

Optimal Learning with Uncertain Bias^{*}

Emiliano Sandri

University College London

Emma Tao

University College London

Yanziyi Zhang

University College London

Abstract

While biased information sources differ in informativeness, people may be more or less confident about the sources' slant. Learning choices are influenced by the trade-off between more informative and more confounded signals. This paper studies a dynamic information acquisition problem, where attention is distributed across sources with different precision and bias uncertainty. We characterise the optimal attention strategy. The set of consulted information sources gradually expands during the learning process. Precision alone determines the order in which sources enter the information diet, whereas bias uncertainty shapes the length of each learning stage. We further discuss how optimal learning adapts to changes in the information environment.

Keywords: Dynamic information acquisition, Bias-precision trade-off.

JEL codes: D83, D90, L82.

^{*}We would like to especially thank Duarte Gonçalves for the continued encouragement and advice on this project. We also thank Philippe Jehiel, Xiaosheng Mu, Ran Spiegler, and Michael Thaler for valuable feedback.

First posted draft: *date*.

1 Introduction

Information sources vary in precision and bias, and people may have different degrees of knowledge about these characteristics. For example, a consumer deciding whether to purchase a product can read online reviews from various experts. While experts with superior testing instruments post more precise information, it is often unclear whether and how their content is sponsored. Similarly, to learn about relevant news events, readers sample from different media platforms. While some platforms transmit information more clearly, the message conveyed is likely to be influenced by the unknown partisan leaning of the writer. Nowadays, choices are often informed by AI’s suggestions. While some LLM models are more consistent than others, biases of more recent models are likely to be less understood. These settings share a common trade-off: a source provides more informative signals conditional on knowing how to interpret it, but the interpretation may be confounded by the poor understanding of the source’s bias.

While the demand for information sources that differ in precision and bias uncertainty has been documented empirically,¹ it is natural to ask how to optimally allocate attention across such sources. For instance, consider a consumer who can learn about the quality of a pair of headphones from two sources: an amateur review with informal sound tests, and a brand-promoted review with professional equipment. The amateur review is rather imprecise but known to be unbiased, while the sponsored review is more precise but affected by an uncertain bias. The consumer could prioritise learning from the unbiased source, mix between the two, or even switch back and forth depending on the signals they receive. Importantly, the choice of attention is dynamic: what they choose to learn shapes what they observe, which in turn affects what they choose to learn.

This paper studies the optimal dynamic information acquisition strategy — i.e., what to learn and when — from information sources that differ in precision and bias uncertainty. Every source is confounded by an idiosyncratic additive bias unknown to the decision

¹The demand for biased information sources is well-documented in many empirical settings, such as financial advice (Bhattacharya et al. 2012, Anagol et al. 2017, Chalmers & Reuter 2020) and political news (Gentzkow & Shapiro 2006, Chopra et al. 2024, 2025). A growing experimental literature studies the preference for information structures with varying accuracy (Charness et al. 2021, Hoffman 2016, Ambuehl & Li 2018) and different degrees of uncertainty (Shishkin & Ortoleva 2023, Liang 2025).

maker, who can split their limited attention at any time and decide when to stop learning, since time is costly. We show that there exists a unique attention allocation strategy optimal for arbitrary prior belief on the payoff-relevant state and for all decision problems. The learner always begins by allocating full attention to the most precise source, and incrementally adds more sources in their information diet in order of decreasing precision. Learning stages have deterministic thresholds that depend on the bias uncertainty of each source. We further show that changing the precision versus the bias uncertainty of an existing source has markedly different effects on the source demand. Moreover, adding a new source delays consumption of less precise sources, whereas introducing a new unbiased source entirely eliminates such demand.

To capture how signal noise and bias uncertainty are intertwined in learning, we model each information source $i \in \{1, \dots, N\}$ as a Brownian motion with biased drift $\mu + b_i$. μ is the payoff-relevant parameter (e.g., the product’s quality, or the effectiveness of a policy), and b_i the unknown (payoff-irrelevant) bias of source i , independent across sources. Sources differ in their level of Brownian noise, which can arise from measurement errors or varying technologies of information delivery. At each instant, the decision maker chooses how to allocate a limited budget of attention across different sources. Additionally, the decision maker endogenously chooses when to stop learning and act. The Bayesian decision maker has Gaussian prior belief on each source’s bias, with different prior uncertainty across sources (e.g., a reader may be more confident about one news channel’s bias than another).² The optimal control problem is solved by dynamically choosing the attention allocation scheme as well as the optimal stopping time.

Our main results characterise the optimal attention allocation strategy for arbitrary prior belief on μ and for all utility functions. The dynamic experiment induced by the optimal strategy is sequentially Blackwell dominant, hence the strategy we characterise is *uniquely* optimal. Under the optimal strategy, the set of active sources — i.e., those receiving strictly positive attention — gradually expands over time. The decision maker always begins by devoting full attention to the most precise source (e.g., the most articulated product review,

²The prior average bias may also vary across sources. However, our emphasis is on bias uncertainty rather than the level, as we consider a Bayesian decision maker.

the news platform with the highest reporting quality), and then follows a sequence of learning stages where the most precise inactive source is added to the active set.

Under the optimal learning strategy, the set of actively learned sources expands at deterministic points in time. The order in which sources enter the information diet does not depend on bias uncertainty, and it is fully determined by the precision of sources. On the other hand, the timing of expansions of the active set depends on bias uncertainty as well. In particular, higher bias uncertainty anticipates the arrival of subsequent stages, whereas an unbiased source fully absorbs all future attention as soon as it enters the active set. The optimal learning strategy highlights the trade-off between precision and bias uncertainty: high precision entails priority in receiving attention, but high bias uncertainty accelerates the arrival of subsequent stages when attention is diverted to other sources.

We then derive comparative statics that shed light on how people revise their learning plan when the information environment changes. We begin by showing that when an existing source's bias becomes more uncertain, cumulative attention is always diverted towards other sources. This is not the case when a source's precision decreases. The key difference is that changes in precision affect both the extensive and intensive margins of learning: a marginal decrease in precision delays the source's consumption, but it will receive more attention once it enters the active set. If the delay effect prevails, attention will be diverted to other sources; however, if the intensity effect dominates, attention will be *drawn* from other sources. We provide a simple characterisation of which effect prevails based on the source's precision.

We also investigate how the attention plan changes when a new source is introduced in the environment. When a new source is available, the decision maker's optimal strategy is unchanged up to the time the new source enters the active set. If a pre-existing source is noisier than the added one, its consumption is delayed, and the size of the delay is intuitively decreasing in the new source's bias uncertainty. Despite potentially delaying their consumption, the new source does not affect the relative attention ratio between pre-existing sources.

One special case of interest is when none of the available information sources is dominated, that is, sources can be equivalently ordered through decreasing precision or decreas-

ing bias uncertainty. This situation allows us to further emphasise the trade-off between precision and bias uncertainty. In this setting, we show that the source that most recently entered the active set captures most of the attention. Hence, sources with lower precision but more predictable bias may eventually receive more attention than sources higher in priority. We show that this is the case whenever the decision maker learns for a sufficient amount of time.

Our final result addresses the optimal stopping time when the decision maker's prior on μ is Gaussian. Given the optimal attention strategy, the stopping time determines which sources are eventually consumed and the total amount of attention they receive. We show that, for a subclass of utility functions (including the quadratic loss utility function), the optimal stopping time is deterministic and strictly decreasing in the cost of time.

Our results concerning the optimal attention allocation strategy hold for any prior distribution on the payoff-relevant parameter μ . This feature allows us to accommodate a wide variety of different subjective beliefs on different state spaces. Under a Gaussian prior, our biased learning model is isomorphic to the information aggregation problem in [Liang et al. \(2022\)](#). However, the general specification of the prior allows our model to encompass a wider range of economic situations. In [Section 8](#), we discuss and illustrate this point by considering the case of a Bernoulli prior.

The rest of the paper is organised as follows. The remainder of this section discusses the related literature. [Section 2](#) presents the model, and [Section 3](#) offers a stylised example. [Section 4](#) states our general results. [Section 5](#) contains comparative statics exercises, while [Section 6](#) describes the optimal stopping time. We illustrate special cases in [Section 7](#), and conclude by discussing variants of the model in [Section 8](#).

1.1 Related Literature

This paper contributes to the literature on dynamic information acquisition. Since [Wald \(1947\)](#) and [Arrow et al. \(1949\)](#), a large strand of this literature has studied sequential learning problems in which the decision maker optimally chooses when to stop learning and act. Much of this work takes the information process as given and focuses on the resulting speed-accuracy trade-off (see, e.g., [Fudenberg et al. 2018](#), [Gonçalves 2024](#)). [Moscarini &](#)

Smith (2001) complement the optimal stopping problem by allowing the decision maker to control the intensity of learning from a fixed information source. Another approach allows the decision maker to choose from a rich set of information structures (see, e.g., Steiner et al. 2017, Hébert & Woodford 2018, Zhong 2022), subject to an information acquisition cost in the spirit of rational inattention (Sims 2003). Our approach differs, as we study environments in which the decision maker allocates costly time and a limited budget of attention across an exogenously fixed set of information sources. This restriction allows to isolate the trade-off between precision and bias uncertainty when evaluating sources.

Our approach is closest to recent work on dynamic attention allocation across a fixed set of sources. One strand of this literature considers problems where the payoff-relevant parameter is *just-identified*, that is, it can be learnt only by consulting all available sources. For instance, considering Gaussian processes, Liang et al. (2018) characterise cases where myopic information acquisition is optimal, while Liang et al. (2022) provide a characterisation of the optimal attention strategy when the prior belief is Gaussian. With Poisson processes, instead, Mayskaya (2022) studies the allocation of attention across sources that test competing hypotheses. Another strand of this literature considers settings with an *over-identified* parameter. Specifically, Che & Mierendorff (2019) study an attention allocation problem with Poisson processes, where each source provides conclusive evidence on the parameter. Relatedly, Liang & Mu (2020) consider a social learning problem and show that the sequential choices among correlated Gaussian processes may lead to learning traps. In contrast, we study dynamic attention allocation across a fixed set of sources when the parameter is *under-identified*. This is the case as bias uncertainty limits the long-run informativeness of the available sources. As we characterise the optimal attention strategy with Gaussian sources, the closest paper to ours is Liang et al. (2022). Although the two models share some special cases, neither nests the other. While we work with a generic prior on the parameter, we can establish a formal link between the two models in the special case of Gaussian prior (see Section 7.1).

Another related literature studies Bayesian learning when the information structure itself is uncertain.³ The key aspect of this literature is that the decision maker jointly updates

³A distinct literature studies learning from ambiguous information sources under ambiguity-sensitive

beliefs about the payoff-relevant parameter and the information structure itself. [Acemoglu et al. \(2016\)](#) study the resulting asymptotic beliefs when the decision maker learns from a single source of uncertain accuracy. Instead, our paper studies the optimal allocation of attention across multiple sources. In this sense, our work is more closely related to [Sethi & Yildiz \(2016\)](#) and [Gentzkow et al. \(2025\)](#), where the information acquisition strategy is guided by the trade-off between information sources’ precision and bias uncertainty. In these papers, the state changes over time while source characteristics are persistent, allowing the decision maker to learn about the source from repeated observations. In contrast, our environment features a fixed state and asks how attention should be optimally distributed across sources. Such attention allocation problem is absent when the decision maker can only observe a single source ([Aromí 2012](#), [Sethi & Yildiz 2016](#), [Gentzkow et al. 2025](#)) or all sources ([Gentzkow et al. 2025](#)).

More generally, our paper relates to work on biased information. Early contributions show that a Bayesian decision maker may value biased sources because they can be especially informative in particular states ([Calvert 1985](#), [Suen 2004](#), [Burke 2008](#)). Related work studies how media outlets compete for attention by choosing slant or quality (see, e.g., [Mullainathan & Shleifer 2005](#), [Gentzkow et al. 2015](#), [Sobbrio 2014](#), [Chen & Suen 2023](#)). Since the set of sources is fixed in our model, we focus instead on the information consumption side. Moreover, our decision maker has no intrinsic preference for biased information: a biased source attracts attention because it is informative for the Bayesian learner. In this respect, our paper is closest to models of attention allocation across biased or heterogeneous sources, such as [Che & Mierendorff \(2019\)](#), [Hu et al. \(2021\)](#), and the biased-news application in [Liang et al. \(2022\)](#). Relative to these papers, we shift attention from the level of bias to uncertainty about bias, isolating a distinct trade-off that shapes the choice of information sources.

Finally, since we study the decision maker’s choice of a dynamic information structure, our work connects to the literature on ranking dynamic experiments ([de Oliveira 2018](#), [Whitmeyer & Williams 2024](#), [de Oliveira & Lamba 2025](#)). We employ a continuous-time

preferences or updating rules (see, e.g., [Epstein & Schneider 2007, 2008](#), [Tang 2020](#), [Chen 2022](#)). Our mechanism is different: the decision maker is Bayesian, and uncertainty about the information structure matters because the payoff-relevant state and source-specific biases must be learnt jointly.

version of adapted garbling in [de Oliveira \(2018\)](#) to define sequential Blackwell dominance in our setting.

2 Setup

Actions and Payoffs. A decision maker chooses an action $a \in A \subseteq \mathbb{R}$. The payoff $u(a, \mu)$ depends on the action and an unknown payoff-relevant parameter $\mu \in \mathbb{R}$. We refer to (A, u) as a **decision problem**. This general setup allows us to cover a variety of settings, such as

- A consumer either buys ($a = 1$) or not ($a = 0$) a product with unknown binary quality $\mu \in \{l, h\}$ or continuous attribute $\mu \in \mathbb{R}$ (e.g. battery duration) and receives utility $u(a, \mu) = \mathbf{1}_{\{a=1\}}f(\mu)$.
- A donor decides a funding level $a \in \mathbb{R}$ to a political campaign based on macroeconomic condition μ .
- A forecaster makes a prediction a on the margin μ of a future election and gets performance pay $u(a, \mu) = W - (a - \mu)^2$.

Information. Before acting, information about μ can be acquired through sources that differ in bias uncertainty and precision. There are N information sources. Signals from source $i \in \{1, \dots, N\}$ are both noisy with **variance** σ_i and confounded by its idiosyncratic and unknown **payoff-irrelevant bias** parameter b_i .

Without loss of generality, throughout the paper we index information sources in order of increasing variance (or, equivalently, decreasing precision), that is, $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_N$. In the consumer example, each source is an online expert with different ability to test products (σ_i) and unclear sponsorship (b_i). In the donor example, each source is a news platform with different reporting quality (σ_i) and unknown partisan bias (b_i).

We represent each source i as a Brownian motion with drift $\mu + b_i$ and variance σ_i^2 . At each t , the decision maker chooses an **attention allocation** $\beta(t) := (\beta_i(t))_{i=1}^N$ and observes signals $X_t^\beta := (X_{it}^\beta)_{i=1}^N$ generated by the **information processes** defined as

$$dX_{it}^\beta = (\mu + b_i)\beta_i(t) dt + \sigma_i \sqrt{\beta_i(t)} dW_{it}$$

for all $i \in \{1, \dots, N\}$. W_{it} is a Brownian motion, and W_{it} and W_{jt} are independent for all $i \neq j$. The vector $\beta(t)$ describes the dynamic allocation of attention across information sources at each instant t .⁴ Increasing the attention allocated to source i increases the flow of information from the source.

Technically, the choice of $\beta(t)$ can depend only on the information available up to time t . Let $\mathcal{G}_t^\beta$ denote the natural filtration generated by $(\beta(s), X_s^\beta)_{s \in [0, t]}$. Then, we require $\beta(t)$ to be measurable with respect to $\mathcal{G}_t^\beta$. For example, the consumer cannot choose which product to buy based on reviews they have not read yet.

Furthermore, we assume that the decision maker faces an attention budget constraint $\sum_{i=1}^N \beta_i(t) = 1$ for all t , capturing a constrained capacity to process information. For instance, donors do not have the capacity to consult infinitely many news platforms to form an opinion on policies.

We call $\beta := (\beta(t))_{t \geq 0}$ a **feasible attention allocation** if $\beta(t)$ is measurable with respect to $\mathcal{G}_t^\beta$ and $\sum_{i=1}^N \beta_i(t) = 1$ for all $t \in \mathbb{R}_+$. We constrain the choice of attention to feasible attention allocations.

Let $q_i^\beta(t) := \int_0^t \beta_i(s) ds$ denote the **cumulative attention** allocated to source i by time t under $\beta(t)$, and $q^\beta(t) := (q_i^\beta(t))_{i=1}^N$. When the dependence is unambiguous, we will often drop the β superscript for notational convenience.

Beliefs. The decision maker engages in Bayesian learning. Let $\pi_t = (\pi_t^\mu, (\pi_t^{b_i})_{i=1}^N)$, where π_t^μ and $\pi_t^{b_i}$ respectively denote the beliefs about μ and b_i at time t . The prior ($t = 0$) belief on the bias term b_i is assumed to be a normal distribution with mean b_{i0} and variance v_i^2 (i.e., $\pi_0^{b_i} = \mathcal{N}(b_{i0}, v_i^2)$), while the prior belief on the payoff-relevant parameter is an arbitrary Borel probability measure on $\mathbb{R}$. This accommodates a wide range of economic settings, where the support and the distribution of μ vary depending on the application of interest.

