# Neutral News as a Bridge to Broader Opinion Exposure*†

Tereza Burýšková‡, CERGE-EI§

**Abstract**

Why do individuals sometimes seek information from sources that oppose their prior beliefs? Existing explanations typically rely on behavioral motives or dynamic information acquisition. This paper shows that such behavior can arise under fully rational information acquisition when neutral news is available.

I develop a static model in which agents with limited attention commit to a small portfolio of news sources before observing signals. When only biased sources are available, they optimally choose sources aligned with their prior beliefs. Introducing a sufficiently reliable neutral source, however, overturns this conclusion. When prior beliefs are not too strong, the optimal portfolio instead combines the neutral source with one opposing the agent's prior beliefs. The two sources are complementary: the opposing source can resolve uncertainty completely, while the neutral source provides reliable information otherwise.

The analysis identifies a novel mechanism through which neutral media reshape the informational value of biased sources. More broadly, the results suggest that neutral media can foster exposure to opposing viewpoints not by directly moderating beliefs, but by creating complementarities across news sources.

JEL classification: D83, L82

Keywords: neutral media, media bias, information acquisition, selective exposure

*I am grateful to my supervisor, Filip Matějka, for invaluable guidance and constant support of this research idea. I would also like to thank Gautam Rao, Margaret Meyer, Jack Stecher, Ole Jann, Jan Zápal, Yiman Sun, Konuray Mutluer, Avner Shaked, Louis Heizlsperger, and six 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.

†During the preparation of this manuscript, I used Grammarly, Microsoft Copilot, ChatGPT, and Google Gemini to assist with language editing, formatting, consistency checks, and improvements to the clarity and presentation of the manuscript. I have critically reviewed and, where appropriate, revised all AI-generated content, independently verified all mathematical results and proofs, and take full responsibility for the content of the published article.

‡Email: tereza.buryskova@cerge-ei.cz

§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 neutral 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. 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.

The mechanism underlying this result is conceptually different from mechanisms emphasized in the existing literature. Most closely related is Che and Mierendorff (2019), who show that opposing information can be valuable because agents acquire information sequentially and can condition future learning on earlier signal realizations. By contrast, I study an environment in which agents commit to their information portfolio before observing any signals and cannot subsequently revise their choices. By shutting down the possibility of contingent continuation, the model isolates a different force: neutral sources can make opposing sources valuable through their informational complementarity.

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.

Importantly, the contribution is not simply that opposing information can be valuable. Under the full-commitment benchmark studied here, opposing sources are never optimal when only biased sources are available. They become optimal only when combined with sufficiently reliable neutral information. Thus, neutral media do not merely provide another source of information; they alter the complementarity structure of discrete information portfolios and thereby create demand for opposing viewpoints even in the absence of dynamic stopping or contingent information acquisition.

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 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 implica-

tions; Section 5 explores the robustness of the results, and Section 6 concludes.

# 2 Related Literature and Contribution

My research contributes to the literature on endogenous information acquisition, often applied to the context of media. The benchmark result in this field is confirmation bias, namely, the tendency to consume information that aligns with prior beliefs. Nickerson (1998) summarizes the earlier psychology literature on this topic. Subsequent contributions 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, 2026; Allon et al., 2021; Montanari and Nunnari, 2025; Novák et al., 2024).[1] Among these, the last class of models aligns most closely with the framework I develop in this paper.

Several studies describe behaviors departing from the general preference for aligned information. Charness et al. (2021) show that such patterns may arise when sources are biased by omission and strategically withhold information unfavorable to their interests. Most closely related to the present paper, Che and Mierendorff (2019) show that opposite-biased learning can arise in a dynamic information-acquisition problem, where agents face a trade-off between obtaining information early and obtaining it more precisely.

The starting point of my paper is similar to that of Che and Mierendorff (2019), but the mechanism generating demand for opposing information differs. In their setting, agents acquire information sequentially and can condition future information acquisition on earlier signal realizations. An opposing signal can therefore be valuable because it creates an option value: after favorable realizations, the agent can stop learning, whereas after unfavorable realizations she continues acquiring information. By contrast, I study simultaneous information-source choice under full commitment. Agents choose their entire information portfolio before observing any signals and cannot subsequently alter their choices. This distinction is important because it isolates a mechanism for exposure to opposing viewpoints that does not rely on dynamic reoptimization or contingent continuation decisions. The contribution of the paper is therefore not simply that opposing information may be valuable, but that neutral information can make opposing information valuable even under full commitment. In this sense, neutral media act as a bridge to broader opinion exposure by creating an asymmetric complementarity with opposing sources.

More broadly, the paper relates to the literature on optimal choice of a costly signal. A com-

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

mon result in models with flexible signal design is that optimal information structures subject to feasibility constraints tend to align with prior beliefs (Matějka and McKay, 2015; Caplin et al., 2022). In contrast, I study an environment in which agents choose from a finite set of discrete media sources. The resulting information portfolios exhibit complementarities across sources, and these complementarities can make a portfolio containing an opposing source optimal even though acquiring the opposing source in isolation is not.

This paper could also help interpret empirical evidence on news consumption patterns and inform the policy debate. Ambuehl and Li (2018) find that when facing biased information sources, agents underrespond to a subjective increase in source informativeness and overvalue sources providing certainty. Calford and Chakraborty (2025) report experimental evidence that, when costs are equal, agents choose diverse portfolios of information sources. In a setting similar to mine, Montanari and Nunnari (2025) show experimentally that even though a Bayesian agent would optimally choose aligned information, individuals sometimes seek out opposing information. Outside the lab, Bursztyn et al. (2022) has documented the demand for pure news and opinion programs in TV broadcasts. Levy (2021) runs a field experiment on social media, finding that users might be satisfied with a combination of opinions. The mechanism proposed in this paper provides one rational explanation for why such behavior may arise.

# 3 Model set-up

Let there be two states of the world $\omega \in \{a, b\}$. The agent (news consumer) has a prior belief represented by the probability of $\omega = a$. 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 $\omega$. 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 statically. The full-commitment assumption is central to the contribution of the paper. Unlike dynamic information-acquisition models, agents cannot condition future information choices on earlier signal realizations and cannot stop learning once uncertainty is resolved. Consequently, any demand for opposing information that emerges in the model cannot be attributed to contingent continuation or information-cost savings. The analysis therefore isolates the informational role of neutral sources in generating exposure to opposing viewpoints.

Table 1: The news sources and their biases

| news sources | | |
|---|---|---|
| $\alpha$ | $\nu$ | $\beta$ |
| biased towards $a$ | neutral | 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 $q, n, e, e_\nu$ are defined above.

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

$$t = \begin{cases} a, & \text{if } r(S|q,W) > 0.5 \\ b, & \text{if } r(S|q,W) < 0.5 \\ a \text{ or } b, & \text{if } r(S|q,W) = 0.5 \end{cases}, \tag{3}$$

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

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

$$\mathbf{E}\left[U(t)|r(S|q,W)\right] = \max\left(r(S|q,W), 1 - r(S|q,W)\right) \tag{4}$$

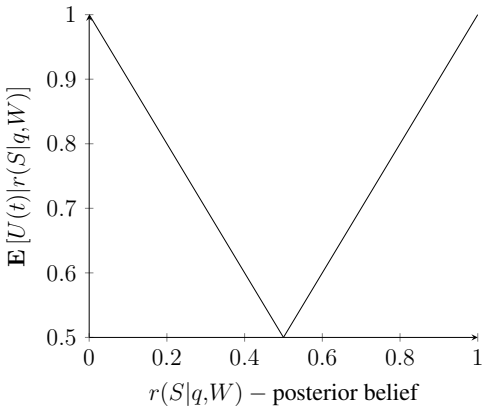


Figure 1: Expected utility as a function of the posterior belief

# 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 $\nu$ 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 $\alpha$ 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 strengthens towards one. Even if the source $\alpha$ 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 $\beta$ does not guarantee 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.

A similar result has already been discussed in the literature (Montanari and Nunnari, 2025; Charness et al., 2021; Che and Mierendorff, 2019). 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 $\nu$ is available, the optimal choice depends on the prior $q$ and the relative precision $e_\nu/e$:*

- *If $e_\nu/e > 0.5$, the aligned source is always optimal.*
- *If $e_\nu/e < 0.5$, the neutral source $\nu$ 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 presents a natural role of neutral media — they become optimal for agents with weak prior beliefs as they offer a strong and balanced belief shift.

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. 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. [2] 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 $\nu$, as demonstrated in Panel (b). Here, the aligned source cannot convince the agent that the state of the world is $a$ as effectively as the neutral source can.

I express the optimality conditions as functions of $e$ and $e_\nu$, showing that agents never opt for the opposing signal. Instead, depending on the strength of their prior belief and the quality of the

[2] This follows from the fact that the expected posterior belief equals the agent's prior belief.

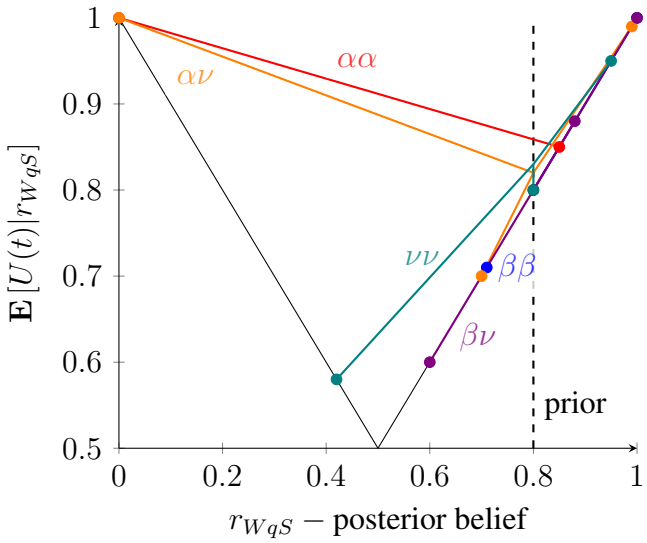


(a) strong prior ($q = 0.8$), optimal choice: $\alpha\alpha$

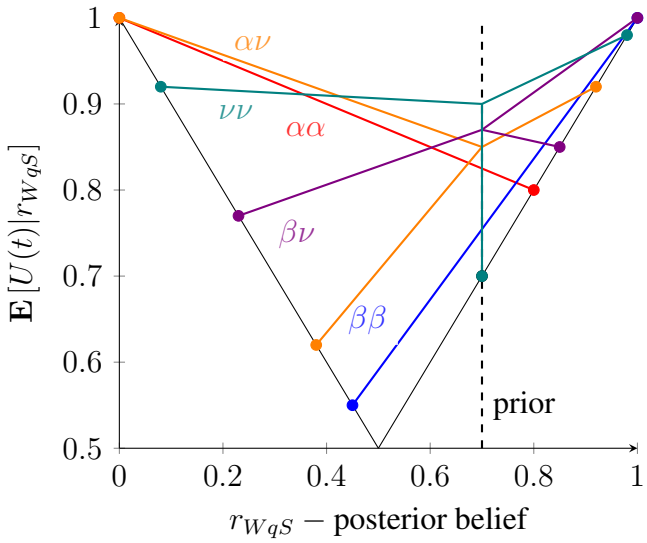


(b) medium prior ($q = 0.7$), optimal choice: $\nu\nu$

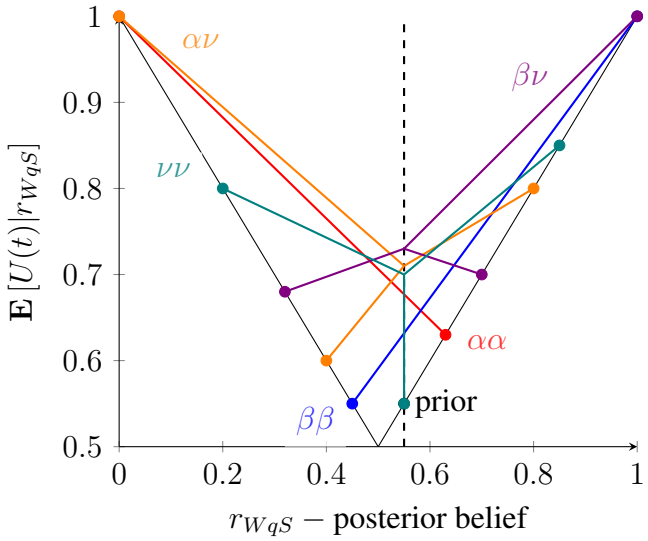


(c) weak prior ($q = 0.55$), optimal choice: $\beta\nu$

Figure 2: Posterior beliefs and expected payoffs for the case of two signals.

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.

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 $e_\nu$ 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 $e_\nu$.

