e_\nu)^2}{2(1-e_\nu)e_\nu-e^2}$, and at the same time $q \geq \frac{e_\nu-e_\nu^2}{e_\nu-2e_\nu^2-ee_\nu}$, $\nu\nu$ is optimal,
    * otherwise $\beta\nu$ is optimal.

- **Case B** $q_{\nu\nu} > q_{\beta\nu H}$
  - For $q \in (q_{\nu\nu}, 1)$: $\alpha\alpha$ is optimal
  - For $q \in (q_{\beta\nu H}, q_{\nu\nu})$
    * if $q < \frac{1-e^2+(1-e_\nu)^2}{2(1-e_\nu)e_\nu-e^2}$, $\alpha\alpha$ is optimal,
    * otherwise $\beta\nu$ is optimal
  - For $q \in (0.5, q_{\beta\nu H})$
    * if $q < \frac{e^2+e_\nu}{e^2+e_\nu-ee_\nu}$, and at the same time $q < \frac{1-e^2+(1-e_\nu)^2}{2(1-e_\nu)e_\nu-e^2}$ $\alpha\alpha$ is optimal,
    * if $q \geq \frac{1-e^2+(1-e_\nu)^2}{2(1-e_\nu)e_\nu-e^2}$, and at the same time $q \geq \frac{e_\nu-e_\nu^2}{e_\nu-2e_\nu^2-ee_\nu}$, $\nu\nu$ is optimal,
    * otherwise $\beta\nu$ is optimal.

$\square$

A.3 Proof of Proposition 4

Next, I examine the case where the agent can receive $n > 2$ signals. Her problem would be to choose source portfolio represented by an ordered triplet $(k_\alpha, k_\nu, k_\beta)$, where $k_\alpha \geq 0, k_\nu \geq 0, k_\beta \geq 0$, and $k_\alpha + k_\nu + k_\beta = n$. That brings $\binom{3+n-1}{n}$ options for a given n.

No neutral source

I start by showing that when a neutral source is unavailable (or very imprecise), and the agent is allowed to choose only from sources α and β, it is never optimal to choose a combination of these sources. Conversely, the agent always chooses to acquire all signals from her aligned source, which means α for agents with $q > 0.5$ and β for agents with $q < 0.5$

In this simplified problem, the agent chooses k_α, and k_β such that $k_\alpha + k_\beta = n$. Then, she gets k_α signals from source α and k_β signals from source β. Let me denote by $k_{\alpha B}$ the number of signals B sent by source α, and by $k_{\beta A}$ the number of signals A sent by source β. Given the design of our sources, we can use these two numbers to classify all the possible signal realizations the agent might get.

There are, in principle, four different cases. First, the case when $k_{\alpha B} > 0$, and $k_{\beta A} = 0$ completely reveals that the state of the world is B. Symmetrically, the case when $k_{\beta A} > 0$, and $k_{\alpha B} = 0$ completely reveals that the state of the world is A. Given the structure of our sources, we can easily compute the probabilities that the agent assigns to these signal realizations.

The case when both $k_{\beta A}$ and $k_{\alpha B}$ are non-zero has probability zero since no state of the world could support this signal realization.

Only when $k_{\beta A}$ and $k_{\alpha B}$ are equal to zero (which means that both sources sent only the signals in line with their bias), the posterior beliefs of the agents could end up being non-trivial.

Table A.2 presents the probabilities of these signal realizations, the posterior beliefs, and expected payoffs.