We assume that, at $t = 0$, the decision maker regards the payoff-relevant parameter and the bias terms to be independent: $\mu \perp\!\!\!\perp b_i$ for all $i \in \{1, \dots, N\}$, and $b_i \perp\!\!\!\perp b_j$ for all $i, j \in \{1, \dots, N\}$ with $i \neq j$. In the case of consumer learning, each online expert's bias

⁴This can be seen as the limiting case of a discrete-time model in which only one source is sampled in each time period.

arises from their independent private interests.⁵

Source i has a **known bias** if $v_i = 0$. For a Bayesian learner, this is equivalent to learning from an unbiased source with variance σ_i . Hence, we will often refer to a source with known bias as an unbiased source.

Optimal Control Problem. Given their utility and information environment, the decision maker chooses what to learn by designing a feasible (dynamic) attention allocation β , potentially mixing between sources and consuming different subsets at different times. We model impatience for action through a constant **flow cost** $c > 0$, accrued at each instant of time in which learning occurs (i.e., before acting). Then, since delaying the action is costly, the decision maker also dynamically decides when to stop learning (and hence act) by choosing a stopping time τ adapted to $\mathcal{G}_t^\beta$.

Let $g(\pi) = \max_{a \in \mathbb{R}} \mathbb{E}_\pi[u(a, \mu)]$ be the expected payoff given a stopping belief π . The **optimal control problem** is thus to choose (i) a feasible attention allocation β , and (ii) a stopping time τ , to maximise their expected payoff net of time costs:

$$V(\pi_0) = \max_{\tau, \beta} \mathbb{E}_{\pi_0}[g(\pi_\tau) - c\tau] \quad (1)$$

3 An Illustrative Example

We begin by considering the case with (i) two available sources ($N = 2$), (ii) a quadratic loss utility function $u(a, \mu) = -(a - \mu)^2$, and (iii) a Gaussian prior on μ , with mean μ_0 and variance σ_μ^2 . We denote with $\sigma_{\mu,t}^2$ the posterior variance at time t , with $\sigma_{\mu,0}^2 = \sigma_\mu^2$. This simple setting is relevant for applications (such as forecasting), and provides the key insights underlying the general results with $N \geq 2$ sources, general utility functions, and generic prior π_0^μ in Section 4.

Fix $\sigma_1 \leq \sigma_2$, so that source 1 is more precise than source 2, while we do not specify how v_1 and v_2 compare. For instance, news platform 1 has better reporting quality, so its intended message is conveyed more clearly. At the same time, one media outlet has a more predictable slant, whereas the other has a more diverse or unfamiliar writer pool.

⁵We discuss a variant with correlated source biases in Section 8.

Recall that β and τ are chosen jointly. Nevertheless, we argue that the optimal β^* can be determined independently of the optimal τ^* . First, since the prior is Gaussian, the posterior belief about μ will also be Gaussian. Moreover, given the quadratic loss preference, the decision maker's expected utility upon stopping $g(\pi_\tau) = -\sigma_{\mu,\tau}^2$. They thus choose an attention allocation strategy that maximises the rate of reduction of posterior variance at any point in time. Hence, the optimal β^* can be pinned down independently from the optimal τ^* .


Next, we turn to analysing the choice of learning. There are two main tensions at play. First, while source 1 is always more precise, it might have a more uncertain bias, creating a static tension between faster learning speed ($\sigma_1 \leq \sigma_2$) and a more confounded fundamental ($v_1 \geq v_2$). The second tension is dynamic: the long-run informativeness of a source is limited by the uncertainty about its bias, hence a source's marginal value decreases with the cumulative amount of attention it received. Consequently, the strength of the trade-off between sources evolves over time, and an *ex ante* inferior source gradually becomes relatively more valuable as gains from better sources are exhausted.

Given these trade-offs, the decision maker needs to consider how to order their learning and whether there is benefit from splitting attention or switching back and forth between sources. In this simple setting, the posterior variance of μ acts as a **sufficient statistic** for the aggregate information flow: if more information is received, posterior variance decreases faster. Hence, to compare two attention allocations it is sufficient to study how much they reduce posterior variance.

Another feature specific to the present environment is that the posterior variance evolves deterministically as a function of time and of the attention allocation scheme. Hence, we can rephrase the optimal control problem as a **static problem** where, given a fixed time horizon $\bar{t}$, the *cumulative* attention allocation $q^*(\bar{t})$ is chosen to minimise posterior variance at $\bar{t}$. Since β^* is determined independently of τ^* , it must be the case that β^* induces $q^*(\bar{t})$ at all $\bar{t} \geq 0$. Hence, analogously to Lemma 1 in [Liang et al. \(2022\)](#), we can derive the optimal attention allocation from the solution of the static problem as $\beta^*(t) = \frac{dq^*(t)}{dt}$.

We now describe the optimal attention allocation and its characteristics in the case of $N = 2$ sources. The optimal attention allocation β^* consists of two stages with a

Figure 1: Optimal Attention Allocation



Notes. This figure plots the cumulative attention allocated to each source (shaded area between the lines) as a function of time, when there are two available sources ($N = 2$). Each panel shows cases with different uncertainty about the biases. (a) source 2 receives attention after t^* ; (b) same time threshold as (a), but source 2 has higher weight in Stage 2; (c) source 2 is unbiased and captures all attention after t^* .

deterministic time threshold $t^* := \frac{\sigma_1(\sigma_2 - \sigma_1)}{v_1^2}$. In the *first stage* ($t < t^*$), the decision maker exclusively learns from the more precise source (i.e., $\beta_1^*(t) = 1$ and $\beta_2^*(t) = 0$, for all $t < t^*$). The most precise source receives full attention even when it has higher bias uncertainty. In *second stage* ($t \geq t^*$), source 2 is added to the information diet. In particular, for $t \geq t^*$, attention is allocated in fixed proportions:

$$\beta_1^*(t) = \frac{\frac{\sigma_1}{v_1^2}}{\frac{\sigma_1}{v_1^2} + \frac{\sigma_2}{v_2^2}} \quad \text{and} \quad \beta_2^*(t) = \frac{\frac{\sigma_2}{v_2^2}}{\frac{\sigma_1}{v_1^2} + \frac{\sigma_2}{v_2^2}}$$

From the definition of t^* , we see that a higher v_1 anticipates the arrival of the second stage, while a larger precision gap ($\sigma_2 - \sigma_1$) delays it.

Given any optimal stopping time τ^* , the more precise source is always learned first regardless of how uncertain the decision maker might be about its bias (see Figure 1). As attention to source 1 accrues, its marginal value decreases and eventually even the less precise receives positive attention. We show in Section 4 that β^* is the uniquely optimal learning strategy for all decision problems when two sources are available.

The trade-off between precision and bias uncertainty is most salient when source 2

is unbiased ($v_1 > v_2 = 0$, see Fig. 1(c)). Even when the bias uncertainty of source 1 is arbitrarily large, it is still always exclusively learnt first, before the decision maker starts to learn from the unbiased source. To see why the preference for precision initially dominates the preference for knowing the bias, consider a scenario in which the biased source is observed for the first dt time. Since b_1 and μ are independent according to the prior, this is equivalent to receiving a signal $s = \mu + b_1 + \epsilon$ where $\epsilon \sim \mathcal{N}(0, v_1^2 + \sigma_\mu^2 / dt)$. The instantaneous rate of reduction in posterior variance does not depend on v_1^2 :

$$\frac{\sigma_{\mu,0}^2 - \sigma_{\mu,dt}^2}{dt} = \frac{\sigma_\mu^4}{dt v_1^2 + \sigma_1^2 + dt \sigma_\mu^2} \rightarrow \frac{\sigma_\mu^4}{\sigma_1^2} \quad \text{as } dt \rightarrow 0$$

Prior uncertainty about the bias does not tamper with the informativeness of the biased source at the beginning, but gradually burdens learning in time.⁶ Thus, no matter how large is the initial uncertainty about the bias of source 1, the optimal attention allocation scheme always begins with full attention on it.

In light of the optimal learning strategy β^* , the stopping time determines whether the less precise source is demanded. With a Gaussian environment and quadratic loss preferences, the optimal stopping time τ^* is deterministic and strictly decreasing with respect to the flow cost c . This setting speaks to economic situations where the decision maker acquires information to reduce the riskiness of their decision. Since their learning strategy is deterministic, uncertainty about the state evolves in a predictable way, so a self-imposed learning deadline (independent of signal realisations) is optimal. In Section 6, we consider a more general class of utility functions and show that the optimal stopping time remains deterministic and decreasing in c .

In this section, we illustrated the key features of the optimal attention strategy. Precision is initially deemed more important than knowledge about the bias, and this trade-off evolves dynamically. This prompts the DM to begin acquiring information from the more precise source and eventually consider also the less precise one.

In the next section, we show that the features of the optimal attention strategy that we just presented remain true also in the general setting with N information sources, general

⁶The larger v_1^2 is, the shorter it takes for its effect on σ_μ^2 to dominate.

prior on μ , and arbitrary preferences.

4 Optimal Attention Allocation

In this section, we formally derive the optimal attention strategy in the general setting with N information sources, and generic preferences $u(a, \mu)$ and prior π_0^μ . Establishing the optimal attention strategy is not trivial since each strategy is differently confounded by biases and involves different rates of raw information arrival. In addition, there is feedback between β and the observed signal process X^β : β shapes the information structure, and the realised path X^β feeds back to the contingent choice of β .

To tackle these issues, we first establish a sufficient statistic of μ at time t given a generic allocation plan β , which consists of two components: (i) an **information score** $S^\beta(t)$ that encodes the aggregate signal on μ after rescaling and removing source-specific biases, and (ii) a (stochastic) **information clock** $\mathcal{I}^\beta(t)$ that captures the amount of information generated by β by time t , measured as the amount of time it would take to generate the same information when observing an unbiased continuous signal with unit variance.

For attention path β , define the **individual information score** at calendar time t as

$$S_i^\beta(t) := \frac{X_{it}^\beta - X_{i0}^\beta - b_{i0}q_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)} \quad (2)$$

The information score $S_i^\beta(t)$ is a normalised version of X_{it}^β , re-centred around μ and with standardised variance.

We also define the **individual information clock** $\mathcal{I}_i^\beta(t)$. The information clock measures the *effective cumulative attention* allocated to source i by calendar time t , accounting for the bias contamination and the different variance of each source. Taking the time derivative of the information clock, we also define $\iota_i^\beta(t)$, interpreted as the *effective instantaneous attention* allocated to source i at time t .

$$\mathcal{I}_i^\beta(t) := \frac{q_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)} \quad \text{and} \quad \iota_i^\beta(t) = \frac{\sigma_i^2 \beta_i(t)}{(\sigma_i^2 + v_i^2 q_i^\beta(t))^2} \quad (3)$$

Since $\beta(t)$ is adapted to $\mathcal{G}_t^\beta$, $q(t)$ will typically be random and path-dependent, and thus standard Itô calculus does not directly apply to the rescaled stochastic process $S_i^\beta(t)$. Nevertheless, with the next lemma we establish that the individual information score is well-behaved as it admits a semimartingale representation.

Lemma 1. *Fix a feasible β and a parameter $\mu = x$. Then, the individual information score $S_i^\beta(t)$ defined in (2) has the semimartingale decomposition*

$$dS_i^\beta(t) = x \iota_i^\beta(t) dt + dM_i^\beta(t), \quad (4)$$

where $M_i^\beta(t)$ are continuous $\mathcal{G}_t^\beta$ -local martingales that satisfy

$$d\langle M_i^\beta, M_j^\beta \rangle_t = \mathbb{1}_{\{i=j\}} \iota_i^\beta(t) dt. \quad (5)$$

Proof. See Appendix B.2. ■

Lemma 1 ensures that the transformed individual information score $S_i^\beta(t)$ is a well-defined stochastic differential equation with an adapted drift term and a martingale noise. The drift and quadratic variation of $S_i^\beta(t)$ share the same clock $\iota_i^\beta(t)$. Hence, after performing a time change from calendar time t to the effective clock $\iota_i^\beta(t)$, $S_i^\beta(t)$ is locally equivalent to a Brownian motion with drift x . Furthermore, since there is no cross-variation between the martingale noise of each source, we can easily aggregate the information scores and clocks across sources.

We define the **aggregate information score** $S^\beta(t)$ as

$$S^\beta(t) := \sum_{i=1}^N S_i^\beta(t) = \sum_{i=1}^N \frac{X_{it}^\beta - X_{i0}^\beta - b_{i0} q_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)} \quad (6)$$

The aggregate information score encodes the collective effective signal about μ based on information learnt so far across all sources. We can directly sum the effective signals of different sources as they are de-biased and contain independent martingale noise.

Similarly, we define the **aggregate information clock** $\mathcal{I}^\beta(t)$ as

$$\mathcal{I}^\beta(t) := \sum_{i=1}^N \mathcal{I}_i^\beta(t) = \sum_{i=1}^N \frac{q_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)} \quad (7)$$

The aggregate information clock measures the *total effective time* spent on all sources by time t , and can therefore be interpreted as a measure of total flow of information under β .

In the next lemma, we prove that $(S^\beta(t), \mathcal{I}^\beta(t))$ is a sufficient statistic for the posterior distribution of μ induced by β at time t . To state the lemma, let P_x^β denote the law of X_t^β induced by β when the true parameter is $\mu = x$.

Lemma 2 (Sufficient Statistic). *Fix a feasible β , a parameter $\mu = x$, and a time $t \leq T$, where $T < +\infty$ is a finite time horizon. Then, the likelihood ratio (relative to the reference state $\mu = 0$) conditional on $\mathcal{G}_t^\beta$ is*

$$L_t^\beta(x) := \frac{dP_x^\beta}{dP_0^\beta} \Big|_{\mathcal{G}_t^\beta} = \exp \left\{ x S^\beta(t) - \frac{x^2}{2} \mathcal{I}^\beta(t) \right\} \quad (8)$$

where $S^\beta(t)$ and $\mathcal{I}^\beta(t)$ are as defined in (6) and (7) respectively. Consequently, for every Borel prior π_0^μ , the marginal posterior of μ is

$$\pi_t^\mu(dx) = \frac{L_t^\beta(x) \pi_0^\mu(dx)}{\int_{\mathbb{R}} L_t^\beta(\tilde{x}) \pi_0^\mu(d\tilde{x})} \quad (9)$$

Proof. See Appendix B.2. ■

Since the bias of each source is unknown, we cannot directly apply Girsanov's theorem to the semimartingale decomposition of $S_i^\beta(t)$ from Lemma 1. Lemma 2 verifies that, despite the uncertainty about the biases, we can average out their effects and obtain a marginal likelihood for μ that depends only on the observed aggregate information score $S^\beta(t)$ and aggregate information clock $\mathcal{I}^\beta(t)$. This is because biases are *a posteriori* independently and normally distributed conditional on μ and the path history.

Example 1 (Section 3 Revisited). *When the prior belief on μ is Gaussian $\mu \sim \mathcal{N}(\mu_0, \sigma_\mu^2)$, the*

posterior belief at any t is also Gaussian. In particular, the posterior variance of μ is

$$\left[\frac{1}{\sigma_\mu^2} + \sum_{i=1}^n \frac{q_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)} \right]^{-1} = \left[\frac{1}{\sigma_\mu^2} + \mathcal{I}^\beta(t) \right]^{-1}$$

Therefore, with a Gaussian prior, minimising posterior variance is equivalent to maximising the aggregate information clock.

In the simplified setting in Section 3, the value of information is exactly its reduction in posterior variance of μ . In general, however, other moments of the posterior distribution may matter, but the aggregate information clock $\mathcal{I}^\beta(t)$ still captures the amount of total effective time spent learning. In what follows, we establish that the optimal learning strategy β^* should maximise $\mathcal{I}^\beta(t)$, regardless of preferences and prior beliefs.

As in Section 3, we first consider a static version of the optimal control problem. Given a fixed time horizon $\bar{t}$, we find the cumulative attention allocation $q(\bar{t})$ that maximises the information clock. Formally, we need to find the feasible $q(\bar{t})$ that solves

$$\max_{q(\bar{t})} \sum_{i=1}^N \frac{q_i(\bar{t})}{\sigma_i^2 + v_i^2 q_i(\bar{t})} \quad (10)$$

Analogous to Section 3, the optimal cumulative attention q^* consists of a series of learning stages. We define the deterministic **time thresholds** in time as $T_1 := 0$, $T_{N+1} := +\infty$, and

$$T_m := \sum_{k=1}^{m-1} \frac{\sigma_k(\sigma_m - \sigma_k)}{v_k^2}, \quad m = 2, \dots, N \quad (11)$$

Further, let **$\bar{t}$ -consumption set** be the set of information sources that receive positive attention by $\bar{t}$. Finally, let J denote the index of the **most precise unbiased source** available, and $N + 1$ if none.

$$J := \min(\{i \in \{1, \dots, N\} : v_i = 0\} \cup \{N + 1\}) \quad (12)$$

We are now ready to describe the optimal cumulative attention $q^*(\bar{t})$ that solves (10).

Proposition 1 (Optimal Cumulative Attention). *Fix a time horizon $\bar{t}$. Let $m \in \{1, \dots, N\}$*

be such that $\bar{t} \in [T_m, T_{m+1})$, where T_m is defined in (11). Let J be the index of the most precise unbiased source defined in (12).