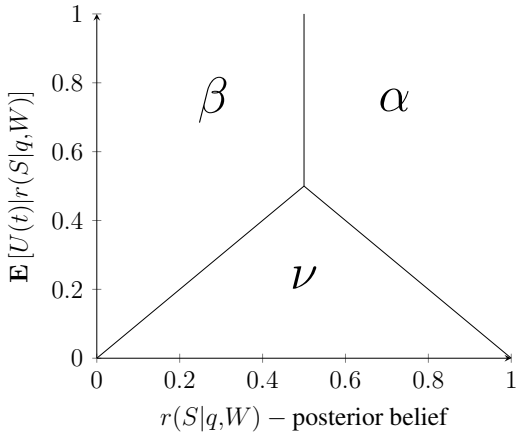


Figure 3: Optimal signal choice for different parameter combinations

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. 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.

## 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)**.** *There exists a threshold $\bar{q} \in (0.5,1)$ such that, for $q \in (0.5,1)$, the optimal choice of a two-signal portfolio ($n = 2$) satisfies the following:*

1. *If the neutral source is unavailable, the unique optimal portfolio is the aligned portfolio $\alpha\alpha$.*

2. *If the neutral source is available, then:*

- *for $q > \bar{q}$, the unique optimal portfolio is $\alpha\alpha$;*
- *for $q < \bar{q}$, the optimal portfolio is one of*

$$\alpha\alpha, \qquad \nu\nu, \qquad \beta\nu.$$

3. *Moreover, if the neutral source is available, and*

$$e_\nu < \frac{e^2}{1+e},$$

*there exists $\varepsilon > 0$ such that $\beta\nu$ is the unique optimal portfolio for all*

$$q \in (0.5, 0.5 + \varepsilon).$$

*Proof.* See Appendix A.2. □

Proposition 3 identifies the central mechanism of the paper. The opposing source does not become valuable by itself; rather, it becomes valuable because it complements neutral information. The optimal portfolio therefore contains an opposing source only together with a neutral source.

**Baseline Intuition**

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 $q > 0.5$. An important implication is that neutral sources are not merely one of several ways of generating demand for opposing information. In fact, every optimal portfolio containing an opposing source also contains a neutral source. Portfolios that combine aligned and opposing sources alone are never optimal.

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 $\beta$ reveals the true state of the world completely. If not (and the agent gets $B$ from $\beta$), 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, for certain error-rate configurations, the portfolio $\beta\nu$ can be more appealing than 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 $\beta$ 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.

This result implies that the availability of a sufficiently precise neutral source is a necessary condition for opposing information to be included in an optimal portfolio under full commitment.

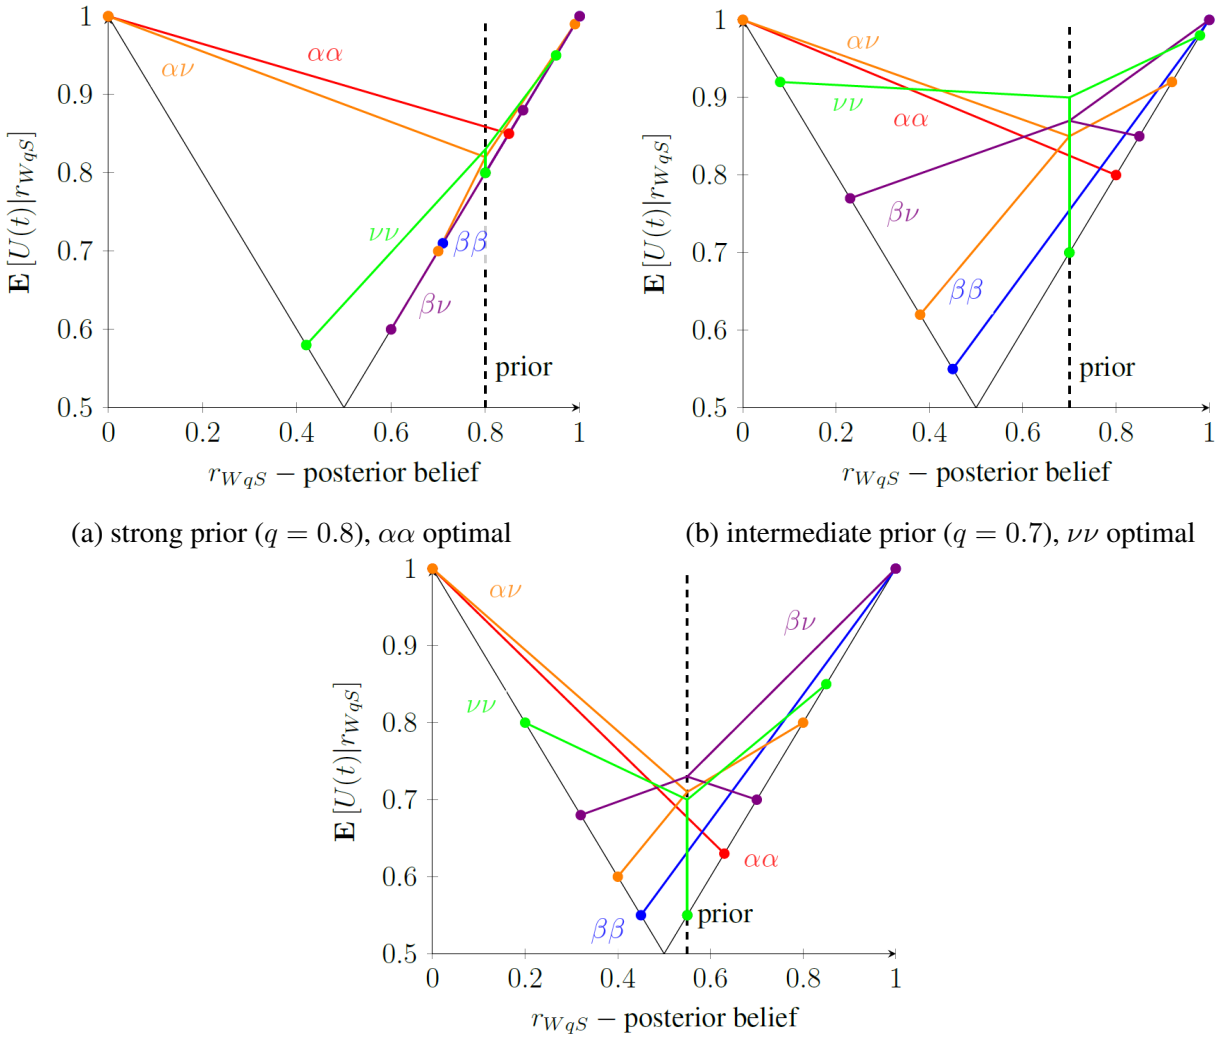


(a) strong prior ($q = 0.8$), $\alpha\alpha$ optimal

(b) intermediate prior ($q = 0.7$), $\nu\nu$ optimal

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

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 portfolio $\alpha\beta$.

### 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 $e$, $e_\nu$, and $q$. I derive the exact optimality conditions in Appendix A.2. 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 $e_\nu$ 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 intermediate priors. However, some agents with weak beliefs are always left with a combination of neutral and opposing sources. Appendix A.3 plots the expected utilities for two selected combinations of $e$ and $e_\nu$ for further illustration.

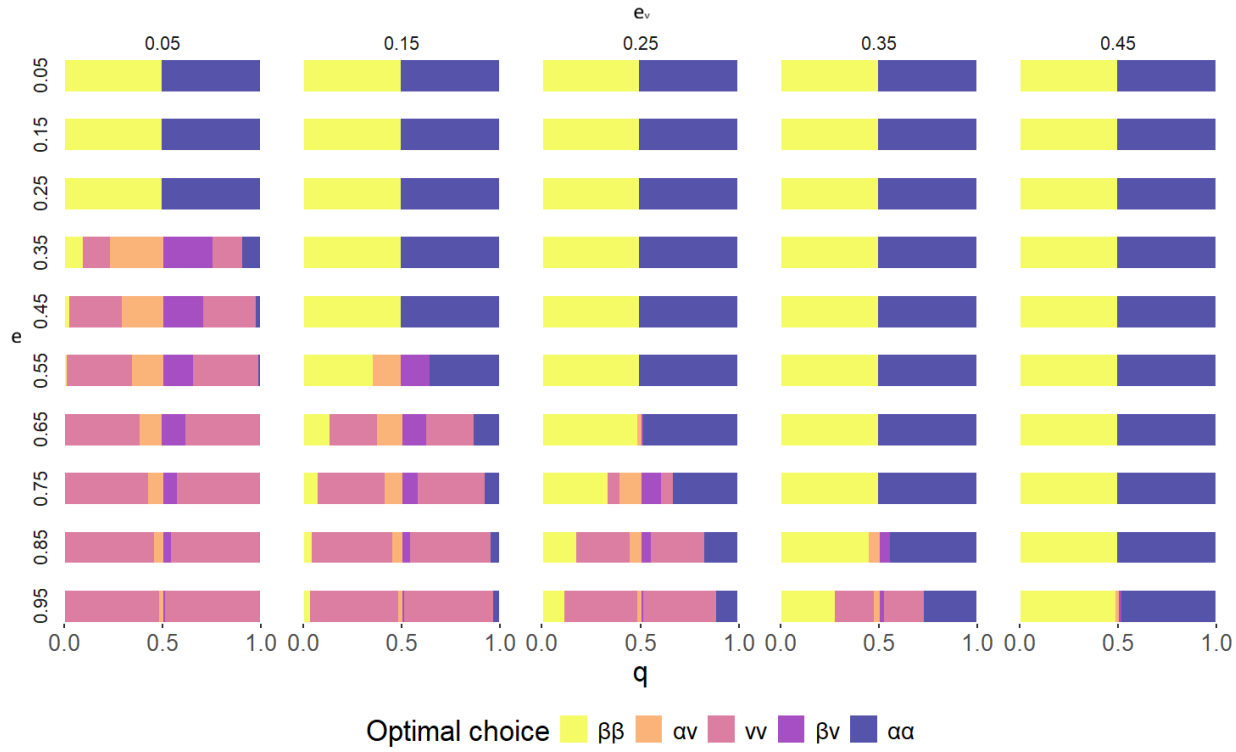


Figure 5: Optimal choice of two sources for different parameter combinations

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

## 4.3 Demand for $n$ Signals

The role of neutral sources as informational stabilizers and complements generalizes beyond two-signal portfolios. The portfolia there can be 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 $\binom{3+n-1}{n}$ different portfolios for a given $n$.

Propositions 4 and 5 establish that the mechanism is not a two-signal artifact. The same structural pattern survives for arbitrary portfolio sizes. Even when agents may acquire many signals,

opposing information enters the optimal portfolio only through its interaction with neutral information.

**Proposition 4** (Optimal Portfolio Candidates with $n$ Signals)**.** *Let the agent 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 $e_\nu$, strong $q$),*
   - *all-neutral (for small $e_\nu$, weak $q$), or*
   - *one opposing, $n-1$ neutral,*
   - *one aligned, $n-1$ neutral.*

*Proof.* See Appendix A.4. □

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.

For every even $n$, the optimality of an opposing-neutral portfolio can be further characterized. Once the neutral source is sufficiently reliable, at least some agents with weak prior beliefs choose this portfolio.

**Proposition 5** (One opposing-rest neutral optimality for even $n$)**.** *Let the agent be able to acquire $n > 2$ signals, where $n$ is even. Then, either*

1. *The only optimal portfolios are all-aligned, or*
2. *There exists a range of priors around $q = 0.5$, where the one opposing-rest neutral portfolio dominates.*

*Proof.* See Appendix A.5. □

The intuition for even $n$ is quite straightforward. With even $n$, and one signal acquired from the opposing source, the agent can further rely on an odd number of neutral sources. Therefore, even in the unfortunate case when the opposing source does not reveal the true state of the world, the agent can rely on decisive information from neutral sources. Since there are an odd number of

them, they must point to one of the states; there cannot be a draw. This provides insurance even against the unfavorable signal realization.

Figure 6a(a) illustrates the optimality regions for an even $n$. Indeed, either everyone chooses only the aligned source, or there is a region of opposing-neutral consumers. In the latter case, the general pattern is that moderate-belief consumers opt for the neutral source, while extreme-belief consumers choose the aligned source. Additionally, for narrow ranges of $q$, symmetric around $q = 0.5$, one aligned-rest neutral, or one opposing-rest neutral, may become optimal.

Conversely, for odd $n$, combining one opposing signal with $n-1$ neutral sources puts the agent into danger of the opposing signal not revealing the true state of the world, and at the same time, the neutral sources delivering an equal number of opposing and aligned signals. This realization would leave the agent quite uncertain; as a result, it is often dominated by other portfolios. Nevertheless, for some parameter configurations, it can still dominate. Figure 6a(b) shows the optimal choice of five sources, demonstrating the general patterns described.