Table A.2: Clasification of all possible signal realizations using $k_{\alpha B}$ and $k_{\beta A}$

$k_{\alpha B}$	$k_{\beta A}$	$P[S\|q,W]$	r_{WqS}	$\mathbf{E}[U\|q,W,S]$
>0	$=0$	$(1-e^{k_\alpha})(1-q)$	0	1
$=0$	>0	$(1-e^{k_\beta})q$	1	1
>0	>0	0	—	—
$=0$	$=0$	$e^{k_\beta}q+e^{k_\alpha}(1-q)$	$\frac{e^{k_\beta}q}{e^{k_\beta}q+e^{k_\alpha}(1-q)}$	$\max\left[\frac{e^{k_\beta}q}{e^{k_\beta}q+e^{k_\alpha}(1-q)},\frac{e^{k_\alpha}(1-q)}{e^{k_\beta}q+e^{k_\alpha}(1-q)}\right]$

Given this, and noticing that it has to be that $k_\alpha + k_\beta = n$, the problem of the agent can be rewritten as

(A.12) $$\max_{k_\alpha\in[0,n]}\left\{\max\left[e^{n-k_\alpha}q, e^{k_\alpha}(1-q)\right]+(1-q)(1-e^{k_\alpha})+q(1-e^{n-k_\alpha})\right\}.$$

Here, for a while, I abstract from the fact that k_α should be an integer and handle it as a real number. I first show that there is no interior solution; therefore, for any n, e, and q, it is optimal to choose $k_\alpha = 0$ or $k_\alpha = n$

I proceed by taking the first derivative of the expected payoff. Since the objective function contains a maximum operator, the derivative of the function for $k_\alpha = log_e\left(\sqrt{\frac{qe^n}{1-q}}\right)$ does not exist. Otherwise, it can be expressed as

(A.13) $$\text{if}\quad k_\alpha < log_e\left(\sqrt{\frac{qe^n}{1-q}}\right):\quad -(1-q)e^{k_\alpha},$$

(A.14) $$\text{if}\quad k_\alpha > log_e\left(\sqrt{\frac{qe^n}{1-q}}\right):\quad qe^{n-k_\alpha}.$$

This shows that for $k_\alpha < log_e\left(\sqrt{\frac{qe^n}{1-q}}\right)$, the function is strictly decreasing, reaches a local minimum in $k_\alpha = log_e\left(\sqrt{\frac{qe^n}{1-q}}\right)$ and then increases again. Therefore, no interior k_α maximizes the objective function.

This implies that the only optimal choice could be $k_\alpha = n$, or $k_\alpha = 0$. Let me now derive conditions for the optimality of each choice. The expected utilities equal

(A.15) $$\mathbf{E}\left[U|q,W,(k_\alpha = 0)\right] = \max\left[e^n q,(1-q)\right] + q(1-e^n)$$

(A.16) $$\mathbf{E}\left[U|q,W,(k_\alpha = n)\right] = \max\left[q,e^n(1-q)\right] + (1-q)(1-e^n)$$

The agent always chooses the k_α that delivers higher expected utility. It turns out that

- for $q < 0.5$: $\mathbf{E}\left[U|q,W,(k_\alpha = 0)\right] > \mathbf{E}\left[U|q,W,(k_\alpha = n)\right]$, and the optimal choice is $k_\alpha = 0$,
- for $q > 0.5$: $\mathbf{E}\left[U|q,W,(k_\alpha = n)\right] > \mathbf{E}\left[U|q,W,(k_\alpha = 0)\right]$, and the optimal choice is $k_\alpha = n$.

That means that all agents that initially believed a is more likely to be the state of the world consume only signals from source α, and all agents who initially believed b is more likely to be the true state of the world consume only from source β.

Including the neutral source

Now, I perform a similar analysis for the case, where the neutral source is available. Here, the agent chooses k_α, k_ν, k_β. Again, I will denote by $k_{\alpha B}$ the number of signals B received from source α, and similarly by $k_{\beta A}$ the number of signals A received from source β. These numbers can be used to classify signal realizations into four cases. Table A.3 summarizes them assuming that the agent chose to consume k_ν signals from the neutral source and received $k_{\nu A}$ signals A from it

Table A.3: Classification of all possible signal realizations using $k_{\alpha B}$ and $k_{\beta A}$, and $k_{\nu A}$