Then, the unique cumulative attention allocation that solves (10) is

$$q_i^*(\bar{t}) = \begin{cases} \mathbf{1}_{\{i \leq m\}} \frac{\sigma_i}{v_i^2} \left(\frac{\bar{t} + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} - \sigma_i \right) & \text{if } m < J \\ \frac{\sigma_i(\sigma_j - \sigma_i)}{v_i^2} & \text{if } m \geq J \text{ and } J > i \\ \mathbf{1}_{\{i=J\}}(\bar{t} - T_J) & \text{if } m \geq J \text{ and } J \leq i \end{cases} \quad (13)$$

Furthermore, the $\bar{t}$ -consumption set is $\{1, \dots, \min\{m, J\}\}$.


Proof. See Appendix B.1. ■

Proposition 1 shows that the most informative cumulative attention path progressively incorporates sources in descending order of precision. First consider the case without an unbiased source, i.e., $J = N + 1$. In this case, when $\bar{t} \in [T_m, T_{m+1})$, the $\bar{t}$ -consumption set contains the m most precise sources $\{1, \dots, m\}$. As $\bar{t}$ increases, the consumption set expands by adding sources in decreasing order of precision (Fig. 2(a)(b)). Hence, the optimal strategy involves a sequence of **learning stages**, with the next most precise source added in each stage.

The expansion of the consumption set occurs at the deterministic time thresholds. The threshold T_{m+1} corresponds to the instant when the speed of information flow attained from the consumption set $\{1, \dots, m\}$ reaches that of the next best source ($m + 1$) and thus the moment source $m + 1$ begins to receive attention. This gradual expansion of the consumption set resembles the result in Liang et al. (2022), where the DM sequentially incorporates attributes according to their correlation structure. In our setting with state-independent biases, the order of expansion is entirely determined by the precision of sources, despite imposing fewer restrictions on the prior on μ . Consequently, both the consumption sets and the associated time thresholds admit closed-form solutions.

While the order in which sources are incorporated in the information diet is determined by their precision, bias uncertainty determines the length of each learning stage. A higher

Figure 2: Optimal Attention Allocation



Notes. This figure plots the cumulative attention allocated to each source (shaded area between the lines) as a function of time, when there are four available sources ($N = 4$). Each panel shows cases with different uncertainty about the biases.

(lower) bias uncertainty anticipates (delays) the arrival of subsequent learning stages, since it control how fast the value of additional attention on the source declines (Fig. 2(a)(b)).

When the information environment includes an unbiased source, $J \in \{1, \dots, N\}$, optimal learning eventually only devotes attention to source J . Before stage J , learning stages unfold exactly as described before. Once the unbiased source enters the consumption set, it fully captures all future attention (Fig. 2(c)). This happens because, contrary to sources with unknown bias, the marginal value of an unbiased source does not decrease with the total amount of attention it receives. Hence, once it starts being consumed, it remains the most valuable source.

We now turn to characterising the optimal attention allocation β^* . Since $q^*(t)$ is monotone, the instantaneous attention allocation $\beta^*(t) := \frac{dq^*(t)}{dt}$ maximises $\mathcal{I}^\beta(t)$ for all t . With general prior belief and preferences, unlike in the case in Section 3, the decision maker's expected payoff could depend on the entire posterior belief distribution of μ , which evolves stochastically. Nevertheless, Proposition 2 shows that β^* sequentially Blackwell dominates any other feasible attention allocation and thus establishes β^* as the optimal attention allocation scheme for generic prior belief on μ and general decision problems, allowing for endogenous stopping times.

Proposition 2 (Optimal Attention). *The universal optimal attention strategy that solves (1) for generic decision problem (A, u) and prior belief π_0^μ is*

$$\beta_i^*(t) = \begin{cases} 0 & \text{if } t < T_i \\ \frac{\sigma_i/v_i^2}{\sum_{k: T_k \leq t} \sigma_k/v_k^2} & \text{if } T_i \leq t < T_J \\ \mathbf{1}_{\{i=J\}} & \text{if } t \geq T_J \end{cases} \quad (14)$$

for all $i \in \{1, \dots, N\}$, where T_m and J are defined in (11) and (12) respectively.

Proof. See Appendix B.3. ■

Proposition 2 shows that β^* is sequentially Blackwell dominant, i.e., it is the unique attention allocation that solves the optimal control problem for any prior on μ and any decision problem. Within each learning stage, attention is allocated in fixed proportions across the **active sources**, i.e., sources that receive strictly positive attention at that time. As time passes, the marginal value of the set of active sources $\{1, \dots, m\}$ falls because bias uncertainty limits the long-run informativeness of repeated observations from the same set of sources. Eventually, the marginal value of $\{1, \dots, m\}$ falls to the *ex ante* marginal value of source $m + 1$: this occurs at the time threshold T_{m+1} , when a new learning stage begins and source $m + 1$ becomes active. By inducing q^* at all t , β^* inherits some of its key features: sources are activated in order of decreasing precision, bias uncertainty regulates the length of learning stages, and the most precise unbiased source absorbs all future attention once it enters the consumption set.

4.1 Proof Sketch of Proposition 2

To prove Proposition 2, we show that the sequential statistical experiment induced by β^* **sequential Blackwell dominates** that induced by any other (potentially random) attention allocation $\tilde{\beta}$. The notion of sequential Blackwell dominance is formally introduced in Appendix A. Sequential Blackwell dominance requires the **existence of a probability kernel** K that maps (garbles), for every state $\mu = x$, the path of observables $H^* := (\beta^*, X^{\beta^*})$ induced by β^* into the distribution over paths of observables $\tilde{H} := (\tilde{\beta}, X^{\tilde{\beta}})$ induced by $\tilde{\beta}$.

Moreover, the kernel must be **adapted**: at any time t , the distribution over time t -paths of observables $\tilde{H}_{\leq t}$ conditional on H^* depends only on $H_{\leq t}^*$. The first requirement is the standard notion of (static) Blackwell dominance. The second requirement, instead, ensures that any stopping time τ feasible for the dynamic experiment $\tilde{\beta}$ is also feasible for β^* .

The first step of the proof is to show that $\mathcal{I}^\beta(t)$ is continuous and strictly increasing in t for all feasible β . This is immediate from the feasibility of β and the definition of $\mathcal{I}^\beta(t)$. This allows us to perform a time change and express the process of the aggregate information score S^β in terms of effective time $\mathcal{I}^\beta(t)$ instead of calendar time t .

In the case of β^* , we introduce $\kappa^*(s) := (\mathcal{I}^{\beta^*})^{-1}(s)$ for $s \leq \lim_{t \rightarrow \infty} \mathcal{I}^{\beta^*}(t)$, which is well-defined since $\mathcal{I}^{\beta^*}$ is deterministic. Under time change and Lemma 1, the aggregate information score under β^* is $C_s^* := S^*(\kappa^*(s))$, which admits a drift-diffusion representation $C_s^* = xs + W_s'$ for $0 \leq s < \lim_{t \rightarrow \infty} \mathcal{I}^{\beta^*}(t)$, for any state $x \in \mathbb{R}$.

Analogously, for a generic and potentially random $\tilde{\beta}$, we define $\kappa^{\tilde{\beta}}(s) := \inf\{t \geq 0 : \mathcal{I}^{\tilde{\beta}}(t) > s\}$ for $s \leq \lim_{t \rightarrow \infty} \mathcal{I}^{\tilde{\beta}}(t)$. Then, the information score $\tilde{C}_s := S^{\tilde{\beta}}(\kappa^{\tilde{\beta}}(s))$ is also Brownian (with drift) up to the random information time $\lim_{t \rightarrow \infty} \mathcal{I}^{\tilde{\beta}}(t)$, by the Dambis–Dubins–Schwarz theorem. If $\lim_{t \rightarrow \infty} \mathcal{I}^{\tilde{\beta}}(t) < \lim_{t \rightarrow \infty} \mathcal{I}^{\beta^*}(t)$, we fill in the path with a standard and independent Brownian motion for $\lim_{t \rightarrow \infty} \mathcal{I}^{\tilde{\beta}}(t) \leq s < \lim_{t \rightarrow \infty} \mathcal{I}^{\beta^*}(t)$. Note that $\tilde{C}_{\mathcal{I}^{\tilde{\beta}}(t)}$ is adapted to $\mathcal{G}_{\mathcal{I}^{\tilde{\beta}}(t)}^{\tilde{\beta}} \subseteq \mathcal{G}_{\mathcal{I}^{\beta^*}(t)}^{\beta^*}$, since $\mathcal{I}^{\tilde{\beta}}(t) \leq \mathcal{I}^{\beta^*}(t)$ for every t thanks to Proposition 1.

Fix the state $\mu = 0$ so that the paths of the aggregate information scores are Brownian without trends. We show that there exists a state-independent adapted kernel $\bar{K}(c^*, \cdot)$ such that: (i) given the aggregate information score path $c^* = (c_s^*)_{s \geq 0}$, $\bar{K}(c^*, \cdot)$ reproduces the distribution over $\tilde{c}$; and (ii) for any calendar time t , its restriction $\bar{K}_t$ reproduces the distribution over $\tilde{c}_{\leq \mathcal{I}(t)}$ conditionally on c^* using only information contained in $c_{\leq \mathcal{I}^*(t)}^*$. Such a kernel exists because the aggregate information score up to effective time $\mathcal{I}^*(t)$ contains weakly more information than the aggregate information score up to effective time $\tilde{\mathcal{I}}(t)$, and by Lemma 2 this is the sufficient statistic to replicate the distribution of signal paths up to calendar time t . Since any observable path h^{β^*} generated by the dynamic experiment β^* can also be expressed as a score path c^* , we can define the garbling kernel $K(h^{\beta^*}, \cdot) := \bar{K}(c^*(h^{\beta^*}), \cdot)$ that maps any path of observables induced by β^* into a

distribution of observable paths under $\tilde{\beta}$. Finally, we use Lemma 2 and likelihood tilting to show that the garbling kernel we found under $\mu = 0$ applies for every state, which implies that β^* sequentially Blackwell dominates $\tilde{\beta}$.

Intuitively, we can convert the information process of any attention strategy β into the shape of standard Brownian motions running on an effective information timescale, under the reference state $\mu = 0$. This reformulation allows us to identify β^* as the strategy that generates the most information for any horizon, which translates to the fastest accumulation of information under calendar time.

5 Comparative Statics on Optimal Attention Allocation

5.1 Changing Existing Sources

While our model focuses on the demand for information, Proposition 1 offers valuable insights for the information supplier who designs source characteristics. An expert may wonder about the effect of revealing that the product review is sponsored. A news editor may be interested in how a change in the platform’s writing quality affects readership. A tech company may want to know whether being more opaque about their existing AI model’s bias could boost or shrink usage. In this section, we explore the comparative statics on cumulative attention when we vary the precision or the bias uncertainty of available sources.

Proposition 3 (Comparative Statics on Bias Uncertainty). *Fix $t \geq 0$ and let $i, j \in \{1, \dots, m\}$ with $i \neq j$, where m is such that $t \in (T_m, T_{m+1})$. Then, the source i ’s optimal cumulative attention $q_i^*(t)$ is decreasing in its own bias uncertainty v_i and increasing in a rival source’s bias uncertainty v_j .*

Proof. See Appendix C.1. ■

Proposition 3 shows that increasing bias uncertainty unambiguously reduces the attention on the source, and diverts it to other sources. Figure 3 illustrates these effects by plotting how cumulative attention varies after increasing v_2 (top panels) or v_1 (bottom

panels). First, consider the effect that increasing v_2 has on the consumption of source 2. By comparing panel (b) with panel (a), we see that increasing the bias uncertainty of source 2 does not affect the time at which the source starts receiving attention, that is, $\frac{\partial T_i}{\partial v_i} = 0$. However, once source 2 enters the consumption set, a higher bias uncertainty reduces the source's relative weight in the information diet, that is, $\frac{\partial \beta_i^*(t)}{\partial v_i} \leq 0$. Since source 2 is consumed starting from the same point in time but at lower intensity, total attention is diverted to other sources (see panel (c)).

Second, consider the effect that increasing v_1 has on the consumption of source 2. Contrasting panel (e) with panel (d), we observe that an increase in bias uncertainty of source 1 anticipates the consumption of source 2, that is, $\frac{\partial T_i}{\partial v_j} \leq 0$ for $j < i$. Moreover, as the higher bias uncertainty of source 1 decreases its relative weight in the information diet, source 2's relative weight increases, that is, $\frac{\partial \beta_i^*(t)}{\partial v_j} \geq 0$. The combination of these two effects increases the total attention devoted to source 2, as shown in panel (f).

The effects of changes in precision are instead different.


Proposition 4 (Comparative Statics on Precision). *Fix $t \geq 0$ and let $i, j \in \{1, \dots, m\}$ with $i \neq j$, where m is such that $t \in (T_m, T_{m+1})$. Then, there exists a precision threshold $\hat{\sigma}_i(t) > 0$ (strictly increasing in t) such that:*

- *the source i 's optimal cumulative attention $q_i^*(t)$ is decreasing in its own variance σ_i if and only if $\sigma_i \geq \hat{\sigma}_i(t)$;*
- *the source j 's optimal cumulative attention $q_j^*(t)$ is increasing in source i 's variance σ_i if and only if $\sigma_i \geq \hat{\sigma}_i(t)$.*

Proof. See Appendix C.1. ■

Proposition 4 shows that when the precision of a source is below a given threshold ($\sigma_i < \hat{\sigma}_i$), making the source less precise will actually attract attention from (rather than divert it towards) other sources. Moreover, this threshold is increasing in the time horizon t . The non-monotone effect of precision on cumulative attention arises from trading-off the extensive versus the intensive margin of learning. On the extensive margin, a less precise source is optimally delayed, that is, $\frac{\partial T_i}{\partial \sigma_i} > 0$. This effect is visualised in Figure 4, where an

Figure 3: Varying Bias Uncertainty




Notes. This figure visualises the comparative statics result when varying bias uncertainty. In the left and centre columns, the figure plots the cumulative attention allocated to each source (shaded area between the lines) as a function of time, when there are three available sources ($N = 3$). Panel (b) uses the same parameters as panel (a), but increases v_2 ; similarly, panel (e) uses the same parameters as panel (d), but increases v_1 . Panel (c) shows the changes in $q^*(t)$ after increasing v_2 (i.e., moving from panel (a) to panel (b)), and panel (f) shows the changes in $q^*(t)$ after increasing v_1 (i.e., moving from panel (d) to panel (e)).

increase in σ_2 delays the time at which source 2 enters the information diet. This effect is counteracted by a higher intensity of learning once the source enters the consumption set, that is, $\frac{\partial \beta_i^*(t)}{\partial \sigma_i} \geq 0$. Again, Figure 4 shows that after the increase in σ_2 source 2 captures more attention. These two opposing forces (delay and higher intensity) create a *back-loading effect*: the source is consumed more intensely but only starting from a later point in time. When the source is already sufficiently precise ($\sigma_i < \hat{\sigma}_i(t)$),⁷ the intensity effect dominates the delay effect, so that a marginal reduction in precision actually attracts attention to the

⁷Or, equivalently, when time is large enough: $t > \hat{\sigma}_i^{-1}(\sigma_i)$.

Figure 4: Varying Precision



Notes. This figure visualises the comparative statics result when varying precision. In the left and centre panels, the figure plots the cumulative attention allocated to each source (shaded area between the lines) as a function of time, when there are three available sources ($N = 3$). Panel (b) uses the same parameters as panel (a), but increases σ_2 . Panel (c) shows the changes in $q^*(t)$ after increasing σ_2 (i.e., moving from panel (a) to panel (b)).

source (see panel (c)).

Beyond describing comparative statics results, this section highlights a key point of the paper: uncertainty about the signal (precision) and uncertainty about the information structure (bias uncertainty) are distinct concepts that shape learning differently. It might be tempting to think about bias uncertainty as some form of “extra noise”, yet the implications of varying σ and v are remarkably different and potentially move in opposite directions.

5.2 Introducing New Sources

Another economically relevant question is about how the demand for information changes when new sources are introduced in the information environment, such as a novel news platform or an additional expert opinion. More specifically, suppose a new information source X' with variance σ' and bias uncertainty v' is added to $(X_i)_{i=1}^N$. Order the $N + 1$ sources in increasing variance and let r denote the index of the added source. Let the updated time thresholds be $(T'_1, \dots, T'_{r-1}, \zeta', T'_r, \dots, T'_N)$,⁸ and the updated optimal

⁸We slightly change the notation here to keep the comparison clean. The r -th source in $(X_i)_{i=1}^N$ is the $r+1$ -th source in the new order, with original time threshold T_r and updated time threshold T'_r . The time threshold

attention allocation be $\beta^{*'}(t)$. Moreover, let $\delta_i := T'_i - T_i$ denote the consumption delay of source i caused by X' . The next result shows how the added information source affects the attention allocation and consumption delay of pre-existing sources.