# 5 Robustness and Scope of the Mechanism

The baseline model makes several simplifying assumptions regarding preferences, signal structures, and the timing of information acquisition. This section investigates which of these assumptions are essential for the emergence of opposing-neutral portfolio optimality. The analysis proceeds from the most fundamental features of the model to more specific modeling choices.

Turns out that under full commitment, demand for opposing information is robust to alternative utility functions, more general signal structures, and imperfectly biased sources. However, it relies on two ingredients: (i) the availability of a source capable of resolving the uncertainty generated by the opposing signal, and (ii) a discrete portfolio environment that allows informational complementarities across sources.

## 5.1 General Preferences

Since the main intuition does not rely on the linearity of the utility function, the results also generalize to more sophisticated functional forms. Intuitively, the news consumption patterns depend on the convexity/concavity of the utility function on each side of the flat belief threshold. If the function is convex on each subinterval, sharpening belief is most valuable when belief already is quite sharp. That means the agent should favor signal structures that ensure at least moderate posteriors and avoid realizations that would leave the agent very uncertain. This results in the agents avoiding the opposing-neutral portfolios, especially when the probability of complete resolution of uncertainty remains quite low. Conversely, for concave utility functions, the opposing-neutral

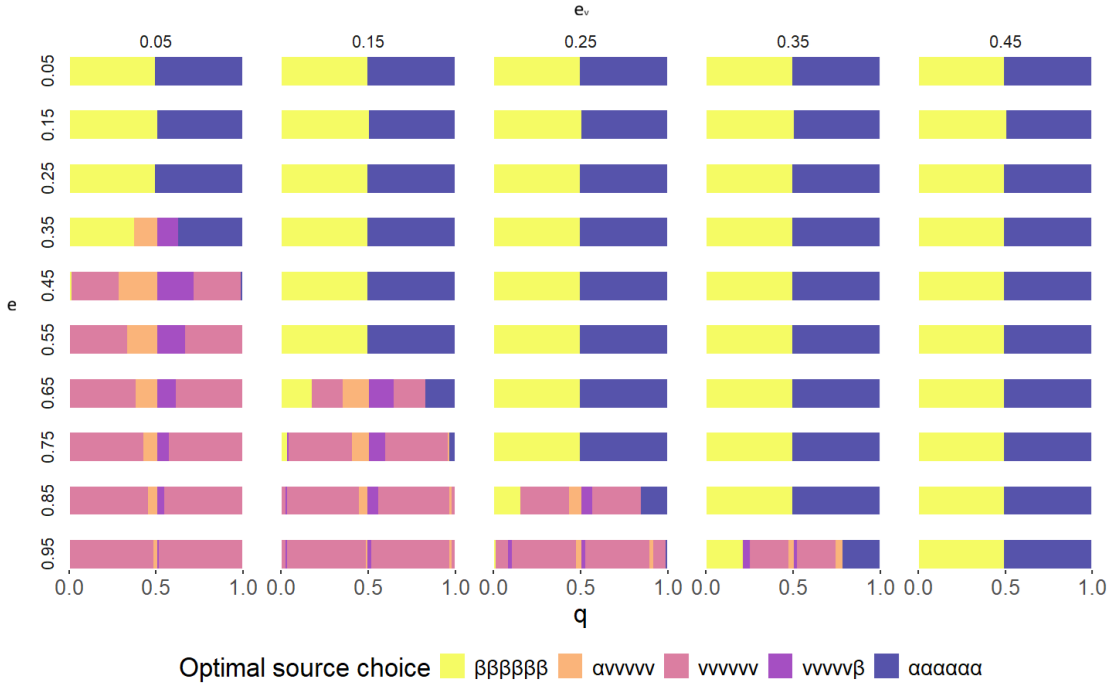


(a) $n = 6$

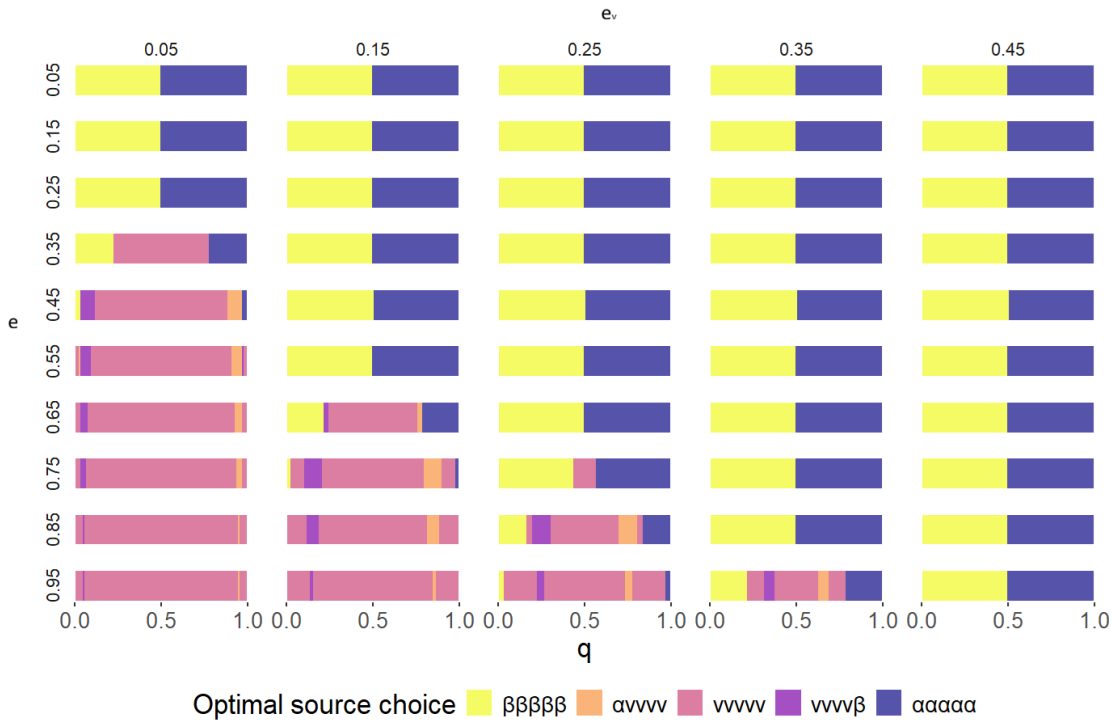


(b) $n = 5$

Figure 6: Optimal source choices for different values of $q$, $e$, and $e_\nu$, and $n$.

portfolios become more appealing, as they offer a significant improvement even when uncertainty is only partially resolved. Figure 7 illustrates this consideration.

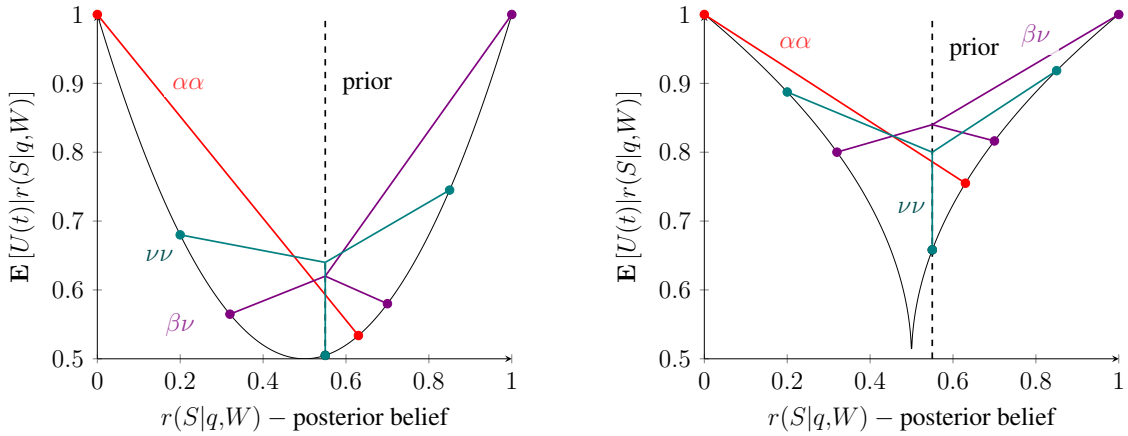


(a) piecewise convex expected utility

(b) piecewise concave expected utility

Figure 7: Optimal choice of two signals with alternative underlying expected utility functions

Note: The dashed black line shows the prior (here $q = 0.55$). The intersections with this line demonstrate the expected utility of the given prior for the given portfolio. The choice of error rates $e = 0.8$, $e_\nu = 0.3$ enables comparison with the benchmark case (Figure 3(c)).

Appendix A.6 explores the optimality regions for a class of polynomial functions, confirming the aforementioned intuition. Overall, the main result of this paper– optimality of the opposing-neutral portfolio– survives as long as the utility function is not too convex.

## 5.2 General Signal Structures

### What Complements an Opposing Source?

The baseline model assumes that the signal complementing the opposing source is neutral. To understand whether neutrality itself is essential, I replace the neutral signal with a generic signal $\kappa$. This signal is neither completely biased nor perfectly neutral. Instead, the error rate in state $a$ is $\varepsilon_a$, and the error rate in state $b$ is $\varepsilon_b$, where in general $\varepsilon_a \neq \varepsilon_b$. Thus, $\kappa$ may be biased toward either state, but its degree and direction of bias are allowed to vary independently.

For fixed prior $q$ and baseline signal accuracy $e$, I evaluate the expected utility of all feasible two-signal portfolios, ($\alpha\alpha$, $\alpha\beta$, $\beta$,$\beta$, $\alpha\kappa$, $\beta\kappa$, and $\kappa\kappa$) and identify the optimal portfolio for every pair $(\varepsilon_a,\varepsilon_b) \in [0,0.5]^2$). Appendix A.7 shows that the signals that optimally complement $\beta$ are not necessarily neutral. Rather, the optimal complementary signals are typically biased towards $a$ (in other words, more precise in state $a$, mathematically $\varepsilon_b > \varepsilon_a$), but only moderately so. As the bias becomes too strong, source $\kappa$ ceases to provide sufficiently reliable information to insure against

the nonrevealing realization of the opposing signal.

More broadly, the role of the second source is not to be exactly neutral. Its role is to resolve residual uncertainty left by the opposing source. Neutral signals provide a particularly natural example of such a complement, but they are not the only possible ones.

**Imperfect Revelation**

A similar generalization can be applied to the biased signals. In the baseline model, each biased source perfectly reveals one of the states of the world. This assumption can be relaxed by allowing biased sources to make mistakes in both states, provided that they remain more accurate in the state toward which they are biased. Specifically, let source $\widehat{\alpha}$ satisfy

$$P(B_{\widehat{\alpha}}|a) = c < P(A_{\widehat{\alpha}}|b) = e,$$

and define $\widehat{\beta}$ symmetrically by

$$P(A_{\widehat{\beta}}|b) = c, \qquad P(B_{\widehat{\beta}}|a) = e,$$

where $0 < c < e < 1$.

Appendix A.8 derives the complete one-source characterization under this more general signal structure. The main qualitative conclusions of Propositions 1 and 2 remain unchanged. For agents with priors above one-half, the generalized aligned source weakly dominates the generalized opposing source, while the neutral source remains optimal only for sufficiently moderate prior beliefs. For sufficiently strong priors, none of the available signals changes the agent's optimal action, so all sources yield the same expected payoff and the agent is indifferent among them.

More importantly, the paper's central portfolio result also survives this generalization. Appendix A.8 establishes that there exist parameter values for which a portfolio combining a generalized opposing source with a neutral source is uniquely optimal. Thus, perfect revelation from biased sources is not the driving force behind the complementarity identified in the paper. Instead, the complementarity arises because opposing and neutral sources provide distinct informational advantages. The generalized opposing source remains relatively informative precisely when the less likely state occurs, whereas the neutral source provides more balanced information across states. Combining the two, therefore, allows the agent to exploit complementary informational strengths that neither source possesses individually.

## 5.3 Discrete versus Continuous Information Acquisition

This subsection connects the portfolio-based mechanism developed in Propositions 3 and 4 with the canonical continuous-signal benchmark in rational inattention and Bayesian persuasion.

In the standard signal design problem (Matějka and McKay, 2015; Caplin et al., 2022), the agent chooses a single binary signal structure $\sigma$ characterized by error rates

$$c := P(B_\sigma|a), \quad e := P(A_\sigma|b),$$

subject to a cost that increases with the informativeness of the signal. In this setting, perfectly informative signals are typically prohibitively costly, so the agent optimally trades off the benefits of more precise information against its acquisition cost.