$k_{\alpha B}$	$k_{\beta A}$	$P[S\|q,W]$	r_{WqS}
>0	$=0$	$(1-e^{k_\alpha})(1-q)e_\nu^{k_{\nu A}}(1-e_\nu)^{k_\nu - k_{\nu A}}$	0
$=0$	>0	$(1-e^{k_\beta})qe_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}}$	1
>0	>0	—	—
$=0$	$=0$	$e^{k_\beta}qe_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}} + e^{k_\alpha}(1-q)e_\nu^{k_{\nu A}}(1-e_\nu)^{k_\nu - k_{\nu A}}$	$\frac{e^{k_\beta}qe_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}}}{\left(e^{k_\beta}qe_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}} + e^{k_\alpha}(1-q)e_\nu^{k_{\nu A}}(1-e_\nu)^{k_\nu - k_{\nu A}}\right)}$

The expected utility for a particular choice of k_α,k_ν, and k_β can be computed by taking an expectation of the expected utilities over all possible realizations of $k_{\nu A}$ — the combination of signals delivered by the neutral source. That yields

$$\begin{aligned}\mathbf{E}[U|q,W] = \mathbf{E}[\mathbf{E}[U|q,W,S]|S] = \\ \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} (1-e^{k_\alpha})(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}} + \\ \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} (1-e^{k_\beta}) q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}} + \\ \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} \left[e^{k_\beta} q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}} + e^{k_\alpha}(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}} \right] \cdot \\ \max\left(\frac{e^{k_\beta} q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}}}{e^{k_\beta} q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}} + e^{k_\alpha}(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}}}, \right. \\ \left. \frac{e^{k_\alpha}(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}}}{e^{k_\beta} q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}} + e^{k_\alpha}(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}}} \right)\end{aligned} \tag{A.17}$$

The first two sums can be modified using the binomial theorem, yielding

$$\begin{aligned}\mathbf{E}[U|q,W] = (1-e^{k_\alpha})(1-q) + (1-e^{k_\beta})q + \\ + \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} \cdot \max\left[e^{k_\beta} q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}}, e^{k_\alpha}(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}} \right].\end{aligned} \tag{A.18}$$

The general problem of the agent could thus be formulated as

$$\begin{aligned}\max_{k_\alpha, k_\beta, k_\nu} \left\{ (1-e^{k_\alpha})(1-q) + (1-e^{k_\beta})q + \right. \\ \left. + \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} \cdot \max\left[e^{k_\beta} q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}}, e^{k_\alpha}(1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}} \right] \right\},\end{aligned} \tag{A.19}$$

$$k_\alpha + k_\beta + k_\nu = n, \tag{A.20}$$

$$k_\alpha, k_\beta, k_\nu \in \mathbf{N}_0. \tag{A.21}$$

We have already established that it is never optimal to combine aligned and opposing sources in one portfolio. Therefore, the only candidates for an optimal portfolio are

1. $(n-k,k,0), \quad k \in 0,..,n$ aligned-neutral family

2. $(0,n-k,k), \quad k \in 0,..,n$ opposing-neutral family

I proceed by showing that $(n-k,k,0)$ is dominated either by $(n,0,0)$ or $(0,n,0)$ for every $q > 0.5$

Lemma A.1. *Assume* $0 < e < 1$, $0 < e_\nu < \frac{1}{2}$, $\frac{1}{2} < q < 1$, *and* $n > 2$. *Let integers* $k_\nu, k_\alpha > 0$ *satisfy* $k_\nu + k_\alpha = n$. *Then at least one of these inequalities holds*

(A.22)
$$(1-e^n)(1-q) + \max\{q,(1-q)e^n\} \geq$$

(A)
$$(1-e^{k_\alpha})(1-q) + \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} \max\left\{q e_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}},\, e^{k_\alpha}(1-q)e_\nu^{k_{\nu A}}(1-e_\nu)^{1-k_{\nu A}}\right\}.$$