Corollary 1. *Let β^* and $(T_i)_{i=1}^N$ denote the optimal attention allocation and time thresholds when the available information sources are $(X_i)_{i=1}^N$. Let $\beta^{*'} and $(T'_1, \dots, T'_{r-1}, \zeta', T'_r, \dots, T'_N)$ denote the optimal attention allocation and time thresholds when the available information sources are $(X_i)_{i=1}^N \cup X'$. Let $\delta_i := T'_i - T_i$ denote the consumption delay of an original source i . Then:$*

- *In earlier stages ($i < r$), optimal attention and time thresholds do not change: $T'_i = T_i$ and $\beta^{*'}(t) = \beta^*(t)$ for $t \leq \zeta'$;*
- *In later stages ($i \geq r$), consumption of original sources is delayed but the relative attention ratio among original sources is unaffected: $\delta_i = \frac{\sigma'(\sigma_i - \sigma')}{(\sigma')^2} > 0$, and $\frac{\beta_i^{*'}(t)}{\beta_j^{*'}(t)} = \frac{\beta_i^*(t)}{\beta_j^*(t)}$ for $i, j \in \{1, \dots, N\}$ if $t > \max\{T'_i, T'_j\}$.*

Moreover, the consumption delay of an original source $i \geq r$ is increasing in the variance of the original source σ_i , decreasing in the bias uncertainty of the new source σ' , and decreasing in the variance of the new source σ' if and only if $\sigma_i < 2\sigma'$.


Proof. Direct computation from Proposition 2 yields the result. ■

Corollary 1 shows that adding a new source has both local and global effects on optimal learning. These are represented graphically in Figure 5. Since the learning stages are ordered by precision, the new source is demanded only after a time threshold, before which the optimal learning strategy is unchanged. The new source creates an additional stage, where the source starts being consumed together with the more precise sources. Additionally, the presence of the new source delays the consumption of other sources in all subsequent stages. Specifically, new sources with lower bias uncertainty more strongly delay the consumption of noisier sources. Furthermore, new sources with higher precision delay the consumption of *slightly noisier* sources more strongly than the consumption of *much noisier* sources.

One special case of Corollary 1 is the addition of an unbiased source, such as a credibly

for the added source X' is ζ' .

Figure 5: Adding a Source



Notes. This figure visualises the comparative statics result when adding a source. Panel (a) plots the cumulative attention allocated to each source (shaded area between the lines) as a function of time, when there are three available sources ($N = 3$) with $\sigma_1 < \sigma_3 < \sigma_4$. Panel (b) plots the cumulative attention after adding a new source (source 2) such that $\sigma_2 \in (\sigma_1, \sigma_3)$. Panel (c) shows the changes in $q^*(t)$ after adding source 2 (i.e., moving from panel (a) to panel (b)).

neutral third-party report. The presence of the unbiased source in the information environment completely blocks the demand for any biased source with lower precision. However, such unbiased source does not alter the duration or consumption ratio of learning stages that are higher in the pecking order. Our model thus illustrates the limited effectiveness of unbiased information interventions.

5.3 Learning from Undominated Sources

Proposition 2 reveals the dynamic value of *ex ante* dominated sources. We say that source j is *dominated* by source i if it has lower precision and a more uncertain bias ($\sigma_i \leq \sigma_j$ and $v_i \leq v_j$). A dominated source can receive strictly positive attention even when the dominating source is available. In this sense, adding a dominated source to the set of available information sources might yield a strict improvement as long as all other information sources have an unknown bias. In fact, when the time cost of acquiring information is sufficiently low, the marginal value of the dominating source will eventually fall below that of the dominated one. In this scenario, both the dominating and the dominated source are eventually consumed.

Next, we consider information environments in which no source is dominated. We say that a set of information sources $(X_i)_{i=1}^N$ is **mutually undominated** if there exists no pair $i, j \in \{1, \dots, N\}$ such that $\sigma_i \leq \sigma_j$ and $v_i \leq v_j$. Given that we ordered sources in decreasing order of precision, a mutually undominated set is characterised by $v_1 > v_2 > \dots > v_N$. This setting, where sources are ordered in decreasing precision *and* increasing bias uncertainty, is particularly useful to highlight the key trade-off between precision and bias uncertainty. The next result shows that, when the set of available information sources is mutually undominated, the optimal learning strategy exhibits a recency effect.

Corollary 2. *Suppose the set of available information sources $(X_i)_{i=1}^N$ is mutually undominated. Suppose $v_i > 0$ for all $i \in \{1, \dots, N\}$. If $t \in [T_m, T_{m+1})$, then $\beta_{i+1}^*(t) > \beta_i^*(t)$ for all $i \in \{1, \dots, m-1\}$.*

Proof. See Appendix C.2. ■

As sources enter in descending precision and ascending bias uncertainty, sources added to the consumption set more recently capture more attention. This apparently counterintuitive result arises from the fact that source i only enters the consumption set when the informativeness of more precise sources (with $j < i$) drops to i 's level.

Since fresher sources are given more relative attention, the total consumption of a noisier source may catch up with the earlier sources. The next result shows that this always occurs if the learning period is long enough.

Corollary 3. *Suppose the set of available information sources $(X_i)_{i=1}^N$ is mutually undominated. Suppose $v_i > 0$ for all $i \in \{1, \dots, N\}$. Then, for $i < j$, there exists a threshold $\tilde{t}_{ij} > 0$ such that*

$$q_i^*(t) < q_j^*(\tau^*) \iff t > \tilde{t}_{ij}$$

That is, the source i 's optimal cumulative attention $q_i^*(t)$ at time t falls below that of source j if and only if time t is beyond a threshold $\tilde{t}_{ij} > 0$.

Proof. See Appendix C.2. ■

To interpret this result, consider the situation where the consumer is patient enough, so that learning time passes the threshold $\tilde{t}_{ij}$. In this case, the consumer will optimally spend

more time consulting product reviews of lower quality and better understood sponsorship, rather than high quality reviews with unclear sponsorship. Similarly, a news platform with lower reporting quality but more predictable bias attracts higher total readership than a news platform with higher quality but more uncertain slant.

6 Optimal Stopping Time

Having characterised the optimal attention strategy for all decision problems and prior beliefs, we now turn to optimal stopping. Throughout this section, we restrict attention to the setting where the prior π_0^μ is Gaussian with $\mu \sim \mathcal{N}(\mu_0, \sigma_\mu^2)$, ensuring that the posterior variance evolves deterministically given any attention strategy. As we have already discussed in Section 3, with a quadratic loss utility function, the optimal stopping time is deterministic and decreasing in the flow cost. In general, however, the choice of stopping might be contingent on the signal realisations.

In the next result, we show that the stopping time is deterministic and decreasing in the flow cost for a subclass of utility functions. Let $\mathcal{U}$ be the set of utility functions of the form $u(a, \mu) = f(a - \mu) + \xi$, where $\xi \in \mathbb{R}$ and $f : \mathbb{R} \rightarrow \mathbb{R}$ is an integrable (i.e., $\mathbb{E}_{\pi_0^\mu}[|f|] < +\infty$) upward-bounded function with well-defined maximum. The set $\mathcal{U}$ includes preferences that depend on the state-action difference, such as the quadratic loss utility function discussed in Section 3.

Proposition 5. *Let $u \in \mathcal{U}$. The earliest optimal stopping time is deterministic and decreasing with respect to c .*

Proof. See Appendix C.3. ■

Analogous to the intuition laid out in Section 3, for any $u \in \mathcal{U}$ and under Gaussian beliefs, the gross value upon stopping is a decreasing function of the posterior variance and does not depend on the posterior mean. The fact that the optimal stopping time is decreasing in the flow cost has implications regarding the set of sources that are consumed. In particular, since the optimal attention β^* involves expanding consumption sets, more impatience leads to earlier stopping and therefore to consuming weakly fewer sources.

7 Applications to Special Prior Distributions

In our model, the payoff-relevant parameter μ can be learnt only by allocating infinite attention to an unbiased source. If, instead, all available information sources have an unknown (independent) bias, μ cannot be fully learnt. In practice, this is often the case when individuals face information providers with unknown and independent private interests, such as sponsored reviews, news platforms, and social media contents. As we have already discussed, the amount of long-run information the decision maker can collect is limited in this context. We can therefore think of decomposing μ into a learnable component (i.e., what can theoretically be learnt about μ if allocating infinite attention to each available source), and a residual component (capturing all the noise induced by the unknown biases that cannot be removed).

In this section, we illustrate how the optimal control problem can be reframed in terms of the learnable component of μ and discuss two special cases. When the prior on μ is Gaussian, the learnable component is simply a linear combination of the rescaled source drifts, which admits an equivalent representation as an information aggregation problem in the framework of [Liang et al. \(2022\)](#). When instead the prior follows a Bernoulli distribution, the learnable component is the inverse-logit of a linear combination of drifts. The discussion of these two cases illustrates how the optimal attention allocation scheme from Proposition 2 applies to a wide range of prior beliefs with potentially different state space's structures.

For ease of exposition, suppose $v_i > 0$ for all $i \in \{1, \dots, N\}$ throughout this section. We start by re-writing source i 's information technology. First, we define the **normalised innovation** of source i as follows:

$$\theta_i := \frac{(\mu - \mathbb{E}_0[\mu]) + (b_i - b_{i0})}{\sigma_i} \quad (15)$$

Then, observing the X^β process is equivalent to observing the Y^β process defined by

$$dY_{it}^\beta = \theta_i \beta_i(t) dt + \sqrt{\beta_i(t)} dW_{it} \quad (16)$$

The maximal amount of information that can be obtained in this environment corresponds to knowing $\theta = (\theta_i)_{i=1}^N$. Hence, we define the **learnable component** of μ (denoted ω) to be the best estimate of μ when θ is known, formally, $\omega := \mathbb{E}[\mu - \mathbb{E}_0[\mu]|\theta]$.

7.1 Learning with Gaussian Prior

Suppose the prior on the underlying state is normally distributed $\mu \sim \mathcal{N}(\mu_0, \sigma_\mu^2)$. In this special case of Gaussian beliefs and Gaussian signals, the posterior variance on μ reduces deterministically and it is bounded from below by the limiting posterior variance $\varphi_\infty := \left(\frac{1}{\sigma_\mu^2} + \sum_{i=1}^N \frac{1}{v_i^2} \right)^{-1}$ when θ is fully learnt.

Lemma 3. Let $\alpha_i := \varphi_\infty \frac{\sigma_i}{v_i^2}$. Then, $\omega = \sum_{i=1}^N \alpha_i \theta_i$. Moreover, μ can be decomposed into the learnable component ω and residual uncertainty ε , with

$$\mu = \mu_0 + \omega + \varepsilon, \quad \varepsilon \sim N(0, \varphi_\infty), \quad \varepsilon \perp \theta. \quad (17)$$

Proof. See Appendix C.4. ■

Lemma 3 illustrates that, under the Gaussian environment, uncertainty on μ can be broken down into two orthogonal components. Uncertainty on the learnable component ω can be reduced by learning from available information sources, whereas the orthogonal irreducible noise ε , arising from source-specific biases, cannot be resolved by any β .

Additionally, Lemma 3 shows that the payoff-relevant learnable component ω is a linear combination of the learnable factors θ_i . Hence, learning μ from idiosyncratically biased sources $(X_i)_{i=1}^N$ is analogous to learning the projected ω from the normalised sources $(Y_i)_{i=1}^N$, using θ_i as attributes and α_i as attributes' weights. In the next lemma, we establish this formal link with the dynamic information aggregation problem in Liang et al. (2022) (henceforth, LMS). Define the residual-averaged utility as

$$\bar{u}(a, \omega) := \mathbb{E}_\varepsilon[u(a, \mu_0 + \omega + \varepsilon)], \quad \varepsilon \sim N(0, \varphi_\infty), \quad (18)$$

Here, $\bar{u}(a, \omega)$ integrates out the irreducible Gaussian noise that is unchanged by the choice

of learning. Hence, an attention strategy that maximises the expected value of the original utility $u(a, \mu)$ equivalently maximises the expected value of $\bar{u}(a, \omega)$.

Lemma 4. *Under Gaussian prior $\mu \sim \mathcal{N}(\mu_0, \sigma_\mu^2)$, the optimal control problem of (1) is equivalent to the LMS problem with attributes θ_i defined in (15), weights $\alpha_i := \varphi_\infty \frac{\sigma_i}{v_i^2}$, information sources defined in (16), payoff state $\omega = \sum_{i=1}^N \alpha_i \theta_i$, and payoff $\tilde{u}(t, a, \omega) := \bar{u}(a, \omega) - ct$. For every corresponding admissible plan (β, τ, a_τ) ,*

$$\mathbb{E}[u(a_\tau, \mu) - c\tau] = \mathbb{E}[\tilde{u}(\tau, a_\tau, \omega)].$$

Hence, the two problems have the same value and the same optimal attention plans.

Proof. See Appendix C.4. ■

Lemma 4 illustrates that when the prior is Gaussian, the biased learning problem can be transformed into an information aggregation problem concerning the learnable component of the original state μ , since the expected value upon stopping is always the same.

7.2 Learning with Bernoulli Prior

Suppose the underlying parameter space is binary with support $\{h, l\}$, and the prior is $p_0 = P(\mu = h)$. Due to the binary support, the posterior log-odds (rather than the posterior mean in the Gaussian case) are a linear combination of drifts θ . This setting cannot be reframed as a linear information aggregation problem in the LMS framework since ω is a non-linear function of attributes.

Lemma 5. *Let the prior be $P(\mu = h) = p_0$. Then,*

$$\omega = (h - l) \left[\text{logit}^{-1} \left(l_0 + (h - l) \sum_{i=1}^N \frac{\sigma_i}{v_i^2} \theta_i \right) - p_0 \right] \quad (19)$$

where $l_0 := \log(\frac{p_0}{1-p_0})$.

Proof. See Appendix C.5. ■

The optimal attention strategy β^* in Proposition 2 does not depend on parameters of the DM's prior on μ such as the values of h and l . Therefore, the most precise source is prioritised even if its bias uncertainty significantly exceeds the range of possible parameter values. This is because precision has first-order importance on the marginal value of information from a source when it first receives attention. We conclude by noting that while β^* is not affected by the prior, the optimal stopping time τ^* might be.⁹ Consequently, the *realised* attention allocation upon stopping might vary with the prior.

8 Discussion

In this paper, we obtain tractable solutions for a rich dynamic information acquisition problem with biased information sources. Our results apply to arbitrary prior beliefs on μ and payoff functions because the optimal information structure can be standardised and ranked under an adaptive-Blackwell order. When the bias of each information source is independent and Gaussian, the aggregate information flow $\mathcal{I}^\beta$ is separable in the cumulative attention of each source, which renders a clean closed-form optimal β^* . We conclude by discussing a few natural variations.

Time Cost. Our model features a linear cost of time. However, the optimality of β^* does not directly depend on the functional form of the cost of time. In fact, the essential assumption is that the decision maker is impatient for action, that is, conditionally on having the same stopping belief they prefer stopping earlier than later. Under such condition, β^* is still adapted-Blackwell dominant and hence it is the unique solution to the attention allocation problem. In other words, all of our main results extend to models with arbitrary non-decreasing waiting costs $c(t)$, as well as to models with geometrically discounted payoffs of the kind $e^{-rt}u(a, \mu)$, where $u(a, \mu) \geq 0$.

Non-Expected Utility. In our model, the criterion that is maximised is a (subjective) expected utility. As the theoretical literature on ambiguity suggests, however, there might be uncertainty about beliefs, as well as non-neutral attitudes towards this type of

⁹In the case of binary-action binary-state decision problems, standard results show that the optimal stopping time will be characterised by two deterministic belief thresholds. One can then compute the expected stopping time and analyse the expected demand of each source.

uncertainty. Typically, uncertainty about beliefs induces actions that differ from those of a subjective expected utility maximiser. Nevertheless, our main results show that β^* is the optimal attention allocation regardless of the prior on the payoff-relevant state μ . In other words, even if multiple priors are entertained as possible, β^* maximises the criterion *pointwise* (at each prior π_0^μ). Hence, our results immediately extend to settings where the decision maker perceives uncertainty about π_0^μ , regardless of the attitudes towards this type of uncertainty.

Correlated Biases. In some settings, biases of certain information sources might be perceived as correlated. For example, news platforms under the same parent company are likely to share similar partisan biases, and LLM models trained on similar datasets may output correlated recommendations. Suppose that the prior on the bias is Gaussian with covariance matrix Σ_b , which is diagonal in the case of independent biases. For general Σ_b , the aggregate information flow can still be computed in closed-form, but is not separable in q_i as in (7). Hence, the marginal benefit of one source will depend on the cumulative attention on all the correlated sources. When Σ_b satisfies certain conditions¹⁰, the cumulative attention that maximises the aggregate information flow is still monotone in t , so an adapted-Blackwell dominant attention allocation exists.

Alternative Signal Structure. In our model, the match between the smooth Brownian information structure and the Gaussian bias uncertainty is designed to capture the confounding effect of signal noise (σ_i) and higher-order uncertainty (v_i) in learning. The presence of bias uncertainty (i) dilutes the effective learning on μ , and (ii) imposes a ceiling on the maximum value of learning from source i . If instead a biased signal ($\mu + b_i$) arrives conclusively via a constant Poisson arrival rate, bias uncertainty only limits the informativeness of the fully revealing signal, but does not slow down the speed of learning. Such alternative signal structure thus will not feature the dynamic bias-variance tension of our model. In this sense, our results do not straightforwardly extend to settings with alternative signal structures.

¹⁰See Assumption 4-6 in [Liang et al. \(2022\)](#) for sufficient conditions.