This framework provides a natural benchmark for my model, which instead assumes costless signals but limits their number. One way to extend my model toward the standard signal-design setting is to allow for flexible signal structures while imposing a feasibility constraint such as $e + c \geq \varepsilon$, without introducing an explicit cost of information. The resulting choice set is convex in $(c,e)$, while the objective is linear in posterior beliefs. Hence, optimal signal structures lie at extreme points of the feasible set. For an agent with prior $q \geq 1/2$, this implies a corner solution with $e = \varepsilon$ and $c = 0$, i.e., a fully aligned experiment. In this sense, the continuous-signal benchmark delivers a sharp prediction: optimal information acquisition involves only signals aligned with the prior.

The portfolio problem studied here differs from this benchmark. Rather than choosing a point in a convex space of signal structures, the agent selects a finite portfolio of discrete sources, each with a fixed error structure, subject to an attention constraint. The induced information structure is determined by the joint distribution of these sources rather than by a single experiment. As a consequence, the set of implementable aggregate information structures is no longer convex. This non-convexity underlies the characterization results in Propositions 3 and 4: introducing a neutral source can generate discontinuous changes in optimal portfolios, and only a small subset of feasible allocations can be optimal.

The portfolio formulation may also provide a natural description of real-world news consumption. Individuals typically allocate attention to a small number of discrete information sources—newspapers, television channels, newsletters, or social media accounts rather than designing statistical experiments. Moreover, these sources are often perceived categorically as broadly neutral or leaning toward one side of an issue, making the portfolio problem a more natural abstraction than continuous signal design. As shown above, the main mechanism is robust to both imperfect bias and imperfect neutrality, suggesting that it arises from the discreteness of available information sources rather than from the particular signal specification.

## 5.4 Static versus Sequential Acquisition

To further distinguish the mechanism studied here from dynamic attention-allocation models such as Che and Mierendorff (2019), I introduce a cost $\gamma > 0$ for acquiring a second signal and compare two timing structures.

In the sequential problem, the agent first chooses and observes one signal, then decides whether to acquire a second signal and, if so, which one, at cost $\gamma$. Acquisition and expenditure can therefore be conditioned on the first signal realization. In the static problem, the agent must instead commit ex ante either to one free signal or to a two-signal portfolio in which the second signal costs $\gamma$. The latter formulation preserves the full-commitment structure of the baseline model while allowing the agent to choose the size of the portfolio.

Figure 8 organizes the comparison along two dimensions. The upper row excludes the neutral source, while the lower row includes it. The left column represents sequential acquisition, whereas the right column imposes static full commitment.

Panel (a) provides a simplified qualitative illustration of the cost-dependent pattern in Che and Mierendorff (2019). It is included as a conceptual benchmark rather than as a quantitative solution to the sequential realization of the present model. For low signal costs, the agent acquires both opposing and aligned sources. At intermediate costs, she switches to aligned information only, while at sufficiently high costs, she acquires no additional signal. Thus, sequential flexibility can make opposite-biased learning optimal when additional information is sufficiently inexpensive.

The remaining three panels report optimal strategies in the model studied here for the representative calibration

$$e = 0.8, \qquad e_\nu = 0.3, \qquad q = 0.55,$$

for which the baseline full-commitment model exhibits $\beta\nu$ optimality. Appendix A.9 provides the corresponding calculations of continuation values and cost thresholds.

Panel (b) removes contingent acquisition while continuing to exclude the neutral source. The opposing-first strategy optimality then disappears. For sufficiently low costs, the agent commits to two aligned signals, $\alpha\alpha$. Once the cost of the second signal exceeds $\gamma_{\text{SIMNO}}$, the agent acquires only one aligned signal, $\alpha$. Thus, in the absence of a neutral source, replacing sequential acquisition with full commitment restores the aligned-only benchmark.

Panel (c) introduces the neutral source into the sequential problem. For sufficiently low costs, the agent begins with $\nu$ and conditions the second source on its realization: after $A$, the agent acquires $\alpha$, whereas after $B$, the agent acquires $\beta$. Importantly, although $\beta$ is opposing relative to the ex-ante prior, it is aligned with the agent's interim belief following the neutral realization $B$.

As costs increase, the aligned signal is dropped before the opposing signal, illustrating the stronger complementarity between the neutral and opposing sources than between the neutral and

aligned sources. Intuitively, after receiving $A$ the agent is already fairly confident that the state is $a$, so an additional signal is unlikely to alter her action sufficiently to justify its cost. By contrast, the opposing signal becomes especially valuable after receiving $B$. In that case, the agent's posterior belief is relatively weak, creating a greater incentive to purchase a second signal and reduce uncertainty. At sufficiently high costs, however, even the opposing signal is no longer worthwhile, and the agent acquires only $\nu$ in the first stage.

Panel (d) illustrates the paper's central mechanism. Under full commitment, the agent cannot condition the second-source choice on the first realization. Because the opposing signal generates greater value than the aligned signal when paired with the neutral source, the agent chooses the portfolio $\beta\nu$ rather than $\alpha\nu$ once this flexibility is removed. When the cost exceeds $\gamma_{\mathrm{SIM}}$, the agent acquires only the neutral signal.

Taken together, the four panels highlight two distinct mechanisms through which an ex-ante opposing source can appear in an optimal information strategy. Sequential acquisition allows the agent to condition subsequent learning and expenditure on earlier signal realizations. Neutral-source complementarity can include an opposing source in an optimal portfolio even when the entire portfolio is chosen ex ante. Because the exercise focuses on a single calibration, its role is primarily illustrative rather than to provide a general characterization of the costly-acquisition environment.

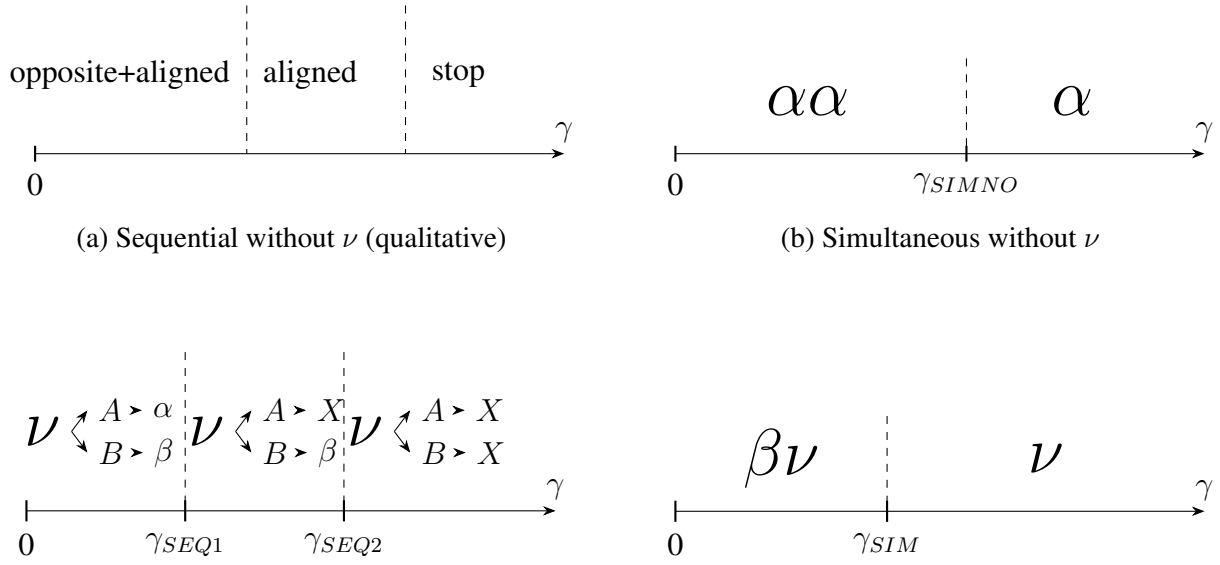


(a) Sequential without $\nu$ (qualitative)

(b) Simultaneous without $\nu$

(c) Sequential with $\nu$

(d) Simultaneous with $\nu$

Figure 8: Optimal strategies as a function of the second signal cost $\gamma$.

Note: Panel (a) provides a simplified qualitative representation of the cost-dependent pattern in Che and Mierendorff (2019). Panels (b)–(d) report optimal strategies in the present model for $e = 0.8$, $e_\nu = 0.3$, and $q = 0.55$. The upper row excludes the neutral source, while the lower row includes it. Sequential acquisition is shown in the left column and static full commitment in the right column. In Panel (c), the first source is chosen before any realization is observed; branches labeled $A$ and $B$ report the continuation choice following the first realization, and $X$ denotes stopping. Vertical dashed lines identify cost thresholds at which the optimal strategy changes. The thresholds are illustrative, and the horizontal axes are not drawn to a common scale. See Appendix A.9 for the calculations.

# 6 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. Neutral news insure against the opposing source failing to reveal while preserving the possibility of immediate certainty. However, when combined with an aligned source, they pose a serious threat of contradicting the non-revealing realization of the aligned source, leaving the agent uncertain.

More broadly, the analysis identifies a mechanism for exposure to opposing viewpoints that is distinct from the dynamic continuation or other mechanisms studied in the literature. Under full commitment, opposing sources are never optimal when only biased sources are available. They become part of an optimal information portfolio only when paired with sufficiently reliable neutral information.

The results are robust to the exact specification of preferences or to imperfect signal structures. Rather, it relies on two central features: the availability of a source that can resolve the residual uncertainty generated by the opposing signal, and a discrete information environment in which complementarities across sources become relevant.

Taken together, these results suggest that exposure to diverse perspectives may emerge endogenously through rational information acquisition, highlighting a potentially underappreciated role of neutral media in facilitating such exposure.

# 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 $\alpha$. 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 $\alpha$ is

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

which can be simplified to

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

Using similar logic, we get

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

and

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

By comparing these expected utilities, we get that

- source $\alpha$ dominates source $\beta$ for all $q \in (0.5,1)$ (and vice versa for $q \in (0,0.5)$)
- for $e_\nu/e < 0.5$ and $q \in \left(\frac{e_\nu}{e}, 1-\frac{e_\nu}{e}\right)$, $\nu$ 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 imply that the portfolio $\alpha\beta$ is dominated by $\alpha\alpha$ (since $e > e^2$)

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

| **Portfolio** | **Expected Utility** |
|---|---|
| —rule $\alpha\alpha$ | $1 - e^2(1-q)$ |
| $\alpha\beta$ | $1 - e(1-q)$ |
| $\beta\beta$ | $\begin{cases} 1 - e^2 q, & \text{if } e < \sqrt{\frac{1-q}{q}} \\ q(e^2+1) - 1, & \text{if } e \geq \sqrt{\frac{1-q}{q}} \end{cases}$ |
| $\alpha\nu$ | $\begin{cases} (1-e_\nu)q + [e(1-e_\nu) + (1-e)], & 0.5 \leq q < \frac{e(1-e_\nu)}{e(1-e_\nu)+e_\nu} \\ q + (1-e)(1-q), & q \geq \frac{e(1-e_\nu)}{e(1-e_\nu)+e_\nu} \end{cases}$ |
| $\beta\nu$ | $\begin{cases} (1-e)q + (1-q), & \text{if } q < q_{\beta\nu L} \\ (1-ee_\nu)q + (1-q)(1-e_\nu), & \text{if } q_{\beta\nu L} \leq q < q_{\beta\nu H} \\ q, & \text{if } q \geq q_{\beta\nu H} \end{cases}$ |
| $\nu\nu$ | $\begin{cases} (1-e_\nu)^2 + 2(1-e_\nu)e_\nu q, & \text{if } 0.5 < q < \frac{(1-e_\nu)^2}{(1-e_\nu)^2+e_\nu^2} \\ q, & \text{if } 1 > q \geq \frac{(1-e_\nu)^2}{(1-e_\nu)^2+e_\nu^2} \end{cases}$ |

and is therefore never optimal. Moreover, $\alpha\alpha$ always dominates $\beta\beta$, because

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

$$\text{for} \quad e > \sqrt{\frac{1-q}{q}}: \quad \mathbf{E}\left[U|\alpha\alpha, q > 0.5\right] = qe^2 + 1 - e^2 \geq qe^2 + q - 1 = \mathbf{E}\left[U|\beta\beta, q > 0.5\right],$$

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

This proves the first part of the proposition: in the absence of the neutral source, the unique optimal portfolio is $\alpha\alpha$. Furthermore, $\alpha\nu$ is always dominated by either $\beta\nu$ or $\alpha\alpha$. Define

$$q_{\alpha\nu} \equiv \frac{e(1-e_\nu)}{e(1-e_\nu) + e_\nu}.$$

If $q \geq q_{\alpha\nu}$, then

$$\mathbf{E}[U|\alpha\nu] = q + (1-e)(1-q),$$

and therefore

$$\mathbf{E}[U|\alpha\nu] - \mathbf{E}[U|\alpha\alpha] = -e(1-e)(1-q) < 0.$$

Hence, $\alpha\nu$ is dominated by $\alpha\alpha$ in this region.