(A.23)
$$\sum_{k=0}^{n} \binom{n}{k} \max\left\{q e_\nu^{n-k}(1-e_\nu)^k,\, (1-q)e_\nu^k(1-e_\nu)^{n-k}\right\} \geq$$

(B)
$$(1-e^{k_\alpha})(1-q) + \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} \max\left\{q e_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}},\, e^{k_\alpha}(1-q)e_\nu^{k_{\nu A}}(1-e_\nu)^{1-k_{\nu A}}\right\}.$$

Proof. Let

$$R = (1-e^{k_\alpha})(1-q) + \sum_{k_{\nu A}=0}^{k_\nu} \binom{k_\nu}{k_{\nu A}} \max\left\{q e_\nu^{k_\nu - k_{\nu A}}(1-e_\nu)^{k_{\nu A}},\, e^{k_\alpha}(1-q)e_\nu^{k_{\nu A}}(1-e_\nu)^{k_\nu-k_{\nu A}}\right\}.$$

Define

$$L_A = (1-e^n)(1-q)+\max\{q,(1-q)e^n\}, \quad L_B = \sum_{k=0}^{n} \binom{n}{k} \max\left\{q e_\nu^{n-k}(1-e_\nu)^k,\, (1-q)e_\nu^k(1-e_\nu)^{n-k}\right\}.$$

Since $e_\nu < \frac{1}{2}$, the likelihood ratio inside the max operator is monotone in k. Hence

$$L_B \geq \max\{q,(1-q)e^n\} = q.$$

Moreover,

$$R \leq (1-q) + \sum_{k=0}^{k_\nu} \binom{k_\nu}{k} \max\left\{q e_\nu^{k_\nu-k}(1-e_\nu)^k,\, (1-q)e_\nu^k(1-e_\nu)^{k_\nu-k}\right\}.$$

If $(1-e^n)(1-q) \geq (1-e^{k_\alpha})(1-q)$, then $L_A \geq R$ and (A) holds. Otherwise, $k_\nu > 0$, and the binomial aggregation in L_B dominates the right-hand side, implying $L_B \geq R$ and (B) holds. □

This shows that portfolios combining neutral and aligned signals are never optimal, thus $(n,0,0)$ and $(0,n,0)$ are the only potential optima in the aligned-neutral family.

As for the opposing-neutral family, I now show that $(0,n-1,1)$ dominates all other portfolios within that family.

Lemma A.2. *Assume $0 < e < 1$, $0 < e_\nu < \frac{1}{2}$, $\frac{1}{2} < q < 1$, and $n > 2$. Let integers $k_\nu, k_\beta > 0$ satisfy $k_\nu + k_\beta = n$. Then*

$$(1-e)q + \sum_{k=0}^{n-1} \binom{n-1}{k} \max\left\{ eq\, e_\nu^{n-1-k}(1-e_\nu)^k, (1-q)e_\nu^k(1-e_\nu)^{n-1-k} \right\}$$
$$\geq (1-e^{k_\beta})q + \sum_{k=0}^{k_\nu} \binom{k_\nu}{k} \max\left\{ e^{k_\beta} q\, e_\nu^{k_\nu-k}(1-e_\nu)^k, e(1-q)e_\nu^k(1-e_\nu)^{k_\nu-k} \right\}.$$

Proof. Consider expressions of the form

$$\max\{Ae_\nu^{m-k}(1-e_\nu)^k,\ Be_\nu^k(1-e_\nu)^{m-k}\}.$$

Define the log–likelihood ratio

$$\Delta(k) = \log\frac{A}{B} + (m-2k)\log\frac{e_\nu}{1-e_\nu}.$$

Since $e_\nu < \frac{1}{2}$, $\Delta(k)$ is strictly decreasing in k, so there exists a unique cutoff k^* such that the first term dominates for $k \leq k^*$ and the second term dominates for $k > k^*$.