References

- Acemoglu, D., Chernozhukov, V. & Woldz, M. (2016), 'Fragility of asymptotic agreement under bayesian learning', *Theoretical Economics* **11**(1), 187–225.
- Ambuehl, S. & Li, S. (2018), 'Belief updating and the demand for information', *Games and Economic Behavior* **109**(C), 21–39.
URL: <https://ideas.repec.org/a/eee/gamebe/v109y2018icp21-39.html>
- Anagol, S., Cole, S. & Sarkar, S. (2017), 'Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market', *The Review of Economics and Statistics* **99**(1), 1–15.
URL: <https://ideas.repec.org/a/tpr/restat/v99y2017i1p1-15.html>
- Aromí, J. D. (2012), 'Learning under limited attention and unknown precision of information', *IIEP-UBA-Conicet.. UdeSA* .
- Arrow, K. J., Blackwell, D. & Girshick, M. A. (1949), 'Bayes and minimax solutions of sequential decision problems', *Econometrica, Journal of the Econometric Society* pp. 213–244.
- Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B. & Meyer, S. (2012), 'Is Unbiased Financial Advice to Retail Investors Sufficient? Answers from a Large Field Study', *The Review of Financial Studies* **25**(4), 975–1032.
URL: <https://doi.org/10.1093/rfs/hhr127>
- Burke, J. (2008), 'Primetime spin: Media bias and belief confirming information', *Journal of Economics & Management Strategy* **17**(3), 633–665.
- Calvert, R. L. (1985), 'The value of biased information: A rational choice model of political advice', *The Journal of Politics* **47**(2), 530–555.
- Chalmers, J. & Reuter, J. (2020), 'Is conflicted investment advice better than no advice?', *Journal of Financial Economics* **138**(2), 366–387.
URL: <https://www.sciencedirect.com/science/article/pii/S0304405X20301537>

- Charness, G., Oprea, R. & Yuksel, S. (2021), 'How do People Choose Between Biased Information Sources? Evidence from a Laboratory Experiment', *Journal of the European Economic Association* **19**(3), 1656–1691.
URL: <https://doi.org/10.1093/jeea/jvaa051>
- Che, Y.-K. & Mierendorff, K. (2019), 'Optimal dynamic allocation of attention', *American Economic Review* pp. 2993–3029.
- Chen, H. & Suen, W. (2023), 'Competition for attention and news quality', *American Economic Journal: Microeconomics* **15**(3), 1–32.
- Chen, J. Y. (2022), 'Biased learning under ambiguous information', *Journal of Economic Theory* **203**, 105492.
- Chopra, F., Haaland, I., Roeben, F., Roth, C. & Sticher, V. (2025), 'News Customization with AI'.
URL: <https://papers.ssrn.com/abstract=5502639>
- Chopra, F., Haaland, I. & Roth, C. (2024), 'The Demand for News: Accuracy Concerns Versus Belief Confirmation Motives'.
URL: <https://papers.ssrn.com/abstract=4797594>
- de Oliveira, H. (2018), 'Blackwell's informativeness theorem using diagrams', *Games and Economic Behavior* **109**, 126–131.
- de Oliveira, H. & Lamba, R. (2025), 'Rationalizing dynamic choices'.
- Epstein, L. G. & Schneider, M. (2007), 'Learning under ambiguity', *The Review of Economic Studies* **74**(4), 1275–1303.
- Epstein, L. G. & Schneider, M. (2008), 'Ambiguity, information quality, and asset pricing', *The Journal of Finance* **63**(1), 197–228.
- Fudenberg, D., Strack, P. & Strzalecki, T. (2018), 'Speed, accuracy, and the optimal timing of choices', *American Economic Review* **108**(12), 3651–3684.

- Gentzkow, M. & Shapiro, J. (2006), 'Media Bias and Reputation', *Journal of Political Economy* **114**(2), 280–316.
URL: <https://www.journals.uchicago.edu/doi/abs/10.1086/499414>
- Gentzkow, M., Shapiro, J. M. & Stone, D. F. (2015), Media bias in the marketplace: Theory, in 'Handbook of media economics', Vol. 1, Elsevier, pp. 623–645.
- Gentzkow, M., Wong, M. B. & Zhang, A. T. (2025), 'Ideological bias and trust in information sources', *American Economic Journal: Microeconomics* **17**(2), 162–213.
- Gonçalves, D. (2024), 'Speed, accuracy, and complexity', *arXiv preprint arXiv:2403.11240* .
- Hébert, B. & Woodford, M. (2018), 'Information costs and sequential information sampling', *National Bureau of Economic Research* .
- Hoffman, M. (2016), 'How is Information Valued? Evidence from Framed Field Experiments', *The Economic Journal* **126**(595), 1884–1911.
URL: <https://doi.org/10.1111/eoj.12401>
- Hu, L., Li, A. & Tan, X. (2021), 'Rationally inattentive echo chambers', *arXiv preprint arXiv:2104.10657* .
- Liang, A. & Mu, X. (2020), 'Complementary information and learning traps', *Quarterly Journal of Economics* pp. 389–448.
- Liang, A., Mu, X. & Skyrgkanis, V. (2018), 'Optimal and myopic information acquisition', *Proceedings of the 2018 ACM Conference on Economics and Computation* .
- Liang, A., Mu, X. & Syrgkanis, V. (2022), 'Dynamically aggregating diverse information', *Econometrica* **90**(1), 47–90.
- Liang, Y. (2025), 'Learning from unknown information sources', *Management Science* **71**(5), 3873–3890.
- Mayskaya, T. (2022), 'Dynamic choice of information sources', *Available at SSRN* 4138333 .

- Moscarini, G. & Smith, L. (2001), 'The optimal level of experimentation', *Econometrica* **69**(6), 1629–1644.
- Mullainathan, S. & Shleifer, A. (2005), 'The market for news', *American economic review* **95**(4), 1031–1053.
- Sethi, R. & Yildiz, M. (2016), 'Communication with unknown perspectives', *Econometrica* **84**(6), 2029–2069.
- Shishkin, D. & Ortoleva, P. (2023), 'Ambiguous information and dilation: An experiment', *Journal of Economic Theory* **208**, 105610.
URL: <https://www.sciencedirect.com/science/article/pii/S0022053123000066>
- Sims, C. A. (2003), 'Implications of rational inattention', *Journal of monetary Economics* **50**(3), 665–690.
- Sobbrio, F. (2014), 'Citizen-editors' endogenous information acquisition and news accuracy', *Journal of Public Economics* **113**, 43–53.
- Steiner, J., Stewart, C. & Matějka, F. (2017), 'Rational inattention dynamics: Inertia and delay in decision-making', *Econometrica* **85**(2), 521–553.
- Suen, W. (2004), 'The self-perpetuation of biased beliefs', *The Economic Journal* **114**(495), 377–396.
- Tang, R. (2020), 'A theory of updating ambiguous information', *arXiv preprint arXiv:2012.13650*.
- Wald, A. (1947), 'Foundations of a general theory of sequential decision functions', *Econometrica* **15**(4), 279–313.
- Whitmeyer, M. & Williams, C. (2024), 'Comparisons of Sequential Experiments for Additively Separable Problems'. arXiv:2405.13709 [econ.TH].
- Zhong, W. (2022), 'Optimal dynamic information acquisition', *Econometrica* **90**(4), 1537–1582.

A Technical Setups

First we formally construct the underlying probability space. Fix a time horizon $T \leq [0, \infty]$, let $C([0, T], \mathbb{R}^N)$ be the space of continuous N -dimensional paths starting at zero. Then we define the state space to be

$$\Omega := \mathbb{R} \times \mathbb{R}^N \times [0, 1] \times C([0, T], \mathbb{R}^N), \quad (20)$$

such that for a state $\omega \in \Omega$ we have

$$\omega := (\mu, b, \zeta, w). \quad (21)$$

Here μ is the payoff-relevant state with $x \in \mathbb{R}$ being it's realisation, b is the vector of biases for each information source, $\zeta \sim \mathcal{U}[0, 1]$ is an independent, uniformly distributed randomisation device and w is the vector of realised Brownian path for $(W_{1t}, W_{2t}, \dots, W_{Nt})$. Let $\mathbf{W}_0$ be the Wiener measure on $C([0, \infty], \mathbb{R})$. Fix a prior belief $\pi_0 = \pi_0^\mu \otimes \pi_0^{b_1} \otimes \pi_0^{b_2} \otimes \dots \otimes \pi_0^{b_N}$,¹¹ we could define the probability space $(\Omega, \mathcal{F}, \mathbb{P}_{\pi_0})$ such that

$$\begin{aligned} \mathcal{F} &:= \mathcal{B}(\Omega) \\ \mathbb{P}_{\pi_0} &:= \pi_0^\mu \otimes_{i=1}^N \pi_0^{b_i} \otimes \mathcal{U}[0, 1] \otimes \mathbf{W}, \end{aligned}$$

where $\mathbf{W} = \otimes_{i=1}^N \mathbf{W}_0$. Moreover, we let $\mathbb{P}_{x,b} := \delta_x \otimes_{i=1}^N \delta_{b_i} \otimes \mathcal{U}[0, 1] \otimes \mathbf{W}$ to denote the measure when the payoff-relevant state and biases are fixed, where δ_x and δ_{b_i} are Dirac measure at x and b_i respectively.

Next we formally define the space and law of observable histories. In our model the observable history consists of realised signal path, the attention allocation scheme as well as realised value of the randomisation device. The later two components are not dispensable as the decision maker is allowed to use mixed strategy. Consequently, we

¹¹Because of the independence assumption, the joint prior is a product measure.

define the canonical space of observables history

$$\mathbb{H} := C([0, T], \mathbb{R}^N) \times C([0, T], \mathbb{R}_+^N) \times [0, 1], \quad (22)$$

and let $H_t := (X_t, q(t), \zeta)$ to denote the observable history up to time t . Notice that H is a random variable (depends on β) mapping from Ω to $\mathbb{H}$ and H_t is the time- t projection of it. Furthermore, we let $H_{\leq t}$ denotes the observable history up to time t and $\mathbb{H}_t$ denotes the set of observable histories at t . Next we define the admissible attention allocation in our framework.

Definition 1 (Regular admissible attention scheme). *Fix $T \leq \infty$. An attention scheme β on $[0, T]$ is regular admissible if:*

(i) *there is a predictable, nonanticipative Borel feedback functional $\Gamma = (\Gamma_i)_{i=1}^N$ such that*

$$\beta_i(t) = \Gamma_i(t, H_{\leq t}, \zeta), \quad \Gamma_i \geq 0, \quad \sum_{i=1}^N \Gamma_i(t, \cdot) = 1 \quad \text{for all } t;$$

(ii) *the same functional Γ is used for every value of the parameters. In particular, it cannot observe μ or b directly.*

(iii) *for every fixed (x, b) , the augmented closed-loop system*

$$\begin{aligned} dX_i(t) &= (x + b_i)\Gamma_i(t, X_{\leq t}, q_{\leq t}, \zeta) dt + \sigma_i \sqrt{\Gamma_i(t, X_{\leq t}, q_{\leq t}, \zeta)} dW_i(t), \\ dq_i(t) &= \Gamma_i(t, X_{\leq t}, q_{\leq t}, \zeta) dt, \quad q_i(0) = 0, \end{aligned}$$

with the fixed initial signal values, admits a nonexplosive strong solution.

The first bullet point in the definition makes sure that admissible scheme cannot use future information while the second bullet point guarantees that the decision maker can introduce correlation between actions and states only thorough observables. Now fix a state $\mu = x$ and an admissible attention allocation β , we let $\mathcal{G}_t^{\beta, x}$ to denote the filtration generated by H_t . In addition, we define $\mathcal{G}_t^\beta$ to be the filtration induced by H_t when the prior about μ is not degenerate. Given $\mu = x$, we further let P_x^β to denote the structure law of observables and $\mathbf{P}_x^\beta$ to denote joint structure law of observables and the bias, namely

(H, b) . In addition to $\mu = x$ and β suppose the bias is also fixed, we then define $P_{x,b}^\beta$ to be the corresponding structural law of the observables.

We are now ready introduce the concept of **adapted garbling** and **adapted Blackwell dominance**.

Definition 2 (Adapted Garbling). *A probability kernel $K : \mathbb{H} \times \mathcal{B}(\mathbb{H}) \rightarrow [0, 1]$ is an adapted garbling from β_1 to β_2 if:*

- (i) $P_x^{\beta_2} = P_x^{\beta_1} K$ for all $x \in \mathbb{R}$;
- (ii) for all $t \geq 0$, there exists $K_t : \mathbb{H}_t \times \mathcal{B}(\mathbb{H}_t) \rightarrow [0, 1]$ such that, for all $A \in \mathcal{G}_t^{\beta_2}$, $K_t(\cdot, A)$ is $\mathcal{G}_t^{\beta_1}$ -measurable and $K(y^{\beta_1}, A) = K_t(y_{\leq t}^{\beta_1}, A)$ $P_{x,t}^{\beta_1}$ -a.s., simultaneously for every $x \in \mathbb{R}$.

If such a kernel exists, β_1 is said to adapted Blackwell dominate β_2 .

The notion of adapted Blackwell dominance extends the concept of Blackwell dominance to sequential experiments. We will use this concept to establish that β^* provides more information than any other feasible β .

B Proof of the Optimal Attention Allocation

In this section, we prove the results concerning the optimal allocation of attention. Our goal here is to show that using information generated by β^* , the attention allocation maximising the information flow, up to time t we can replicated the distribution of time- t observable history induced by any competing attention allocation strategy $\tilde{\beta}$.

Throughout this section, we will use some further notation introduced here. For $i \in \{1, \dots, N\}$, let $Y_i^\beta(t) := X_{it} - X_{i0} - b_{i0}q_i^\beta(t)$ be the source i 's signal corrected for the bias accumulated until time t . The information score for source i at calendar time t was already defined in the main body of the paper, and can here be re-written as $S_i^\beta(t) = \frac{Y_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)}$. Recall the aggregate information score $S^\beta(t) = \sum_{i=1}^N S_i^\beta(t)$. Recall also the effective time spent on source i is $\mathcal{I}_i^\beta(t) = \frac{q_i^\beta(t)}{\sigma_i^2 + v_i^2 q_i^\beta(t)}$ and in aggregate $\mathcal{I}^\beta(t) = \sum_{i=1}^N \mathcal{I}_i^\beta(t)$. Finally, the effective instantaneous attention accumulation of source i given $\beta(t)$ is $\iota_i^\beta(t) = \frac{\sigma_i^2 \beta_i(t)}{(\sigma_i^2 + v_i^2 q_i^\beta(t))^2}$. When clear from the context, we will omit the dependence on β . If $\beta = \beta^*$, we will write S^* instead of S^{β^*} , $\mathcal{I}^*$ instead of $\mathcal{I}^{\beta^*}$, etc.; similarly, we will shorten the notation

when $\beta = \tilde{\beta}$. Finally, let b_{it} denote the posterior average of the bias term of source i at time t . We denote with $\delta b_i := b_i - b_{i0}$ the change from the prior. Furthermore, we write $b_0 + \delta b := (b_{i0} + \delta b_i)_{i=1}^N$.

The rest of Appendix B is organised as follows. In Section B.1, we prove Proposition 1 which shows that q^* is the uniquely optimal cumulative attention allocation that maximises information flow $\mathcal{I}$ at any given point in time. In Section B.2 we prove few technical lemmas. Specifically, we show that the information score S_i admits a convenient semimartingale representation in Lemma 1. Then, Lemma 2 allows us to characterise the likelihood of $\mu = x$ (relative to $\mu = 0$) in terms of the aggregated information score S and the information flow $\mathcal{I}$. Finally, we state and prove Lemma 6, which establishes the existence of an adapted disintegration. Armed with these results, we prove Proposition 2 in Section B.3.

B.1 Proof of Proposition 1

Proof of Proposition 1. Fix $t \geq 0$. The fixed-horizon problem we consider is to choose a feasible cumulative attention strategy q to maximise $\mathcal{I}(t) = \sum_{i=1}^N \frac{q_i(t)}{v_i^2 q_i(t) + \sigma_i^2}$ such that $\sum_{i=1}^N q_i(t) = t$. When implicitly clear, we will omit the dependence of $\mathcal{I}$ and q on t . We divide the proof in two cases: (i) $v_i > 0$ for all $i \in \{1, \dots, N\}$, and (ii) there exists $i \in \{1, \dots, N\}$ such that $v_i = 0$.

Case (i). Suppose $v_i > 0$ for all $i \in \{1, \dots, N\}$. For each source i , define $f_i(q_i) := \frac{q_i}{v_i^2 q_i + \sigma_i^2}$. Then, the first and second order derivatives are $f'_i(q_i) = \frac{\sigma_i^2}{(v_i^2 q_i + \sigma_i^2)^2}$ and $f''_i(q_i) = -\frac{2v_i \sigma_i^2}{(v_i^2 q_i + \sigma_i^2)^3} < 0$. Thus the fixed-horizon problem is a concave resource-allocation problem with a separable and strictly concave objective function in each q_i when $v_i > 0$. Therefore, the Karush–Kuhn–Tucker (KKT) conditions are necessary and sufficient.

Let λ denote the Lagrange multiplier on the resource constraint. The marginal value of allocating additional attention to source i is given by $f'_i(q_i)$. For any source i in the consumption set, the KKT condition is $\frac{\sigma_i^2}{(v_i^2 q_i + \sigma_i^2)^2} = \lambda$, or, equivalently, $\frac{\sigma_i}{v_i^2 q_i + \sigma_i^2} = \sqrt{\lambda}$.