Now consider $0.5 < q < q_{\alpha\nu}$. In this region,

$$\mathbf{E}[U|\alpha\nu] = (1 - e_\nu)q + [e(1 - e_\nu) + (1 - e)]\,(1 - q),$$

while

$$\mathbf{E}[U|\beta\nu] = (1 - ee_\nu)q + (1 - q)(1 - e_\nu).$$

Their difference is

$$\mathbf{E}[U|\alpha\nu] - \mathbf{E}[U|\beta\nu] = e_\nu(e - 1)(2q - 1) < 0.$$

Thus, $\alpha\nu$ is dominated by $\beta\nu$ in this region. Consequently, $\alpha\nu$ is never optimal. By this, I have excluded $\alpha\nu$, $\beta\beta$, and $\alpha\beta$ from the candidates for the optimal portfolio for $q > 0.5$. The only remaining candidates are $\alpha\alpha$, $\nu\nu$, and $\beta\nu$.

To characterize the relevant regions, define the switching points

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

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

and

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

Since $0 < e < 1$ and $0 < e_\nu < 0.5$, we have

$$q_{\beta\nu L} < q_{\beta\nu H}, \qquad q_{\nu\nu} > 0.5, \qquad q_{\beta\nu H} > 0.5.$$

For $q > q_{\beta\nu H}$, we have

$$\mathbf{E}[U|\beta\nu] = q < 1 - e^2(1 - q) = \mathbf{E}[U|\alpha\alpha],$$

where the inequality follows from

$$(1 - q)(1 - e^2) > 0.$$

Thus, $\beta\nu$ cannot be optimal for $q > q_{\beta\nu H}$.

Similarly, $\nu\nu$ can only be optimal if $q \leq q_{\nu\nu}$. For $q > q_{\nu\nu}$,

$$\mathbf{E}[U|\nu\nu] = q < 1 - e^2(1 - q) = \mathbf{E}[U|\alpha\alpha],$$

so $\nu\nu$ is dominated by $\alpha\alpha$. It follows that, for

$$q > \bar{q} \equiv \max\{q_{\nu\nu}, q_{\beta\nu H}\},$$

both $\nu\nu$ and $\beta\nu$ are strictly dominated by $\alpha\alpha$. Since the other three portfolios have already been excluded, $\alpha\alpha$ is the unique optimal portfolio for $q > \bar{q}$. Moreover, $q_{\nu\nu}, q_{\beta\nu H} \in (0.5, 1)$, and hence $\bar{q} \in (0.5, 1)$.

For $q < \bar{q}$, the elimination arguments above imply that the optimal portfolio must be one of

$$\alpha\alpha, \qquad \nu\nu, \qquad \beta\nu.$$

Finally, to prove the third part of the proposition, consider a right neighborhood of $q = 0.5$. Suppose that

$$e_\nu < \frac{e^2}{1+e}.$$

This condition implies

$$q_{\beta\nu L} < 0.5 < q_{\beta\nu H},$$

so, in a right neighborhood of $q = 0.5$,

$$\mathbf{E}[U|\beta\nu] = (1 - ee_\nu)q + (1-q)(1-e_\nu).$$

At $q = 0.5$,

$$\mathbf{E}[U|\beta\nu] - \mathbf{E}[U|\alpha\alpha] = \frac{e^2 - e_\nu(1+e)}{2} > 0,$$

while

$$\mathbf{E}[U|\beta\nu] - \mathbf{E}[U|\nu\nu] = \frac{e_\nu(1-e)}{2} > 0.$$

Thus, $\beta\nu$ strictly dominates both $\alpha\alpha$ and $\nu\nu$ at $q = 0.5$. By continuity of the expected utilities, there exists $\varepsilon > 0$ such that $\beta\nu$ is the unique optimal portfolio for every

$$q \in (0.5, 0.5 + \varepsilon).$$

□

## A.3 Expected utility as a function of $q$ for $n = 2$

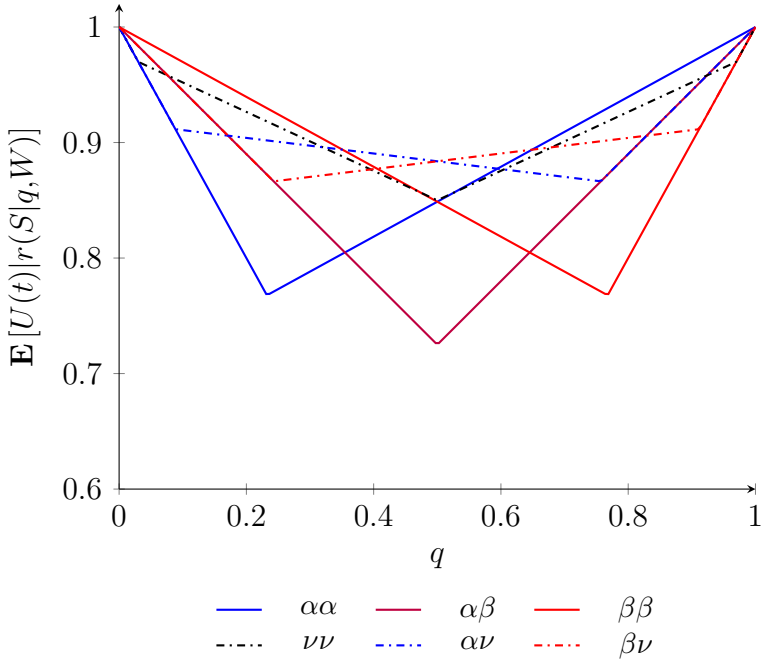


(a) $e = 0.55$, $e_\nu = 0.15$

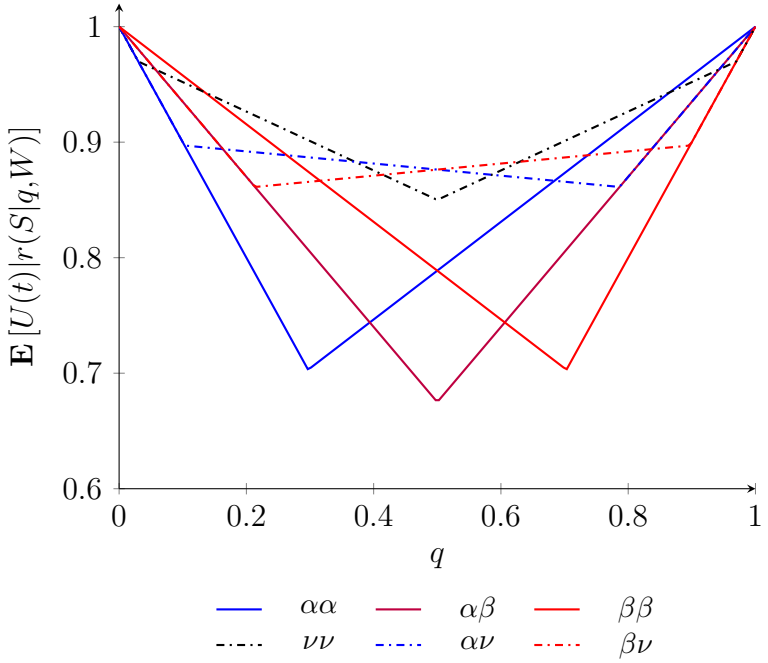


(b) $e = 0.65$, $e_\nu = 0.15$

Figure A.1: Expected utility of the different source choices for different values of $q$.

## A.4 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 $\alpha$ and $\beta$, 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 $\alpha$ for agents with $q > 0.5$ and $\beta$ for agents with $q < 0.5$

In this simplified problem, the agent chooses $k_\alpha$, and $k_\beta$ such that $k_\alpha + k_\beta = n$. Then, she gets $k_\alpha$ signals from source $\alpha$ and $k_\beta$ signals from source $\beta$. Let me denote by $k_{\alpha B}$ the number of signals $B$ sent by source $\alpha$, and by $k_{\beta A}$ the number of signals $A$ sent by source $\beta$. 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 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 $k_{\beta A}$

| $k_{\alpha B}$ | $k_{\beta A}$ | $P[S\|q,W]$ | $r(S\|q,W)$ | $\mathbf{E}[U\|q,W,S]$ |
|---|---|---|---|---|
| $> 0$ | $= 0$ | $(1-e^{k_\alpha})(1-q)$ | 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

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

Here, for a while, I abstract from the fact that $k_\alpha$ 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

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

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

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 $k_\alpha$ 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

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

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

The agent always chooses the $k_\alpha$ 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 $\alpha$, and all agents who initially believed $b$ is more likely to be the true state of the world consume only from source $\beta$.

**Including the neutral source**

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

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

The expected utility for a particular choice of $k_\alpha$,$k_\nu$, and $k_\beta$ 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

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

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

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

The general problem of the agent could thus be formulated as

$$\max_{k_\alpha,k_\beta,k_\nu}U(k_\alpha,k_\beta,k_\nu)$$

$$U(k_\alpha, k_\beta, k_\nu) = (1 - e^{k_\alpha})(1 - q) + (1 - e^{k_\beta})q +$$

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

$$k_\alpha + k_\beta + k_\nu = n,$$

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

We have already shown that, when only sources $\alpha$ and $\beta$ are available, no interior allocation with $k_\alpha > 0$ and $k_\beta > 0$ is optimal. The presence of neutral signals does not alter this result. For any fixed realization of the neutral signals, the utility contribution of the non-neutral sources has exactly the same form as in the two-source problem, differing only by a positive multiplicative factor inside the maximum operator. Hence, for every realization of the neutral signals, utility is maximized by assigning all non-neutral signals to a single biased source. Taking expectations over neutral-signal realizations preserves this ordering. Therefore, no optimal portfolio contains both $\alpha$ and $\beta$, and it is sufficient to consider portfolios containing at most one biased source.

Therefore, the only candidates for an optimal portfolio are

1. $(n,0,0)$ all-aligned
2. $(0,0,n)$ all-neutral
3. $(k,0,n-k), \quad k \in 1,..,n-1$ aligned-neutral family
4. $(0,k,n-k), \quad k \in 1,..,n-1$ opposing-neutral family

**Lemma A.1.** *Assume* $0 < e < 1$, $0 < e_\nu < \frac{1}{2}$, $\frac{1}{2} < q < 1$, *and* $n \in \mathbf{N}$, $n > 2$ *Then for every* $k = 2, \ldots, n-1$,

$$U(k,0,n-k) \leq \max\{U(n,0,0), U(0,0,n)\}.$$

*Proof.* Let

$$m = n - k, \qquad X \sim \mathrm{Bin}(m, 1 - e_\nu).$$

Since $0 < e_\nu < \frac{1}{2}$, the ratio

$$\frac{q e_\nu^{m-i}(1 - e_\nu)^i}{e^k(1 - q) e_\nu^i (1 - e_\nu)^{m-i}} = \frac{q}{e^k(1 - q)} \left(\frac{e_\nu}{1 - e_\nu}\right)^{m-2i}$$

is strictly increasing in $i$. Hence there exists a unique threshold

$$t_k = \left\lceil \frac{1}{2}\left(m - \frac{\log\left(\frac{e^k(1-q)}{q}\right)}{\log\left(\frac{e_\nu}{1-e_\nu}\right)}\right)\right\rceil$$

such that

$$qe_\nu^{m-i}(1-e_\nu)^i \geq e^k(1-q)e_\nu^i(1-e_\nu)^{m-i} \iff i \geq t_k.$$

Therefore,

$$\begin{aligned}
&\sum_{i=0}^{m}\binom{m}{i}\max\left[qe_\nu^{m-i}(1-e_\nu)^i, e^k(1-q)e_\nu^i(1-e_\nu)^{m-i}\right] \\
&= q\,\Pr(X \geq t_k) + e^k(1-q)\Pr(X < t_k) \\
&= e^k(1-q) + \left(q - e^k(1-q)\right)\Pr(X \geq t_k).
\end{aligned}$$

Substituting into $U(k,0,n-k)$ gives

$$\begin{aligned}
U(k,0,n-k) &= (1-e^k)(1-q) + e^k(1-q) + \left(q - e^k(1-q)\right)\Pr(X \geq t_k) \\
&= (1-q) + \left(q - e^k(1-q)\right)\Pr(X \geq t_k).
\end{aligned}$$

Since $q > \frac{1}{2}$,

$$U(n,0,0) = (1-e^n)(1-q) + q = 1 - (1-q)e^n.$$

Thus

$$U(n,0,0) - U(k,0,n-k) = q - (1-q)e^n - \left(q - e^k(1-q)\right)\Pr(X \geq t_k).$$

Because $k < n$ and $0 < e < 1$,

$$e^n < e^k,$$

and hence

$$q - (1-q)e^n > q - (1-q)e^k.$$

Since $0 \leq \Pr(X \geq t_k) \leq 1$,

$$\left(q - e^k(1-q)\right)\Pr(X \geq t_k) \leq q - e^k(1-q) < q - (1-q)e^n.$$

Therefore

$$U(n,0,0) - U(k,0,n-k) > 0,$$

so that

$$U(k,0,n-k) \leq U(n,0,0) \qquad (k = 2,\ldots,n-1).$$

Finally,

$$U(n,0,0) \leq \max\{U(n,0,0), U(0,0,n)\},$$

which yields

$$U(k,0,n-k) \leq \max\{U(n,0,0), U(0,0,n)\}.$$

□

This shows that portfolios that combine the neutral source with several aligned signals are never optimal; thus, $(1,0,n-1)$ is the only potential optimum in the aligned-neutral family.