On the left-hand side, the cutoff satisfies

$$\log\frac{eq}{1-q} + (n-1-2k)\log\frac{e_\nu}{1-e_\nu} = 0,$$

while on the right-hand side it satisfies

$$\log\frac{e^{k_\beta}q}{e(1-q)} + (k_\nu - 2k)\log\frac{e_\nu}{1-e_\nu} = 0.$$

Because $q > \frac{1}{2}$, $e^{k_\beta} > e$, and $n-1 > k_\nu$, the cutoff on the left-hand side is weakly larger than the cutoff on the right-hand side. Hence, the left-hand side places weakly more probability mass on the larger branch of the max operator.

Let $X_m \sim \text{Bin}(m,1-e_\nu)$. Since $n-1 > k_\nu$ and $e_\nu < \frac{1}{2}$, X_{n-1} first-order stochastically dominates X_{k_ν}. Therefore, for any decreasing function f,

$$\sum_k \binom{n-1}{k} f(k) \;\geq\; \sum_k \binom{k_\nu}{k} f(k).$$

Applying this inequality to the post-max payoff functions yields dominance of the binomial sums.

Finally, since $0 < e < 1$ and $k_\beta > 0$,

$$(1-e)q \;\geq\; (1-e^{k_\beta})q.$$

Combining the dominance of the unconditional term with the dominance of the binomial components establishes the desired inequality. □

By that, I have shown that only all-aligned, all-neutral, and opposing-neutral are the candidates for optimum. I conclude with a discussion of when each of them is optimal

The payoff from the all-aligned portfolio is

$$U_A = (1-e^n)(1-q) + \max\{q,(1-q)e^n\}.$$

Since $q > \frac{1}{2}$ and $e^n < 1$

$$U_A = 1 - (1-q)e^n.$$

As $e \to 0$, $U_A \to 1$. By contrast, if $e_\nu \to \frac{1}{2}$, neutral signals become uninformative and the payoffs from the all-neutral and opposing–neutral portfolios converge to at most $q < 1$. Hence, for sufficiently small e and e_ν close to $\frac{1}{2}$,

$$U_A > \max\{U_N,U_{ON}\}.$$

The payoff from the all-neutral portfolio is

$$U_N = \sum_{k=0}^{n} \binom{n}{k} \max\left\{q e_\nu^{n-k}(1-e_\nu)^k,\, (1-q)e_\nu^k(1-e_\nu)^{n-k}\right\}.$$

As $e_\nu \to 0$, neutral signals become arbitrarily informative, and posterior beliefs converge almost surely to the true state. Consequently,

$$\lim_{e_\nu \to 0} U_N = 1.$$

For any $e \in (0,1)$ and $q < 1$, both the all-aligned and opposing–neutral portfolios yield strictly less than full learning. Hence, for sufficiently small e_ν,

$$U_N > \max\{U_A, U_{ON}\}.$$

The payoff from the opposing–neutral portfolio is

$$U_{ON} = (1-e)q + \sum_{k=0}^{n-1} \binom{n-1}{k} \max\left\{eq\, e_\nu^{n-1-k}(1-e_\nu)^k, (1-q)e_\nu^k(1-e_\nu)^{n-1-k}\right\}.$$

Choose e close to 1, q close to $\frac{1}{2}$, and e_ν in an intermediate range bounded away from both 0 and $\frac{1}{2}$. In this region, aligned signals perform poorly, while neutral-only signals have low persuasive power. By contrast, the opposing–neutral portfolio preserves almost the full informational value of neutral learning plus a chance of immediate resolution Thus, there exists a set of parameters such that

$$U_{ON} > \max\{U_A, U_N\}.$$

Corollary A.2.1. *Each discussed portfolio—all aligned, all neutral, and opposing–neutral—is strictly optimal for some nonempty subset of the parameter space.*