Hence, $v_i^2 q_i + \sigma_i^2 = \frac{\sigma_i}{\sqrt{\lambda}}$. Let $h := \frac{1}{\sqrt{\lambda}}$. Then, the optimal attention allocation satisfies

$$q_i = \frac{\sigma_i}{v_i^2} (h - \sigma_i). \quad (23)$$

Since $q_i > 0$ if and only if $h > \sigma_i$, source i is active exactly when its sampling noise σ_i lies below h . Hence, the t -consumption set must be of the form $\{1, \dots, m\}$ for some $m \in \{1, \dots, N\}$. Conditional on the consumption set being $\{1, \dots, m\}$, the resource constraint implies $\sum_{i=1}^m \frac{\sigma_i}{v_i^2} (h - \sigma_i) = t$. Expanding the left-hand side gives $h \sum_{i=1}^m \frac{\sigma_i}{v_i^2} - \sum_{i=1}^m \frac{\sigma_i^2}{v_i^2} = t$. Thus, the “water level” when the t -consumption set is $\{1, 2, \dots, m\}$ is

$$h_{m+1}(t) = \frac{t + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} \quad (24)$$

Substitute (24) into (23), we obtain $q_i^*(t)$ for $i \in \{1, \dots, m\}$. For $i \in \{m+1, \dots, N\}$, instead, we have $h_{m+1} \leq \sigma_i$, so $q_i^*(t) = 0$.

The t -consumption set is pinned down by the requirement that the sources in the t -consumption set have positive attention while the inactive sources do not. Equivalently, the t -consumption set is $\{1, \dots, m\}$ if and only if $\sigma_m < h_{m+1}(t) \leq \sigma_{m+1}$, with the convention that $\sigma_{N+1} = +\infty$. The time threshold at which source m becomes active is determined by $h_m(t) = \sigma_m$. Using (24), this gives $\frac{t + \sum_{k=1}^{m-1} \sigma_k^2 / v_k^2}{\sum_{k=1}^{m-1} \sigma_k / v_k^2} = \sigma_m$. Solving for t , the time threshold at which source m enters the consideration set is $T_m = \sum_{k=1}^{m-1} \frac{\sigma_k(\sigma_m - \sigma_k)}{v_k^2}$ for all $m \in \{2, \dots, N\}$. Set $T_1 := 0$ and $T_{N+1} := +\infty$. Then, we conclude that for all $m \in \{1, \dots, N\}$ and $t \in [T_m, T_{m+1})$, the optimal cumulative allocation is $q_i^*(t)$.

Case (ii). $v_J = 0$ implies

$$\mathcal{I}_J(q) = \frac{q}{\sigma_J^2}, \quad \mathcal{I}'_J(q) = \frac{1}{\sigma_J^2}.$$

For $i < J$, one has $v_i > 0$ by the definition of J . Before T_J , the water-filling calculation in case 1 uses only those active sources. Every source $i \geq J$ is inactive and satisfies the required KKT inequality $\mathcal{I}'_i(0) = 1/\sigma_i^2 \leq 1/h_m(t)^2$. Hence the displayed stage- m formulas remain globally optimal for every $m < J$, even though sources with known biases may

appear among the inactive sources.

If $J = 1$, then $T_J = 0$, there are no earlier sources, and $q_J^*(t) = t - T_J$ simply reads $q_1^*(t) = t$. Suppose now that $J > 1$. Approaching T_J through the stage $J - 1$ formula gives

$$h_{J-1}(T_J) = \sigma_J.$$

This follows from the same direct substitution used in the threshold calculation in case 1. Therefore the left limit of the selected cumulative path is

$$q_i^*(T_J) = \frac{\sigma_i}{v_i^2}(\sigma_J - \sigma_i), \quad i < J,$$

because the water level equals σ_J . The marginal information of each earlier source with $q_i^*(T_J) > 0$ is then

$$\begin{aligned} \mathcal{I}'_i(q_i^*(T_J)) &= \frac{\sigma_i^2}{(\sigma_i^2 + v_i^2 q_i^*(T_J))^2} \\ &= \frac{\sigma_i^2}{\{\sigma_i^2 + \sigma_i(\sigma_J - \sigma_i)\}^2} \\ &= \frac{\sigma_i^2}{(\sigma_i \sigma_J)^2} = \frac{1}{\sigma_J^2} = p'_J(q). \end{aligned}$$

Any further attention to an old biased source $i < J$ would make its marginal information strictly less than $1/\sigma_J^2$, whereas the known-bias source J retains the constant marginal information $1/\sigma_J^2$. A source $i > J$ has

$$\mathcal{I}'_i(0) = \frac{1}{\sigma_i^2} \leq \frac{1}{\sigma_J^2}$$

because $\sigma_i \geq \sigma_J$. Therefore all marginal attention after T_J may optimally be assigned to source J .

The cumulative allocation is feasible because

$$\sum_{i < J} q_i^*(T_J) = \sum_{i < J} \frac{\sigma_i(\sigma_J - \sigma_i)}{v_i^2} = T_J,$$

and $q_J^*(t) = t - T_J$, which means the total cumulative attention sums to t .

For completeness, global optimality after T_J follows directly from argument below. Fix $t \geq T_J$, write $q^* = q^*(t)$, and set $\lambda := 1/\sigma_J^2$. For every earlier source with $q_i^* > 0$, the preceding calculation gives $\mathcal{I}'_i(q_i^*) = \lambda$; if an earlier source has $\sigma_i = \sigma_J$, the selected value is $q_i^* = 0$ and $\mathcal{I}'_i(0) = \lambda$. The selected source satisfies $p'_J(q_J^*) = \lambda$. Finally, for every inactive $i > J$,

$$\mathcal{I}'_i(0) = \frac{1}{\sigma_i^2} \leq \frac{1}{\sigma_J^2} = \lambda.$$

Now let $\bar{q} \in \mathcal{Q}(t)$ be arbitrary. Concavity and differentiability of every $\mathcal{I}_i$ imply

$$\mathcal{I}_i(\bar{q}_i) - \mathcal{I}_i(q_i^*) \leq \mathcal{I}'_i(q_i^*)(\bar{q}_i - q_i^*).$$

For every coordinate with $q_i^* > 0$, the right-hand slope equals λ . For every coordinate with $q_i^* = 0$, one has $\bar{q}_i - q_i^* = \bar{q}_i \geq 0$ and $\mathcal{I}'_i(0) \leq \lambda$. Summing yields

$$\begin{aligned} \mathcal{I}(\bar{q}) - \mathcal{I}(q^*) &\leq \lambda \sum_{i=1}^N (\bar{q}_i - q_i^*) \\ &= \lambda(t - t) = 0. \end{aligned}$$

Thus $q^*(t)$ is a global maximizer for every $t \geq T_J$. ■

B.2 Proofs of Technical Lemmas

Proof of Lemma 1. The information about μ is encoded in $P_{x, b_0 + \delta b}^\beta$ (more specifically the law of X_t given $\beta, \mu = x$, and the bias vector changed by δb from the prior). When all increments δb_i are zero, we have P_{x, b_0}^β . Girsanov's formula gives

$$\frac{dP_{x, b_0 + \delta b}^\beta}{dP_{x, b_0}^\beta} \Big|_{\mathcal{H}_t} = \prod_{i=1}^N \exp \left\{ \frac{\delta b_i}{\sigma_i^2} (Y_i(t) - x q_i^\beta(t)) - \frac{\delta b_i^2}{2\sigma_i^2} q_i^\beta(t) \right\} \quad (25)$$

To see this we first consider the state where the bias is the prior mean, which is just b_0 . In this case, the demeaned signal process, denoted by Y_{i0} is

$$dY_{i0} = x\beta_i(t)dt + \sigma_i\sqrt{\beta_i(t)}dW_{it} \quad (26)$$

On the other hand, in the state where the bias is b_i , the demeaned signal process is

$$dY_i = (x\beta_i(t) + \delta b_i\beta_i(t))dt + \sigma_i\sqrt{\beta_i(t)}dW_{it} \quad (27)$$

Notice that the two processes differ in drift and the difference is $\delta b_i\beta_i(t)$. Moreover, we have $\delta b_i\beta_i(t) = \sigma_i\sqrt{\beta_i(t)} \cdot \frac{\delta b_i\sqrt{\beta_i(t)}}{\sigma_i}$. Therefore by the Girsanov theorem, there exist a process

$$M_t := \exp \left\{ \sum_{i=1}^N \int_0^t \frac{\delta b_i\sqrt{\beta_i(s)}}{\sigma_i} dW_{is} - \frac{1}{2} \sum_{i=1}^N \int_0^t \left(\frac{\delta b_i\sqrt{\beta_i(s)}}{\sigma_i} \right)^2 ds \right\} \quad (28)$$

such that Y_i admits the same law as Y_{i0} with respect to the measure $dQ := M_T d\mathbb{P}_{x,b_0}$, if M_t is a martingale. Here $T \in [0, \infty)$ is any finite time horizon. Since Novikov's condition is satisfied, M_t is indeed a martingale. We know that the law of Y_{i0} is the push forward of $\mathbb{P}_{x,b_0}$. Moreover ζ is independent and admissible attention allocations are essentially functions of $(Y(t), \zeta)$. As a result, 28 is the Radon-Nikodym derivative of $P_{x,b_0+\delta b}^\beta$ with respect to P_{x,b_0}^β . Now because $Y_i(t) - xq_i^\beta(t) = \sigma_i \int_0^t \sqrt{\beta_i(s)} dW_{is}$ and $\left\langle \frac{1}{\sigma_i} \int_0^\cdot \sqrt{\beta_i(s)} dW_{is} \right\rangle_t = \frac{q_i^\beta(t)}{\sigma_i^2}$, 25 holds.

Conditional on $\mu = x$, we have $dY_i(t) = (x + \delta b_i)\beta_i(t)dt + \sigma_i\sqrt{\beta_i(t)}dW_{it}$ for all $i \in \{1, \dots, N\}$. Notice that if $v_i = 0$, then $\delta b_i = 0$ and we have,

$$dS_i(t) = \frac{x\beta_i(t)}{\sigma_i^2} dt + \frac{\sqrt{\beta_i(t)}}{\sigma_i} dW_{it}$$

Since $\iota^\beta(t) = \frac{\beta_i(t)}{\sigma_i^2}$, this proves the result for $v_i = 0$. In the remainder of the proof, we set $v_i > 0$.

If $v_i > 0$, the prior density of δb_i is $\mathcal{N}(0, v_i^2)$. Recall that, even under generic attention allocation, the posterior remains Gaussian. Furthermore, because both the Gaussian prior density of $(\delta b_i)_i$ and the likelihood (25) factorise, the posterior density also factorises across

sources. Hence, we can find the posterior density by multiplying the prior density with the source i 's factor of (25), which gives

$$\exp \left\{ \frac{\delta b_i}{\sigma_i^2} (Y_i(t) - x q_i(t)) - \frac{\delta b_i^2}{2\sigma_i^2} q_i(t) - \frac{\delta b_i^2}{2v_i^2} \right\}$$

Collecting the linear and quadratic terms in δb_i gives posterior precision $\frac{1}{v_i^2} + \frac{q_i(t)}{\sigma_i^2}$, and hence the posterior mean is

$$\frac{\left(\frac{Y_i(t) - x q_i(t)}{\sigma_i^2} \right)}{\left(\frac{1}{v_i^2} + \frac{q_i(t)}{\sigma_i^2} \right)} = \frac{v_i^2 (Y_i(t) - x q_i(t))}{\sigma_i^2 + v_i^2 q_i(t)} =: m_i(t)$$

The controlled likelihood is valid under admissible attention allocation because the same nonanticipative functional is used under every parameter, namely the attention allocation $\beta(t)$ does not depend directly on δb . As a result, we can treat $q(t)$ as constant when collecting terms.

Next, define the innovation.

$$N_i(t) := Y_i(t) - \int_0^t (x + m_i(s)) \beta_i(s) ds \quad (29)$$

Standard results about innovation processes imply that N_i is a continuous local martingale in the state- x observation filtration $\mathcal{G}_x^\beta$. Notice that

$$N_i(t) = \sigma_i \int_0^t \sqrt{\beta(s)}_i dW_{is} + \int_0^s (\delta b_i - m_i(s)) \beta_i(s) ds \quad (30)$$

Let's define $L(t) := \int_0^s (\delta b_i - m_i(s)) \beta_i(s) ds$ and $R(t) := \sigma_i \int_0^t \sqrt{\beta(s)}_i dW_{is}$. Since $L(t)$ is just a ordinary integration with nice integrant, it has finite total variation on any bounded time interval $[0, t]$. Thus it has zero quadratic variation. Moreover, the cross variation of $L(t)$ and $R(t)$ is also zero as

$$\sum_j |\Delta_j R \Delta_j L| \leq (\max_j |\Delta_j R|) \sum_j |\Delta_j L| \rightarrow 0, \quad (31)$$

where $\Delta_j L = L(t_j) - L(t_{j-1})$ (similarly $\Delta_j R = R(t_j) - R(t_{j-1})$) for a partition $\{t_0 = 0, t_1, t_2, \dots, t_n = t\}$ on $[0, t]$. Thus, the filtering error does not contribute to the quadratic variation of $N_i(t)$.

$$d\langle N_i \rangle_t = \sigma_i^2 \beta_i(t) dt \quad (32)$$

Rearranging (29) gives $dY_i(t) = (x + m_i(t))\beta_i(t) dt + dN_i(t)$. Notice that $d[\sigma_i^2 + v_i^2 q_i(t)] = v_i^2 dq_i(t) = v_i^2 \beta_i(t) dt$. Since $[\sigma_i^2 + v_i^2 q_i(t)]$ has finite quadratic variation, Itô's formula gives

$$\begin{aligned} dS_i(t) &= \frac{1}{\sigma_i^2 + v_i^2 q_i(t)} dY_i(t) - \frac{Y_i(t)}{[\sigma_i^2 + v_i^2 q_i(t)]^2} d[\sigma_i^2 + v_i^2 q_i(t)] \\ &= \frac{1}{\sigma_i^2 + v_i^2 q_i(t)} [(x + m_i(t))\beta_i(t) dt + dN_i(t)] - \frac{Y_i(t)}{[\sigma_i^2 + v_i^2 q_i(t)]^2} v_i^2 \beta_i(t) dt \\ &= \left[\frac{x + m_i(t)}{\sigma_i^2 + v_i^2 q_i(t)} - \frac{v_i^2 Y_i(t)}{[\sigma_i^2 + v_i^2 q_i(t)]^2} \right] \beta_i(t) dt + \frac{1}{\sigma_i^2 + v_i^2 q_i(t)} dN(t) \end{aligned}$$

Using the definition of $m_i(t)$, the drift term simplifies to

$$\left[\frac{x + m_i(t)}{\sigma_i^2 + v_i^2 q_i(t)} - \frac{v_i^2 Y_i(t)}{[\sigma_i^2 + v_i^2 q_i(t)]^2} \right] \beta_i(t) = x \underbrace{\frac{\sigma_i^2}{[\sigma_i^2 + v_i^2 q_i(t)]^2} \beta_i(t)}_{=l_i^\beta(t)}$$

For the diffusion term, define $M_i(t) := \int_0^t [\sigma_i^2 + v_i^2 q_i(s)]^{-1} dN(s)$. By (32), we conclude that the quadratic variation of $M_i(t)$ is $d\langle M \rangle_t = \frac{\sigma_i^2 \beta_i(t)}{[\sigma_i^2 + v_i^2 q_i(s)]^2} dt = l_i^\beta(t) dt$. This proves equation (4).

To conclude the proof, notice that the original Brownian process W_{it} have zero cross-variation, and finite-variation filtering errors do not alter cross-variation. Hence, $\langle N_i, N_j \rangle = \langle M_i, M_j \rangle = 0$ for $i \neq j$. This proves equation (5). $\blacksquare$

Proof of Lemma 2. Fix $x \in \mathbb{R}$ and $T < +\infty$. Fix $t \leq T$. We start by deriving the relative likelihood while holding the vector of changes in bias levels δb fixed. By Girsanov theorem

we get the following expression for the likelihood.

$$\ell_t(x, \delta b) = \prod_{i=1}^N \exp \left\{ \frac{x}{\sigma_i^2} (Y_i(t) - \delta b_i q_i(t)) - \frac{x^2}{2\sigma_i^2} q_i(t) \right\} \quad (33)$$

Indeed, under $P_{0, b_0 + \delta b}^\beta$,

$$Y_i - \delta b_i q_i = \sigma_i \int_0^t \sqrt{\beta_i(s)} dW_{is}.$$

The log density for adding drift $x\beta_i$ is consequently

$$\frac{x}{\sigma_i} \int_0^t \sqrt{\beta_i(s)} dW_{is} - \frac{x^2}{2\sigma_i^2} \int_0^t \beta_i(s) ds,$$

which is the i th exponent in (33). Its total quadratic-variation is bounded by

$$x^2 \sum_{i=1}^N \frac{q_i}{\sigma_i^2} \leq \frac{x^2 t}{\sigma_1^2}.$$

Thus Novikov's condition holds. Then Girsanov theorem and the fact that β is a function of observables only together establish the density captured by equation (33).

Next we integrate out the latent biases. From the proof of Lemma 1, we know that, under $\mathbf{P}_0^\beta$ and conditionally on $\mathcal{G}_t^\beta$, the δb_i variables are independent and normally distributed with mean $m_i(t) = \frac{v_i^2 Y_i(t)}{\sigma_i^2 + v_i^2 q_i(t)}$ and variance $s_i^2(t) := \frac{v_i^2 \sigma_i^2}{\sigma_i^2 + v_i^2 q_i(t)}$. Notice that these formulas also cover the case of $v_i = 0$, where both quantities are zero.