The proof of a similar claim for the opposing-neutral family follows.

**Lemma A.2** (Opposing-neutral boundary dominance)**.** *Assume* $0 < e < 1$, $0 < e_\nu < \frac{1}{2}$, $\frac{1}{2} < q < 1$, *and* $n > 2$. *Then for every* $k = 2, \ldots, n-1$,

$$U(0,k,n-k) \le \max\{U(n,0,0), U(0,0,n), U(1,0,n-1), U(0,1,n-1)\}.$$

*Proof.* Let

$$m = n - k,$$

and define

$$X \sim \mathrm{Bin}(m, 1-e_\nu).$$

Since $0 < e_\nu < \frac{1}{2}$, the ratio

$$\frac{e^k q\, e_\nu^{m-i}(1-e_\nu)^i}{(1-q)e_\nu^i(1-e_\nu)^{m-i}} = \frac{e^k q}{1-q}\left(\frac{e_\nu}{1-e_\nu}\right)^{m-2i}$$

is strictly increasing in $i$. Therefore there exists a threshold $t_k$ such that

$$e^k q\, e_\nu^{m-i}(1-e_\nu)^i \ge (1-q)e_\nu^i(1-e_\nu)^{m-i} \iff i \ge t_k.$$

Hence

$$\sum_{i=0}^{m} \binom{m}{i} \max\left[e^k q\, e_\nu^{m-i}(1-e_\nu)^i, (1-q)e_\nu^i(1-e_\nu)^{m-i}\right] = (1-q) + \left(e^k q - (1-q)\right)\Pr(X \ge t_k).$$

Substituting into $U(0,k,n-k)$ yields

$$U(0,k,n-k) = 1 - e^k q + \left(e^k q - (1-q)\right)\Pr(X \ge t_k).$$

Now distinguish two cases. *Case 1:* $e^k q \le 1-q$. Then,

$$U(0,k,n-k) \le 1 - e^k q.$$

Since $k < n$, $0 < e < 1$, and $q > 1-q$,

$$e^k q > e^n(1-q),$$

and therefore

$$1 - e^k q < 1 - (1-q)e^n = U(n,0,0).$$

Hence

$$U(0,k,n-k) < U(n,0,0).$$

*Case 2:* $e^k q > 1 - q$. There,

$$U(0,k,n-k) \leq 1 - e^k q + e^k q - (1-q) = q.$$

On the other hand,

$$U(0,1,n-1) = (1-e)q + \sum_{i=0}^{n-1} \binom{n-1}{i} \max\left[eq\, e_\nu^{n-1-i}(1-e_\nu)^i, (1-q)e_\nu^i(1-e_\nu)^{n-1-i}\right].$$

Since $\max(a,b) \geq a$,

$$U(0,1,n-1) \geq (1-e)q + eq \sum_{i=0}^{n-1} \binom{n-1}{i} e_\nu^{n-1-i}(1-e_\nu)^i = q.$$

it is that

$$U(0,k,n-k) \leq q \leq U(0,1,n-1).$$

Thus, for every $k \geq 2$,

$$U(0,k,n-k) \leq \max\{U(n,0,0), U(0,1,n-1)\}.$$

This completes the proof. □

By this, it is established that only all-aligned, all-neutral, one aligned-rest neutral, and one opposing-rest neutral can become the optimal portfolio choice.

## A.5 Proof of Proposition 5

*Proof.* I prove the statement for $1 > q > 0.5$, the logic applies symmetrically to $0 < q < 0.5$. First, I show that for even $n$, $(k_\alpha, k_\beta, k_\nu) = (0,1,n-1)$ dominates $(k_\alpha, k_\beta, k_\nu) = (0,0,n)$ for some interval $(0.5, 0.5 + t)$. All expected utility expressions are piecewise linear in $q$. Moreover, the expected utility for $\nu\nu$ is symmetric around $q = 0.5$. That yields that there exists $0.5 > t > 0$ such

that the expected utility inside $q \in [0.5,0.5+t]$ for a given even $n$ is

$$\begin{aligned}
&\mathbf{E}[U|q \in (0.5,0.5+t),(0,0,n)] = \\
&= \sum_{k_{\nu A}=0}^{n} \binom{n}{k_{\nu A}} \max\left\{ q e_\nu^{n-k_{\nu A}} (1-e_\nu)^{k_{\nu A}},\, (1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{n-k_{\nu A}} \right\} = \\
&= \binom{n}{0}(1-q)(1-e_\nu)^n + (1-q)\binom{n}{1} e_\nu (1-e_\nu)^{n-1} + .. + q\binom{n}{\frac{n}{2}} e_\nu^{\frac{n}{2}} (1-e_\nu)^{\frac{n}{2}} + \\
&\quad ... + \binom{n}{n-1} q e_\nu (1-e_\nu)^{n-1} + \binom{n}{0} q (1-e_\nu)^n = \\
&= \binom{n}{\frac{n}{2}} e_\nu^{\frac{n}{2}} (1-e_\nu)^{\frac{n}{2}} q + \sum_{i=0}^{\frac{n}{2}-1} \binom{n}{n-i} (1-e_\nu)^{n-i} e_\nu^i
\end{aligned}$$

Next, I demonstrate that at $q = 0.5$, the oppposing-neutral portfolio at least weakly dominates $\nu\nu$. Notice that the expected utility from the opposing-neutral portfolio at $q = 0.5$ is decreasing in $e$. Therefore, for a given $e_\nu$, the minimal expected utility is attained for $e = 1$ and equals

$$\begin{aligned}
&\mathbf{E}[U|q = 0.5,(0,1,n-1)] = \\
&= \sum_{k_{\nu A}=0}^{n-1} \binom{n-1}{k_{\nu A}} \frac{1}{2} \max\left\{ q e_\nu^{k_\nu - k_{\nu A}} (1-e_\nu)^{k_{\nu A}},\, (1-q) e_\nu^{k_{\nu A}} (1-e_\nu)^{k_\nu - k_{\nu A}} \right\} = \\
&= \sum_{i=0}^{\frac{n}{2}} \binom{n-1}{n-i} (1-e_\nu)^{n-i} e_\nu^i
\end{aligned}$$

Comparing the two expected utilities and applying standard rules for sums and binomials[A.1] yields

$$\binom{n}{\frac{n}{2}} e_\nu^{\frac{n}{2}} (1-e_\nu)^{\frac{n}{2}} \frac{1}{2} + \sum_{i=0}^{\frac{n}{2}-1} \binom{n}{n-i} (1-e_\nu)^{n-i} e_\nu^i = \sum_{i=0}^{\frac{n}{2}} \binom{n-1}{i} (1-e_\nu)^{n-1-i} e_\nu^i.$$

That means that in the extreme case of $e = 1$, the all-neutral and one opposing-rest neutral deliver the same expected utility. Given that the one opposing-rest neutral expected utility decreases with $e$, any value $1 > e > 0$ makes this portfolio dominate the all-neutral one.

At $q = 0.5$, by symmetry, the expected utility from a one aligned-rest neutral portfolio should

[A.1] Using $\binom{n}{n-i} = \binom{n}{i}$, the left-hand side equals $\sum_{i=0}^{\frac{n}{2}-1} \binom{n}{i} e_\nu^i (1-e_\nu)^{n-i} + \frac{1}{2}\binom{n}{\frac{n}{2}} e_\nu^{\frac{n}{2}} (1-e_\nu)^{\frac{n}{2}}$. Applying Pascal's identity and splitting the middle term $\frac{1}{2}\binom{n}{\frac{n}{2}} = \frac{1}{2}\binom{n-1}{\frac{n}{2}} + \frac{1}{2}\binom{n-1}{\frac{n}{2}-1}$ one verifies that all terms combine pairwise into $\sum_{i=0}^{\frac{n}{2}} \binom{n-1}{i} e_\nu^i (1-e_\nu)^{n-1-i}$, which is the right-hand side.

always be equal to the one opposing-rest neutral. However, at the same time, the one aligned-rest neutral utility is decreasing in $q$, whereas the one opposing-rest neutral utility is increasing in $q$.

Taken together, that implies that there is a region around $q = 0.5$ where the one opposing-rest neutral dominates both all-neutral and one aligned-rest neutral cases.

What remains is to show that the only other optimum for given $e$, $e_\nu$, $n$ is that the whole continuum of agents opts for an all-aligned portfolio. This is easily demonstrated by recalling that $\mathbf{E}[U|q,(n,0,0)]$ is strictly increasing in $q$ on $q \in (0.5,1)$. Therefore, it remains to show that there exists such a parameter combination where $\mathbf{E}[U|q = 0.5,(n,0,0)] \geq \mathbf{E}[U|q = 0.5,(0,1,n-1)]$

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

The sum can be bound using the fact that $e < 1$ by

$$\begin{aligned}
&\sum_{k_{\nu A}=0}^{n-1} \binom{n-1}{k_{\nu A}} \max\left\{e e_\nu^{n-1-k_{\nu A}}(1-e_\nu)^{k_{\nu A}},\, e_\nu^{k_{\nu A}}(1-e_\nu)^{n-1-k_{\nu A}}\right\} \leq \\
&\leq \sum_{k_{\nu A}=0}^{n-1} \binom{n-1}{k_{\nu A}} \max\left\{e_\nu^{n-1-k_{\nu A}}(1-e_\nu)^{k_{\nu A}},\, e_\nu^{k_{\nu A}}(1-e_\nu)^{n-1-k_{\nu A}}\right\} = \\
&= \sum_{i=0}^{\frac{n}{2}} \binom{n-1}{n-1-i} (1-e_\nu)^{n-1-i} e_\nu^i \leq 1
\end{aligned}$$

That yields

$$e - e^n \geq \sum_{i=0}^{\frac{n}{2}} \binom{n-1}{n-i} (1-e_\nu)^{n-i} e_\nu^i$$

This separates $e$ and $e_\nu$ into different sides of the equation, giving a condition for many admissible $e$ and $e_\nu$ to satisfy this equation for any $n$. That yields the conclusion that the all-aligned portfolio can indeed become optimal for all agents. □

## A.6 General Polynomial Expected Utility

This exercise further develops the generalization to other utility functions as sketched in Section 5.1. I start by defining the requirements on the utility function given by the underlying economic problem

- $U(r(S|q,W))$ - the expected utility depends on the posterior belief only,
- $U(r(S|q,W))$ is symmetric around $r(S|q,W) = \frac{1}{2}$,
- $U(1) = U(0) = 1$ - certainty delivers the highest possible utility,
- $U\left(\frac{1}{2}\right) = \frac{1}{2}$ - flat posterior belief delivers utility equal to that of choosing randomly,
- $U(r(S|q,W))$ is decreasing on $(0,\frac{1}{2})$ and increasing on $(\frac{1}{2},1)$ - stronger beliefs bring higher payoffs.

Henceforth, I restrict my analysis to piecewise polynomial functions. To satisfy the assumptions and parametrize the convexity/concavity with a single number, I choose the following functional form

$$U(r(S|q,W)) = 2^{p-1} \cdot \left| r(S|q,W) - \frac{1}{2} \right|^p + \frac{1}{2}, \quad p > 0,$$

where parameter $p$ governs the (piecewise) convexity/concavity of the function ($p > 1$ convex, $p < 1$ concave, $p = 1$ linear benchmark case).

Intuition suggests that opposing-neutral optimality should persist at least for some $e$ and $e_\nu$ as long as the utility function is *not too convex*. I demonstrate this with the following exercise. I fix values of $e$ and $e_\nu$ and vary the parameter $p$. For each $p$, I show the optimal portfolio choice as a function of $q$. Figure A.2 presents the result of this exercise. Indeed, the prevalence of opposing-neutral choice decreases with $p$, confirming the intuition presented in the text.