Denote with ν the Gaussian distribution of the vector δb . Let f be a generic $\mathcal{G}_t^\beta$ -measurable bounded function. Notice that the expectation of f can be written averaging over ν : $\mathbb{E}_{P_x^\beta}[f] = \mathbb{E}_{\mathbf{P}_x^\beta}[f] = \int \mathbb{E}_{P_{x, b_0 + \delta b}^\beta}[f] \nu(d\delta b)$. Within the integral, δb is fixed, so we can use the relative likelihood from equation (33) to express the integral in terms of the law under $\mu = 0$: $\int \mathbb{E}_{P_{x, b_0 + \delta b}^\beta}[f] \nu(d\delta b) = \int \mathbb{E}_{P_{0, b_0 + \delta b}^\beta}[f \ell_t(x, \delta b)] \nu(d\delta b) = \mathbb{E}_{\mathbf{P}_0^\beta}[f \ell_t(x, \delta b)]$. Then, using the law of iterated expectations, $\mathbb{E}_{\mathbf{P}_0^\beta}[f \ell_t(x, \delta b)] = \mathbb{E}_{\mathbf{P}_0^\beta} \left[f \mathbb{E}_{\mathbf{P}_0^\beta}[\ell_t(x, \delta b) \mid \mathcal{G}_t^\beta] \right]$.

To summarise, we have

$$\begin{aligned}
\mathbb{E}_{P_x^\beta}[f] &= \int \mathbb{E}_{P_{x,b_0+\delta b}^\beta}[f] \nu(\mathrm{d}\delta b) \\
&= \int \mathbb{E}_{P_{0,b_0+\delta b}^\beta}[f \ell_t(x, \delta b)] \nu(\mathrm{d}\delta b) \\
&= \mathbb{E}_{\mathbf{P}_0^\beta}[f \ell_t(x, \delta b)] \\
&= \mathbb{E}_{\mathbf{P}_0^\beta} \left[f \mathbb{E}_{\mathbf{P}_0^\beta}[\ell_t(x, \delta b) \mid \mathcal{G}_t^\beta] \right] \\
&= \mathbb{E}_{P_0^\beta} \left[f \mathbb{E}_{\mathbf{P}_0^\beta}[\ell_t(x, \delta b) \mid \mathcal{G}_t^\beta] \right].
\end{aligned}$$

Since this holds for all $\mathcal{G}_t^\beta$ -measurable bounded f , we get

$$\left. \frac{\mathrm{d}P_x^\beta}{\mathrm{d}P_0^\beta} \right|_{\mathcal{G}_t^\beta} = \mathbb{E}_{\mathbf{P}_0^\beta}[\ell_t(x, \delta b) \mid \mathcal{G}_t^\beta]. \quad (34)$$

We can further evaluate $\mathbb{E}_{P_0^\beta}[\ell_t(x, \delta b) \mid \mathcal{G}_t^\beta]$. Since the components of δb are independent conditionally on $\mathcal{G}_t^\beta$, the conditional expectation in 34 is the product of the conditional expectations of the sources. For source i , the moment-generating function of a normal variable gives

$$\begin{aligned}
\mathbb{E}_{\mathbf{P}_0^\beta} \left[\exp \left\{ \frac{x}{\sigma_i^2} (Y_i(t) - \delta b_i q_i(t)) - \frac{x^2 q_i(t)}{2\sigma_i^2} \right\} \middle| \mathcal{G}_t^\beta \right] &= \\
&= \exp \left\{ \frac{x Y_i(t)}{\sigma_i^2} - \frac{x q_i(t) m_i(t)}{\sigma_i^2} + \frac{x^2 q_i^2(t) s_i^2(t)}{2\sigma_i^4} - \frac{x^2 q_i(t)}{2\sigma_i^2} \right\}
\end{aligned}$$

Using the definition of $m_i(t)$, the terms linear in x simplify to

$$x \left(\frac{Y_i(t)}{\sigma_i^2} - \frac{q_i(t) m_i(t)}{\sigma_i^2} \right) = x \frac{Y_i(t)}{\sigma_i^2 + v_i^2 q_i(t)} = x S_i^\beta(t)$$

Using the definition of $s_i^2(t)$, the terms quadratic in x reduce to

$$x^2 \left(\frac{q_i^2(t)s_i^2(t)}{2\sigma_i^4} - \frac{q_i(t)}{2\sigma_i^2} \right) = -\frac{x^2}{2} \frac{q_i(t)}{\sigma_i^2 + v_i^2 q_i(t)} = -\frac{x^2}{2} \mathcal{I}_i(t)$$

Hence, the i -th factor in (34) is $\exp\left\{xS_i^\beta(t) - \frac{x^2}{2}\mathcal{I}_i(t)\right\}$. Using the definitions of S^β and $\mathcal{I}^\beta$, we conclude

$$L_t^\beta(x) := \frac{dP_x^\beta}{dP_0^\beta} \Big|_{\mathcal{G}_t^\beta} = \exp\left\{xS^\beta(t) - \frac{x^2}{2}\mathcal{I}^\beta(t)\right\}$$

Since $L_t^\beta(x) > 0$, P_x^β and P_0^β are equivalent on every finite-horizon observation field. Bayes' formula yields equation (9). The denominator is strictly positive. The denominator is also finite for every Borel probability measure π_0^μ . To see this, notice that, if $\mathcal{I}^\beta(t) > 0$, then completing the square yields $\tilde{x}S^\beta(t) - \frac{\tilde{x}^2}{2}\mathcal{I}^\beta(t) \leq \frac{(S^\beta(t))^2}{2\mathcal{I}^\beta(t)}$ for all $\tilde{x} \in \mathbb{R}$. Hence, the integrand is bounded as a function of $\tilde{x}$. If, instead, $\mathcal{I}^\beta(t) = 0$, we must have $t = 0$ and hence $S^\beta(0) = \mathcal{I}^\beta(0) = 0$; in this case, the denominator equals one. ■

Next we introduce a useful disintegration Lemma that factors the law of observables thorough Brownian path. Before stating the Lemma, let's introduce the **truncation operator** R_t , which maps correspondences and functions defined over $\mathbb{R}$ into their restrictions over the domain $[0, t]$. Similarly, for $s \leq t$, we define $R_{t,s}$ to be the operator restrict the domain $[0, t]$ to $[0, s]$. Given a correspondence or function f , we denote with $f_{\leq t}$ the it's truncated, that is, $f_{\leq t} = R_t(f)$.

Lemma 6. Fix $U \in (0, \infty]$. Let $C \in C([0, U], \mathbb{R})$ be defined such that $C(s) := xs + W_s$, where $W \in C([0, U], \mathbb{R})$ is a standard Brownian motion defined relative to a right-continuous filtration $(\mathcal{G}_s)_{0 \leq s \leq U}$. Denote with $\mathbf{W}_x$ the distribution over sample paths of C given x , and with $\mathbf{W}_0$ the Wiener measure as before. With a slight abuse of notation, let $Y \in \mathbb{H}$ be a history. Denote with $\mathbf{P}_0$ the joint distribution over sample paths of C and histories Y given $x = 0$. Let $\rho : [0, \infty) \rightarrow [0, U]$ be a deterministic, continuous, and strictly increasing function, with $\rho(0) = 0$, $\rho(t) \rightarrow U$ as $t \rightarrow +\infty$, and $\sigma(Y_{\leq t}) \subseteq \mathcal{G}_{\rho(t)}$ for $t < +\infty$.

There exists a Borel probability kernel $K : C([0, U], \mathbb{R}) \times \mathcal{B}(\mathbb{H}) \rightarrow [0, 1]$ such that, for all

$A \in \mathcal{B}(C([0, U], \mathbb{R}))$ and $B \in \mathcal{B}(\mathbb{H})$,

$$\mathbf{P}_0(C \in A, Y \in B) = \int_A K(c, B) \mathbf{W}_0(dc)$$

Moreover, for every $t < +\infty$, there exist a Borel probability kernel $K_t : C([0, \rho(t)], \mathbb{R}) \times \mathcal{B}(\mathbb{H}_t) \rightarrow [0, 1]$ and a Borel set $D_t \subseteq \mathcal{B}(C([0, U], \mathbb{R}))$ both independent of x , such that $\mathbf{W}_x(D_t) = 1$ for all $x \in \mathbb{R}$, and $K(c, R_t^{-1}(B_t)) = K_t(c_{\leq \rho(t)}, B_t)$ for all $c \in D_t$ and $B_t \in \mathcal{B}(\mathbb{H}_t)$.

Proof of Lemma 6. Step 1. We start by showing that, when predicting the history until time s (i.e., $Y_{\leq s}$) using observations from the Brownian motion $C \in C([0, U], \mathbb{R})$, the past path $C_{\leq \rho(s)}$ contains all relevant information and subsequent Brownian increments of C do not provide additional information. Define a binary random variable $\chi = \mathbf{1}(Y_{\leq s} \in B_s)$ for whether the history until time s belongs to a set $B_s \in \mathcal{B}(\mathbb{H}_s)$, where $\mathcal{B}(\mathbb{H}_s)$ denotes the Borel σ -algebra generated by H_s . Let $\mathcal{F}_a = \sigma(\{C_u : u \leq a\})$ for $a \in [0, U]$. In this step, we want to show that, if $0 \leq s \leq t < +\infty$, then $\mathbf{E}_0[\chi \mid \mathcal{F}_{\rho(t)}] = \mathbf{E}_0[\chi \mid \mathcal{F}_{\rho(s)}]$ $\mathbf{P}_0$ -a.s., that is, the probabilistic assessment of $Y_{\leq s}$ using information in $\mathcal{F}_{\rho(s)}$ does not change if we use (more) information in $\mathcal{F}_{\rho(t)}$.

To show this, first notice that, by definition, ρ is such that $\sigma(Y_{\leq s}) \subseteq \mathcal{G}_{\rho(s)}$ for all $s < +\infty$. Hence, Z is $\mathcal{G}_{\rho(s)}$ -measurable. Further, since ρ is increasing, $\mathcal{G}_{\rho(s)} \subseteq \mathcal{G}_{\rho(t)}$, that is, $\mathcal{G}_{\rho(s)}$ -measurability implies $\mathcal{G}_{\rho(t)}$ -measurability. Let $\hat{\chi} = \mathbf{E}_0[\chi \mid \mathcal{F}_{\rho(s)}]$, which is clearly $\mathcal{F}_{\rho(s)}$ -measurable. Let f be a real-valued bounded function of $C_{\leq \rho(s)}$, which is clearly $\mathcal{F}_{\rho(s)}$ -measurable. Let g be a real-valued bounded function of increments $(C_r - C_{\rho(s)})_{\rho(s) \leq r \leq \rho(t)}$. By the properties of the Brownian motion, g is independent of $\mathcal{G}_{\rho(s)}$. Hence,

$$\begin{aligned} \mathbf{E}_0[f g(\chi - \hat{\chi})] &= \mathbf{E}_0[\mathbf{E}_0[f g(\chi - \hat{\chi}) \mid \mathcal{G}_{\rho(s)}]] \\ &= \mathbf{E}_0[f(\chi - \hat{\chi}) \mathbf{E}_0[g \mid \mathcal{G}_{\rho(s)}]] \\ &= \mathbf{E}_0[f(\chi - \hat{\chi}) \mathbf{E}_0[g]] \\ &= \mathbf{E}_0[f(\chi - \hat{\chi})] \mathbf{E}_0[g] \end{aligned}$$

where the first equality is the law of iterated expectations; the second equality uses the fact that f , χ and $\hat{\chi}$ are $\mathcal{G}_{\rho(s)}$ -measurable; the third equality uses the fact that g is independent

of $\mathcal{G}_{\rho(s)}$. Then,

$$\begin{aligned}\mathbf{E}_0[f(\chi - \hat{\chi})] &= \mathbf{E}_0[\mathbf{E}_0[f(\chi - \hat{\chi}) \mid \mathcal{F}_{\rho(s)}]] \\ &= \mathbf{E}_0[f \mathbf{E}_0[(\chi - \hat{\chi}) \mid \mathcal{F}_{\rho(s)}]] \\ &= 0\end{aligned}$$

where the first equality is again the law of iterated expectations; the second equality uses the fact that f is $\mathcal{F}_{\rho(s)}$ -measurable; the third equality uses the fact that $\mathbf{E}_0[(\chi - \hat{\chi}) \mid \mathcal{F}_{\rho(s)}] = 0$ by definition of $\hat{\chi}$. Hence, $\mathbf{E}_0[f(\chi - \hat{\chi})] = 0$.

By the monotone-class theorem, $\mathbf{E}_0[h(\chi - \hat{\chi})] = 0$ for all $\mathcal{F}_{\rho(t)}$ -measurable bounded h . Notice that $\hat{\chi}$ is $\mathcal{F}_{\rho(t)}$ -measurable. Furthermore, for all $F \in \mathcal{F}_{\rho(t)}$, $\mathbf{E}_0[\hat{\chi} \mathbf{1}(F)] = \mathbf{E}_0[\chi \mathbf{1}(F)]$. Hence, $\hat{\chi}$ is a version of the conditional expectation of χ with respect to $\mathcal{F}_{\rho(t)}$, that is, $\hat{\chi} = \mathbf{E}_0[\chi \mid \mathcal{F}_{\rho(t)}]$. We conclude that $\mathbf{E}_0[\chi \mid \mathcal{F}_{\rho(s)}] = \mathbf{E}_0[\chi \mid \mathcal{F}_{\rho(t)}]$.

Step 2. For all $t < +\infty$, both $C_{\leq \rho(t)}$ and $Y_{\leq t}$ live in standard-Borel spaces. Hence, for all $t < +\infty$, there exists a Borel kernel $K_t : C(\rho(\mathbb{R}), \mathbb{R}) \times \mathcal{B}(\mathbb{H}) \rightarrow [0, 1]$ such that

$$K_t(C_{\leq \rho(t)}, B_t) = \mathbf{E}_0[Y_{\leq t} \in B_t \mid \mathcal{F}_{\rho(t)}] \quad \mathbf{P}_0\text{-a.s.}$$

Given the existence of the Borel kernels for each $t < +\infty$, in this step we verify the mutual consistency of such kernels across t and therefore construct a kernel K whose marginals are the K_t kernels.

In order to show the mutual consistency of the kernels over t , we need to show that $K_t(C_{\leq \rho(t)}, R_{t,s}^{-1}(B_s)) = K_s(C_{\leq \rho(s)}, B_s)$ $\mathbf{P}_0$ -a.s., that is, the probabilistic assessment about history Y until time s is the same whether one directly evaluates the conditional probability of $Y_{\leq s}$ with the information available at s , or evaluates the conditional probability of $Y_{\leq t}$ with the information available at t and then restricts to time s . This can be easily verified using the definition of $R_{t,s}$ (first and second equalities below) and the result of step 1 (third

equality below):

$$\begin{aligned}
\mathbf{E}_0[\mathbf{1}(Y_{\leq t} \in R_{t,s}^{-1}(B_s)) \mid \mathcal{F}_{\rho(t)}] &= \mathbf{E}_0[\mathbf{1}(R_{t,s}(Y_{\leq t}) \in B_s) \mid \mathcal{F}_{\rho(t)}] \\
&= \mathbf{E}_0[\mathbf{1}(Y_{\leq s} \in B_s) \mid \mathcal{F}_{\rho(t)}] \\
&= \mathbf{E}_0[\mathbf{1}(Y_{\leq s} \in B_s) \mid \mathcal{F}_{\rho(s)}]
\end{aligned}$$

Hence, $K_t(C_{\leq \rho(t)}, R_{t,s}^{-1}(B_s)) = K_s(C_{\leq \rho(s)}, B_s)$ $\mathbf{P}_0$ -a.s. for all $B_s \in \mathcal{B}(H_s)$ and all $s < +\infty$.

This proves the mutual consistency of the kernels between t and s .

For each $m \in \mathbb{N}$, there exists a countable algebra $\mathcal{A}_m$ that generates $\mathcal{B}(\mathbb{H}_m)$ (since $\mathbb{H}_m$ is a standard Borel space). Hence, we can define the set of functions for which any pair of kernels is mutually consistent:

$$D := \bigcap_{\substack{m, n \in \mathbb{N} \\ m \leq n}} \bigcap_{B_m \in \mathcal{A}_m} \left\{ c \in C(\rho(\mathbb{R}), \mathbb{R}) : K_n(c_{\leq \rho(n)}, R_{n,m}^{-1}(B_m)) = K_m(c_{\leq \rho(m)}, B_m) \right\}$$

Notice $D \in \mathcal{B}(C(\mathbb{R}, \mathbb{R}))$ and $\mathbf{W}_0(D) = 1$. For $m \leq n$, whether $K_n(c_{\leq \rho(n)}, R_{n,m}^{-1}(B_m)) \neq K_m(c_{\leq \rho(m)}, B_m)$ is determined by $c_{\rho(n)}$. On $\mathcal{F}_{\rho(n)}$, the relative likelihood of $\mathbf{W}_x$ with respect to $\mathbf{W}_0$ is $\frac{d\mathbf{W}_x}{d\mathbf{W}_0} = \exp\{xc_{\rho(n)} - \frac{x^2\rho(n)}{2}\} > 0$. This implies that $\mathbf{W}_x$ and $\mathbf{W}_0$ are mutually absolutely continuous, and hence $\mathbf{W}_0(D) = 1$ implies $\mathbf{W}_x(D) = 1$ for all $x \in \mathbb{R}$. Hence, we have established that, for $c \in D$, the Borel kernels are consistent over all elements of $\{\mathcal{A}_m\}_{m \in \mathbb{N}}$. Invoking the π - λ theorem, we conclude that, for $c \in D$, the Borel kernels are consistent over all elements of $\{\mathcal{B}(\mathbb{H}_m)\}_{m \in \mathbb{N}}$.