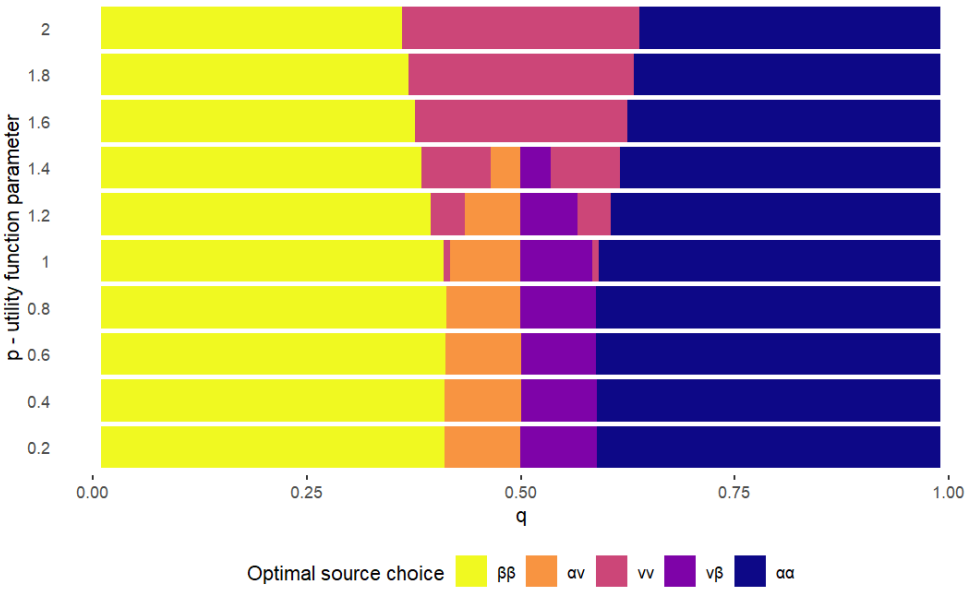


Figure A.2: The impact of utility function shape (as measured by $p$) on optimal portfolio choice

Note: $e = 0.8$, $e_\nu = 0.3$, $n = 2$; Higher values of $p$ are associated with piecewise convexity, lower values with piecewise concavity, $p = 1$ represents the benchmark linear case.

## A.7 Imperfect Neutral Signal

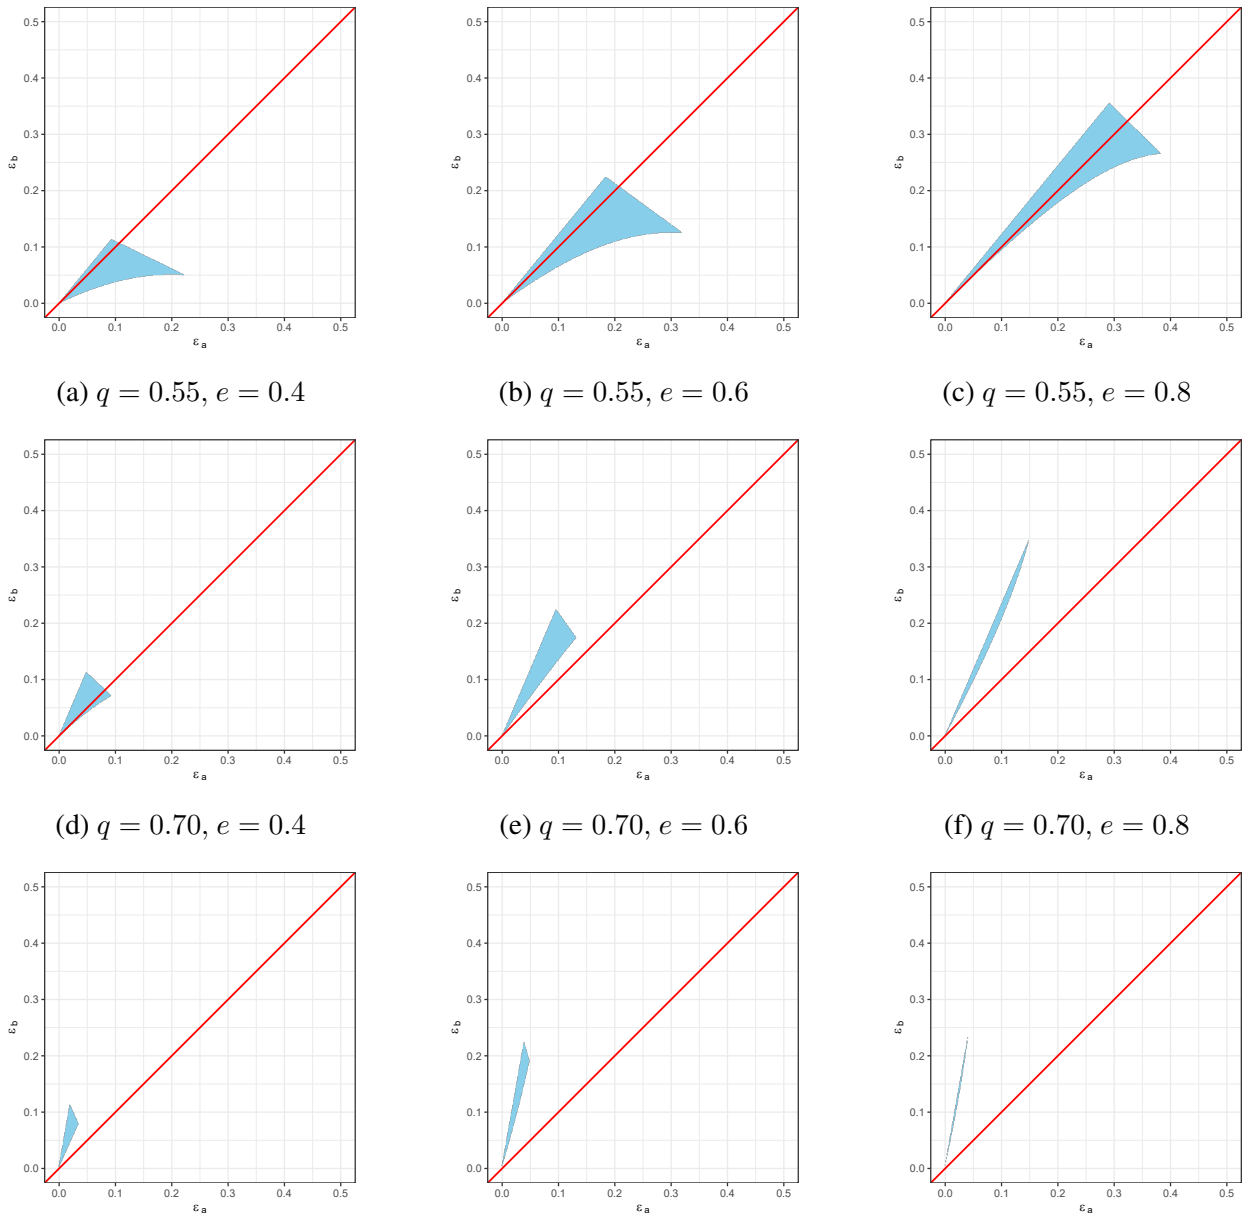


(a) $q = 0.55$, $e = 0.4$ (b) $q = 0.55$, $e = 0.6$ (c) $q = 0.55$, $e = 0.8$

(d) $q = 0.70$, $e = 0.4$ (e) $q = 0.70$, $e = 0.6$ (f) $q = 0.70$, $e = 0.8$

(g) $q = 0.85$, $e = 0.4$ (h) $q = 0.85$, $e = 0.6$ (i) $q = 0.85$, $e = 0.8$

Figure A.3: Region of general signals $\kappa$ for which the portfolio $(\beta,\kappa)$ is optimal.

Note: The horizontal axis reports the error rate of source $\kappa$ in state $a$ ($\varepsilon_a$), while the vertical axis reports its error rate in state $b$ ($\varepsilon_b$). The blue region identifies parameter combinations for which $(\beta,\kappa)$ yields the highest expected utility among all feasible two-signal portfolios. The red diagonal corresponds to $\varepsilon_a = \varepsilon_b$, i.e., neutral signal.

## A.8 General Biased Signals

Assume the following set of general biased signals and a neutral signal.

Table A.4: General biased signals

| news sources | | |
|---|---|---|
| $\hat{\alpha}$ | $\nu$ | $\hat{\beta}$ |
| biased towards $a$ | neutral | biased towards $b$ |
| $P(A\|b) = e$ | $P(A\|b) = e_\nu$ | $P(A\|b) = c$ |
| $P(B\|a) = c$ | $P(B\|a) = e_\nu$ | $P(B\|a) = e$ |

Assume throughout this section that

$$0 < c < e < 1, \qquad c + e < 1, \qquad 0 < e_\nu < \frac{1}{2}.$$

The condition $c + e < 1$ ensures that the signal realizations are informative in the direction suggested by their labels.

I formulate the result for $q \in [1/2,1]$. The corresponding result for $q \in [0,1/2]$ follows symmetrically.

**Proposition** (Optimal one-source choice under general biased signals)**.** *Suppose that*

$$0 < c < e, \qquad c + e < 1, \qquad 0 < e_\nu < \frac{1}{2},$$

*and let* $q \in [1/2,1]$. *Define*

$$q_A := \frac{1-e}{1-e+c},$$

$$q_N := \min\left\{1 - e_\nu, \frac{e - e_\nu}{e - c}\right\},$$

*and*

$$\overline{q} := \max\left\{q_A, 1 - e_\nu\right\}.$$

*Then:*

1. *In the absence of the neutral source,* $\widehat{\alpha}$ *weakly dominates* $\widehat{\beta}$ *for every* $q \in [1/2,1]$.

2. *If the neutral source is available, then:*

   (a) *if*

   $$\frac{1}{2} < q < q_N,$$

   *the neutral source is uniquely optimal;*

*(b) if*

$$\max\left\{\frac{1}{2},q_N\right\} < q < \overline{q},$$

*the aligned source $\widehat{\alpha}$ is uniquely optimal;*

*(c) if*

$$q \geq \overline{q},$$

*all three sources yield expected payoff $q$, and the agent is indifferent among them.*

*If $q_N \leq 1/2$, the region in which the neutral source is uniquely optimal is empty.*

*Proof.* For any binary signal $M$, expected utility can be written directly as the sum, over the two signal realizations, of the larger of the two state-contingent joint probabilities:

$$V_M(q) = \sum_{s\in\{A,B\}} \max\left\{qP(s|a,M), (1-q)P(s|b,M)\right\}.$$

For the generalized aligned source,

$$V_{\widehat{\alpha}}(q) = \max\left\{(1-c)q, e(1-q)\right\} + \max\left\{cq, (1-e)(1-q)\right\}.$$

Because $q \geq 1/2$ and $1-c > e$, the first maximum is always attained by $(1-c)q$. Therefore,

$$V_{\widehat{\alpha}}(q) = (1-c)q + \max\left\{cq, (1-e)(1-q)\right\}.$$

Equivalently,

$$V_{\widehat{\alpha}}(q) = \max\left\{1-e+(e-c)q, q\right\}.$$

The two expressions intersect at

$$q_A = \frac{1-e}{1-e+c}.$$

Hence,

$$V_{\widehat{\alpha}}(q) = \begin{cases} 1-e+(e-c)q, & \frac{1}{2} \leq q < q_A, \\ q, & q_A \leq q \leq 1. \end{cases}$$

Similarly,

$$V_{\widehat{\beta}}(q) = \max\left\{(1-e)q, c(1-q)\right\} + \max\left\{eq, (1-c)(1-q)\right\}.$$

Since $q \geq 1/2$ and $1-e > c$, the first maximum is always attained by $(1-e)q$.

Thus,

$$V_{\widehat{\beta}}(q) = \max\left\{1-c-(e-c)q, q\right\}.$$

For every $q \geq 1/2$,

$$[1 - e + (e - c)q] - [1 - c - (e - c)q] = (e - c)(2q - 1) \geq 0.$$

Because both expected-utility expressions also contain the common alternative $q$, it follows that

$$V_{\widehat{\alpha}}(q) \geq V_{\widehat{\beta}}(q)$$

for every $q \in [1/2,1]$.

This proves the no-neutral result.

For the neutral source,

$$V_\nu(q) = \max\{(1 - e_\nu)q, e_\nu(1 - q)\} + \max\{e_\nu q, (1 - e_\nu)(1 - q)\}.$$

Since $q \geq 1/2$ and $e_\nu < 1/2$, the first maximum is attained by $(1 - e_\nu)q$.

Therefore,

$$V_\nu(q) = \max\{1 - e_\nu, q\},$$

or equivalently,

$$V_\nu(q) = \begin{cases} 1 - e_\nu, & \frac{1}{2} \leq q < 1 - e_\nu, \\ q, & 1 - e_\nu \leq q \leq 1. \end{cases}$$

Because $\widehat{\alpha}$ weakly dominates $\widehat{\beta}$, it remains to compare $V_{\widehat{\alpha}}$ and $V_\nu$. In the region in which both sources potentially improve on the prior payoff $q$, the neutral source is preferred if and only if

$$1 - e_\nu \geq 1 - e + (e - c)q.$$

Rearranging gives

$$q \leq \frac{e - e_\nu}{e - c}.$$

The neutral source must also satisfy $q \leq 1 - e_\nu$ in order to produce more than the prior payoff.

It is therefore optimal for

$$q \leq q_N := \min\left\{1 - e_\nu, \frac{e - e_\nu}{e - c}\right\}.$$

Above this threshold, the aligned source is optimal until all sources yield only the prior payoff $q$. The aligned source yields $q$ for $q \geq q_A$, while the neutral source yields $q$ for $q \geq 1 - e_\nu$.

Moreover,

$$\frac{1-c}{1-c+e} < \frac{1-e}{1-e+c} = q_A,$$

because

$$\frac{1-e}{1-e+c} - \frac{1-c}{1-c+e} = \frac{(e-c)(1-e-c)}{(1-e+c)(1-c+e)} > 0.$$

Hence, by the time $q \geq q_A$, the opposing source also yields expected payoff $q$. It follows that all three sources yield payoff $q$ whenever

$$q \geq \overline{q} := \max\{q_A, 1-e_\nu\}.$$

This establishes the stated characterization. □

Allowing the biased sources to make mistakes in both states creates a region of sufficiently strong priors in which no available signal changes the agent's optimal action. In this region, each source yields expected payoff $q$, and the agent is indifferent among them. The signal realizations may still change posterior beliefs; the point is that this belief updating has no decision value because the agent chooses action $a$ after every realization.

Once we allow the agent to choose more signals, the main intuition applies again. Instead of deriving the full optimality conditions for an arbitrary $n$ $e$, $e_\nu$,$c$, $q$, I show that for some combination of parameters, there exists a region of opposing-neutral optimality for the case of two signals.

**Proposition** (Opposing-neutral optimality with imperfectly biased signals)**.** *There exist parameters satisfying*

$$0 < c < e, \qquad c+e < 1, \qquad 0 < e_\nu < \frac{1}{2},$$

*and a nonempty interval of priors immediately above* $1/2$ *for which the portfolio* $\widehat{\beta}\nu$ *is uniquely optimal among all two-signal portfolios.*

*Proof.* Consider

$$e = 0.8, \qquad c = 0.05, \qquad e_\nu = 0.3.$$

These parameters satisfy

$$0 < c < e, \qquad c+e < 1, \qquad e_\nu < \frac{1}{2}.$$

There are six unordered two-signal portfolios:

$$\widehat{\alpha}\widehat{\alpha}, \quad \widehat{\alpha}\widehat{\beta}, \quad \widehat{\beta}\widehat{\beta}, \quad \widehat{\alpha}\nu, \quad \widehat{\beta}\nu, \quad \nu\nu.$$

Let

$$\delta := q - \frac{1}{2}.$$

For $q$ in a sufficiently small right neighborhood of $1/2$, direct evaluation of the four signal realizations gives

$$V_{\widehat{\beta}\nu}(q) = \frac{57}{80} + \frac{19}{200}\delta,$$

$$V_{\widehat{\alpha}\nu}(q) = \frac{57}{80} - \frac{19}{200}\delta,$$

$$V_{\nu\nu}(q) = \frac{7}{10} + \frac{21}{50}\delta,$$

$$V_{\widehat{\alpha}\widehat{\alpha}}(q) = \frac{101}{160} + \frac{217}{400}\delta,$$

$$V_{\widehat{\beta}\widehat{\beta}}(q) = \frac{101}{160} - \frac{217}{400}\delta,$$

and

$$V_{\widehat{\alpha}\widehat{\beta}}(q) = \frac{23}{40} + \frac{77}{100}\delta.$$

At $q = 1/2$,

$$V_{\widehat{\beta}\nu}\left(\frac{1}{2}\right) = V_{\widehat{\alpha}\nu}\left(\frac{1}{2}\right) = \frac{57}{80},$$

while

$$\frac{57}{80} > \frac{7}{10} > \frac{101}{160} > \frac{23}{40}.$$

Thus, the two mixed biased–neutral portfolios are tied for the highest expected payoff at the symmetric prior. For every $\delta > 0$,

$$V_{\widehat{\beta}\nu}(q) - V_{\widehat{\alpha}\nu}(q) = \frac{19}{100}\delta > 0.$$

Moreover,

$$V_{\widehat{\beta}\nu}(q) - V_{\nu\nu}(q) = \frac{1}{80} - \frac{13}{40}\delta.$$

This difference remains strictly positive whenever

$$0 < \delta < \frac{1}{26}.$$

The remaining portfolios are separated from $\widehat{\beta}\nu$ by strictly larger gaps at $q = 1/2$; the expressions above show that these gaps also remain positive for all sufficiently small $\delta > 0$. In particular, they remain positive for

$$0 < \delta \leq \frac{1}{100}.$$

Consequently,

$$V_{\widehat{\beta}\nu}(q) > V_W(q)$$

for every other two-signal portfolio $W$ whenever

$$q \in \left(\frac{1}{2}, \frac{51}{100}\right].$$

Thus, the opposing-neutral portfolio is uniquely optimal on a nonempty interval of priors even though both biased sources make mistakes in both states. □

This example establishes that perfect revelation by the biased sources is not necessary for opposing–neutral optimality. The result is an existence claim rather than a complete characterization for arbitrary $(q,e,c,e_\nu)$. Nevertheless, it shows that the central portfolio-complementarity mechanism survives when both biased sources make mistakes in both states.

## A.9 Separating Dynamic Attention from Portfolio Choice

This Appendix derives the quantitative strategies reported in Panels (b)–(d) of Figure 8. Panel (a) is a simplified qualitative representation of the cost-dependent dynamic mechanism highlighted by Che and Mierendorff (2019); it is not obtained by solving the present model under the calibration below.

Throughout this section, I consider the calibration

$$e = 0.8, \qquad e_\nu = 0.3, \qquad q = 0.55.$$

Under the baseline full-commitment model, this calibration yields $\beta\nu$ as the optimal two-signal portfolio.

**Sequential with $\nu$** Figure A.4 represents the sequential problem as an extensive-form game for this given calibration. The agent first chooses a signal, observes its realization, and then decides whether to purchase a second signal at cost $\gamma$. The terminal nodes report the expected utility associated with each continuation strategy. Strategies that are never optimal for this parameterization are omitted for clarity.

The optimal policy follows from comparing the continuation values at each decision node. For the purpose of this paper, I only report the equilibrium paths.

The agent always begins by acquiring the neutral signal, which provides a balanced first assessment of the state. When information acquisition is inexpensive ($\gamma \leq 0.08 := \gamma_{SEQ1}$), the agent then acquires a second, biased signal tailored to the initial realization: after observing $A$, she

acquires $\alpha$, whereas after observing $B$, she acquires $\beta$. Thus, the neutral signal serves as a bridge, directing attention to the source expected to be most informative given the initial evidence.

As the cost of additional information increases, the value of refining beliefs declines. For intermediate costs ($0.08 < \gamma \leq 0.09 := \gamma_{SEQ2}$), the agent acquires a second signal only after one realization of the neutral signal, while for sufficiently high costs ($\gamma > \gamma_{SEQ2}$) she stops after the initial neutral signal regardless of its realization. Throughout, the first signal remains neutral, reflecting its role as the most valuable source for reducing overall uncertainty before committing to a more specialized source. The example illustrates how optimal information acquisition naturally combines broad initial learning with increasingly selective follow-up as information becomes more costly.

**Simultaneous with $\nu$** For comparison, consider the static version of the problem in which the agent commits ex ante to her complete information portfolio. Without the option to condition future acquisition on observed realizations, the optimization problem simplifies considerably. The agent either purchases the two-signal portfolio $\beta\nu$ or acquires only the neutral signal. The switching point occurs exactly when the marginal benefit of $\beta\nu$ is equal to the cost of the signal. In this calibration, this occurs for$\gamma_{SIM}(= 0.033)$.

**Simultaneous without $\nu$** By a similar argument, when the neutral source is not available, the agent consumes from the aligned source only. The threshold between acquiring one or two such signals is exactly where the cost $\gamma$ equals the benefit of the additional signal. For the calibration in question, this occurs at $\gamma_{SIMNO} = 0.072$

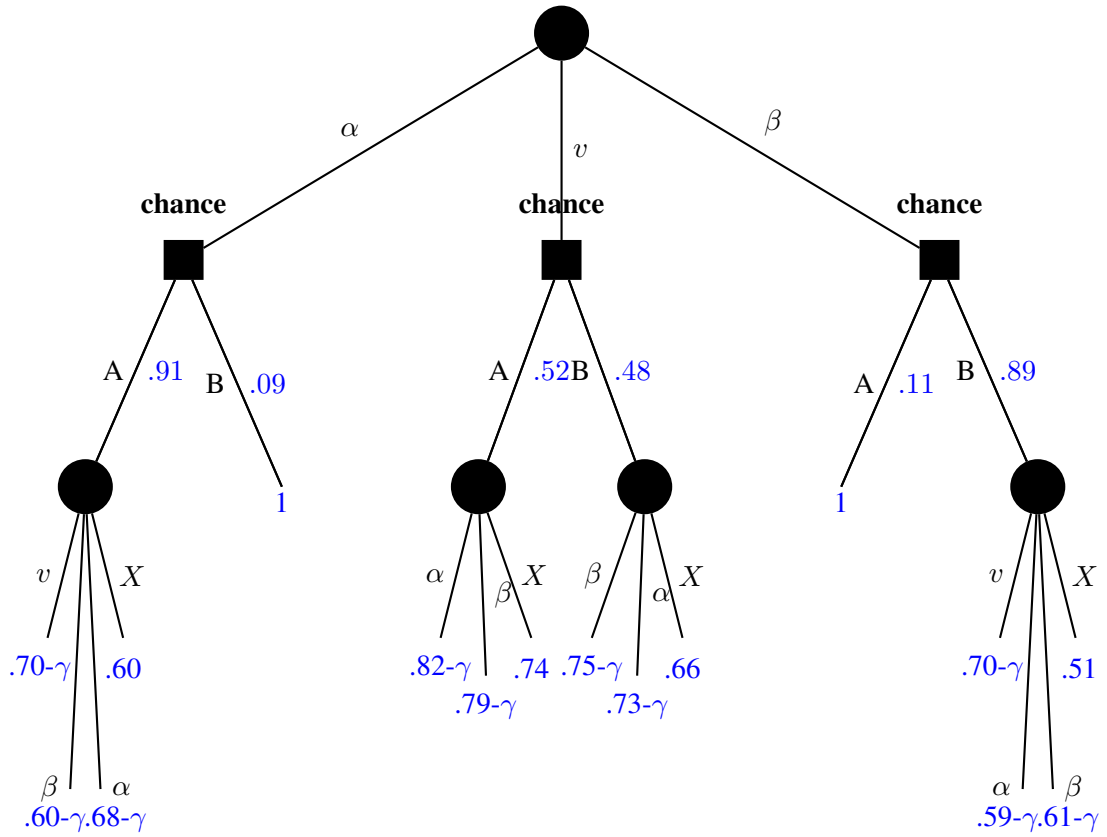


Figure A.4: Tree representation of the sequential problem.

Note: The circles denote the nodes where the agent chooses a signal. The square nodes are where chance chooses the realization of the first signal. After the second signal is realized, the agent can no longer decide on the information acquisition strategy; therefore, the game ends. $X$ denotes the stopping option. The blue numbers denote the expected payoffs and probabilities for $q = 0.55$, $e = 0.8$, $e_\nu = 0.3$; $\gamma$ is the cost of acquiring the second signal. The never-optimal strategies for the given parameter combinations are omitted.

# 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 (2025). The value of and demand for diverse news sources. *Games and Economic Behavior 151*, 24–41.

Caplin, A., M. Dean, and J. Leahy (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy 130*(6), 1676–1715.

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 (2026). When are echo chambers useful?

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.

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–298.

Montanari, G. and S. Nunnari (2025). Audi alteram partem: An experiment on selective exposure to information. *Journal of the Economic Science Association 11*(1), 139–150.

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.