For every $c \in D$, the laws $\{K_n(c_{\leq \rho(n)}, \cdot)\}_{n \geq 1}$, have compatible marginals. The standard-Borel projective-limit theorem gives a unique law $K(c, \cdot)$ on $\mathbb{H}$ with these integer marginals.¹² To complete the definition of K , fix $y^0 \in \mathbb{H}$ and set $K(c, \cdot) = \delta_{y^0}$ for $c \notin D$ (the choice of y^0 does not matter as $\mathbf{W}_x(D^c) = 0$). We have

$$K(c, R_n^{-1}(B_n)) = \mathbf{1}(c \in D) K_n(c, B_n) + \mathbf{1}(c \notin D) \mathbf{1}(y_{\leq n}^0 \in B_n)$$

Notice that, for fixed B_n , the map $c \mapsto K(c, R_n^{-1}(B_n))$ is measurable. $\{R_n^{-1}(B_n) : n \in$

¹²O. Kallenberg, *Foundations of Modern Probability*, 3rd ed., Springer, 2021, Corollary 8.22, pp. 178–179.

$\mathbf{N}, B_n \in \mathcal{B}(\mathbb{H}_n)\}$ is a π -system generating $\mathcal{B}(\mathbb{H})$, and it is contained in the λ -system $\{B \in \mathcal{B}(\mathbb{H}) : c \mapsto K(c, B) \text{ is measurable}\}$. By the π - λ theorem, we can extend measurability of $c \mapsto K(c, B)$ to all $B \in \mathcal{B}(\mathbb{H})$. Hence, we conclude that K is a Borel kernel.

Step 3. Next, we verify that, for all $A \in \mathcal{B}(C([0, U], \mathbb{R}))$ and $B \in \mathcal{B}(\mathbb{H})$,

$$\mathbf{P}_0(C \in A, Y \in B) = \int_A K(c, B) \mathbf{W}_0(dc)$$

To verify this, fix $m \geq 1$ and $B_m \in \mathcal{B}(\mathbb{H}_m)$. Let $B = R_m^{-1}(B_m)$. On non-null sets (i.e., outside D), we have $K(C, B) = K_m(C_{\leq \rho(m)}, B_m)$. Recall that, for $n \geq m$, $K_m(C_{\leq \rho(m)}, B_m) = K_n(C_{\leq \rho(n)}, R_{n,m}^{-1}(B_m))$. Moreover, using the fact that $R_m(Y) = R_{n,m}(Y_{\leq n})$,

$$\begin{aligned} K_n(C_{\leq \rho(n)}, R_{n,m}^{-1}(B_m)) &= \mathbb{E}[\mathbf{1}(Y_{\leq n} \in R_{n,m}^{-1}(B_m)) \mid \mathcal{F}_{\rho(n)}] \\ &= \mathbb{E}[\mathbf{1}(R_{n,m}(Y_{\leq n}) \in B_m) \mid \mathcal{F}_{\rho(n)}] \\ &= \mathbb{E}[\mathbf{1}(R_m(Y) \in B_m) \mid \mathcal{F}_{\rho(n)}] \\ &= \mathbb{E}[\mathbf{1}(Y \in B) \mid \mathcal{F}_{\rho(n)}] \end{aligned}$$

Hence, $K(C, B) = \mathbb{E}[\mathbf{1}(Y \in B) \mid \mathcal{F}_{\rho(n)}]$ for all $n \geq m$. Since $\rho(n) \rightarrow U$ as $n \rightarrow +\infty$, we have $\mathcal{F}_{\rho(n)} \rightarrow \sigma(C)$. By Lévy's upward theorem, $K(C, B) = \mathbf{E}_0[\mathbf{1}(Y \in B) \mid \sigma(C)]$.

Fix $A \in \mathcal{B}(C([0, U], \mathbb{R}))$. Since $\mathbf{1}(C \in A)$ is $\sigma(C)$ -measurable, we have

$$\begin{aligned} \mathbf{P}_0(C \in A, Y \in B) &= \mathbf{E}_0[\mathbf{1}(C \in A) \mathbf{1}(Y \in B)] \\ &= \mathbf{E}_0[\mathbf{E}_0[\mathbf{1}(C \in A) \mathbf{1}(Y \in B) \mid \sigma(C)]] \\ &= \mathbf{E}_0[\mathbf{1}(C \in A) \mathbf{E}_0[\mathbf{1}(Y \in B) \mid \sigma(C)]] \\ &= \mathbf{E}_0[\mathbf{1}(C \in A) K(C, B)] \\ &= \int_A K(c, B) \mathbf{W}_0(dc) \end{aligned}$$

This holds for $B \in \{\tilde{B} \in \mathcal{B}(\mathbb{H}) : \tilde{B} = R_m^{-1}(B_m), m \geq 1\}$. A final π - λ argument extends this to all $B \in \mathcal{B}(\mathbb{H})$.

Step 4. Finally, fix $t < +\infty$ and an integer $n > t$. As in step 2, let $\mathcal{A}_t$ be a countable

algebra generating $\mathcal{B}(\mathbb{H}_t)$. Define

$$D_t := D \cap \bigcap_{B_t \in \mathcal{A}_t} \left\{ c : K_n(c_{\leq \rho(n)}, R_{n,t}^{-1}(B_t)) = K_t(c_{\leq \rho(t)}, B_t) \right\}$$

Arguing as in step 2, we have $\mathbf{W}_x(D_t) = 1$ for every $x \in \mathbb{R}$. For $c \in D_t$, $K(c, R_t^{-1}(B_t)) = K_t(c_{\leq \rho(t)}, B_t)$ for all $B_t \in \mathcal{A}_t$. Again, the π - λ theorem extends the equality to every $B_t \in \mathcal{B}(\mathbb{H}_t)$. $\blacksquare$

B.3 Proof of Proposition 2

Proof of Proposition 2. Step 1. Consider a feasible $\tilde{\beta}$. By Proposition 1, $\mathcal{I}^*(t) \geq \tilde{\mathcal{I}}(t)$.

Step 2. Fix $T \geq 0$. Define $\underline{a}_T := \min_{1 \leq i \leq N} \frac{\sigma_i^2}{(\sigma_i^2 + v_i^2 T)^2} > 0$. For any feasible β , we have $q_i(T) \leq T$. Using the definition of $\mathcal{I}^\beta(t)$,

$$\left. \frac{d\mathcal{I}^\beta(t)}{dt} \right|_{t=T} = \sum_{i=1}^N \frac{\sigma_i^2}{(\sigma_i^2 + v_i^2 q_i(T))^2} \beta_i(T) \geq \underline{a}_T \sum_{i=1}^N \beta_i(T) = \underline{a}_T > 0$$

Hence, $\mathcal{I}^\beta$ is continuous and strictly increasing. Further, notice that $\mathcal{I}^\beta(0) = 0$, so that $\mathcal{I}^\beta$ maps $[0, +\infty]$ into $[0, \mathcal{I}^\beta(\infty)]$, where $\mathcal{I}^\beta(\infty) := \lim_{t \rightarrow \infty} \mathcal{I}^\beta(t) \in (0, \infty]$. Notice that this applies to any feasible β , including $\tilde{\beta}$ and β^* .

Step 3. Fix $\mu = x$. Define $S^*(t) := \sum_{i=1}^N S_i^{\beta^*}(q_i(t))$. By Lemma 1,

$$\mathbb{E}_x[S^*(t)] = x \sum_{i=1}^N \int_0^t \iota_i^{\beta^*}(s) ds = x \sum_{i=1}^N \mathcal{I}_i^{\beta^*}(t) = x \mathcal{I}^*(t)$$

For $s \leq t$, coordinatewise monotonicity of q^* and Lemma 1 give

$$\text{Cov}_x(S^*(s), S^*(t)) = \sum_{i=1}^N \text{Cov}_x(S_i^{\beta^*}(s), S_i^{\beta^*}(t)) = \sum_{i=1}^N \mathcal{I}_i^{\beta^*}(s) = \mathcal{I}^*(s)$$

Therefore, S^* is a continuous Gaussian process with mean $x \mathcal{I}^*(t)$ and covariance $\mathcal{I}^*(s \wedge t)$. For $s \in [0, \mathcal{I}^*(\infty)]$, define $\kappa^*(s) := (\mathcal{I}^*)^{-1}(s)$, which is strictly increasing in t . Then, define

the **aggregate score in information flow** as $C^*(s) := S^*(\kappa^*(s))$. We have

$$\mathbb{E}_x[C^*(s)] = \mathbb{E}_x[S^*(\kappa^*(s))] = x\mathcal{I}^*(\kappa^*(s)) = xs$$

For $s \leq r$ (hence, $\kappa^*(s) \leq \kappa^*(r)$), we have

$$\text{Cov}_x(C^*(s), C^*(r)) = \text{Cov}_x(S^*(\kappa^*(s)), S^*(\kappa^*(r))) = \mathcal{I}^*(\kappa^*(s)) = s$$

Hence, C^* is a continuous Gaussian process with mean xs and covariance $s \wedge r$, that is, $C^*(t) = xs + \tilde{W}_s$ where $\tilde{W}$ is standard Brownian motion. In short, C^* is a random variable mapping from Ω to $C([0, \infty], \mathbb{R})$, with standard Wiener measure as it's law when $\mu = 0$.

Step 4. Work under $P_0^{\tilde{\beta}}$. By Lemma 1, the aggregate score $\tilde{S}$ is a continuous local martingale with quadratic variation $\tilde{\mathcal{I}}(t)$. Again, for $s \in [0, \tilde{\mathcal{I}}(\infty)]$, define the inverse of information flow $\tilde{\kappa}(s) := \inf\{t \geq 0 : \tilde{\mathcal{I}}(t) > s\}$. Since $\tilde{\mathcal{I}}(t)$ is continuous and strictly increasing, this is the ordinary inverse on $[0, \tilde{\mathcal{I}}(\infty))$. For $s \in [0, \tilde{\mathcal{I}}(\infty)]$, define $\tilde{C}(s) := \tilde{S}(\tilde{\kappa}(s))$. By the Dambis-Dubins-Schwarz theorem, $\tilde{C}(s)$ is Brownian up to the random information time $\mathcal{I}^\beta(\infty)$. To be comparable with C^* , $\tilde{C}$ needs to be defined over $[0, \mathcal{I}^*(\infty)]$. Hence, if $[0, \tilde{\mathcal{I}}(\infty)] \subset [0, \mathcal{I}^*(\infty)]$, we introduce a standard Brownian motion $\bar{W}$ independent of $\mathcal{G}^{\tilde{\beta}}$ and define

$$\tilde{C}_s := \begin{cases} \tilde{S}(\kappa(s)), & s < \tilde{\mathcal{I}}(\infty), \\ \tilde{S}(\infty) + B_{s-\tilde{\mathcal{I}}(\infty)}, & s \geq \tilde{\mathcal{I}}(\infty). \end{cases}$$

We enlarge the time-changed filtration by revealing the auxiliary Brownian path only as it is used after $\tilde{\mathcal{I}}(\infty)$, and then take the usual augmentation. Thus, at information time s , the additional information consists of the path of B up to time $s - \tilde{\mathcal{I}}(\infty)$. The time-change theorem and independence of B ensure that the extended process is a continuous local martingale in this filtration with quadratic variation s , and hence a standard Brownian motion by Lévy's characterisation. As a result, similar to C^* when $\mu = 0$, $\tilde{C}$ is also a random variable mapping from Ω to $C([0, \infty], \mathbb{R})$, with standard Wiener measure as it's law.

Step 5. We are now ready to invoke Lemma 6. Let $U = \mathcal{I}^*(\infty)$, we consider the

product space $C([0, U], \mathbb{R}) \times \mathbb{H}$. Let $J_0^{\tilde{\beta}}$ denote the joint law of $(\tilde{C}, \tilde{H}) \in C([0, U], \mathbb{R}) \times \mathbb{H}$, where $\tilde{C}$ is given by step 4 under $P_0^{\tilde{\beta}}$ and $\tilde{H}$ is the observable history induced by $\tilde{\beta}$. Then, $\tilde{C} : [0, U] \rightarrow \mathbb{R}$ is a standard Brownian motion under $J_0^{\tilde{\beta}}$, defined relatively to a right-continuous filtration $\mathcal{G}_{\kappa(s)}^{\tilde{\beta}}$. Moreover, let $\rho = \mathcal{I}^* : [0, +\infty) \rightarrow [0, U)$, which we have verified being a deterministic, continuous, strictly increasing function, with $\rho(0) = 0$, $\lim_{t \rightarrow +\infty} \rho(t) = U$ and $\tilde{\mathcal{I}}(t) \leq \rho(t)$ for all t .

Next we want to show $\sigma(\tilde{H}_{\leq t}) \subseteq \mathcal{G}_{\rho(t)}$. Since $\tilde{\mathcal{I}}(t)$ is continuous and strictly increasing,

$$\{\tilde{\mathcal{I}}(t) \leq s\} = \{\kappa(s) \geq t\}.$$

The event on the right-hand side belongs to $\mathcal{G}_{\kappa(s)}^{\tilde{\beta}}$, thus $\tilde{\mathcal{I}}(t)$ is a $\mathcal{G}_{\kappa(s)}^{\tilde{\beta}}$ -stopping time. Now fix a $s \geq 0$. For an event $B \in \mathcal{G}_t^{\tilde{\beta}}$ suppose $\kappa(s) < t$, then

$$B \cap \{\tilde{\mathcal{I}}(t) \leq s\} = B \cap \{\kappa(s) \geq t\} = \emptyset \in \mathcal{G}_{\kappa(s)}^{\tilde{\beta}}.$$

On the other hand, suppose $\kappa(s) \geq t$, we define $E := B \cap \{\kappa(s) \geq t\}$. For a $r > 0$

$$E \cap \{\kappa(s) \leq r\} = \begin{cases} \emptyset & r < t, \\ B \cap \{t \leq \kappa(s) \leq r\} & r \geq t. \end{cases}$$

Clearly the empty set is in $\mathcal{G}_r^{\tilde{\beta}}$. Moreover, as $\mathcal{G}_t^{\tilde{\beta}} \subseteq \mathcal{G}_r^{\tilde{\beta}}$ the intersection of B and $\{t \leq \kappa(s) \leq r\}$ are also in $\mathcal{G}_r^{\tilde{\beta}}$. Since this is true for all r , $E \in \mathcal{G}_{\kappa(s)}^{\tilde{\beta}}$. Thus for all s

$$B \cap \{\tilde{\mathcal{I}}(t) \leq s\} = B \cap \{\kappa(s) \geq t\} \in \mathcal{G}_{\kappa(s)}^{\tilde{\beta}}.$$

Because $\tilde{\mathcal{I}}(t)$ is a $\mathcal{G}_{\kappa(s)}^{\tilde{\beta}}$ -stopping time, $B \in \mathcal{G}_{\tilde{\mathcal{I}}(t)}^{\tilde{\beta}}$. Since $\tilde{\mathcal{I}}(t) \leq \rho(t)$, we establish $\sigma(\tilde{H}_{\leq t}) \subseteq \mathcal{G}_t^{\tilde{\beta}} \subseteq \mathcal{G}_{\rho(t)}^{\tilde{\beta}}$. Hence, by Lemma 6, there exists an adaptive kernel $\bar{K} : C([0, U], \mathbb{R}) \times \mathcal{B}(\mathbb{H}) \rightarrow [0, 1]$ defined by¹³

$$\bar{K}(c, A) := J_0^{\tilde{\beta}}(\tilde{H} \in A \mid \tilde{C} = c) \quad (35)$$

¹³We suppress the superscript $\tilde{\beta}$ on $\bar{K}$ to simplify notation, however the kernel still depends on $\tilde{\beta}$.

For each t , there is $\bar{K}_t(c, A) = \bar{K}(c_{\leq \rho(t)}, A_{\leq t})$. In other words, the restriction of this kernel to the space of paths before t has a version depending only on $c_{\leq \rho(t)}$.

Given β^* , we define $K(h^*, A) := \bar{K}(c^*, A)$, where c^* is the aggregate score of information flow and h^* denotes the realisation of a history when $\beta = \beta^*$. Notice that, when the attention allocation strategy is β^* , $c_{\leq \mathcal{I}^*(t)}^*$ is observed by time t , which is just $c_{\leq \rho(t)}^*$. Hence, this is a same-calendar adapted kernel in the sense of Definition 2. It remains to prove that the kernel, although constructed under $\mu = 0$, reproduces the target law under every $\mu = x$.

Step 6. Under $P_0^{\tilde{\beta}}$, Lemma 2 gives the finite-date ($t \leq T$) likelihood $L_t(x)$ of $\mu = x$ relative to $\mu = 0$. Under the experiment induced by $\tilde{\beta}$, this is $\tilde{L}_t(x) = \exp\left\{x\tilde{S}(t) - \frac{x^2}{2}\tilde{\mathcal{I}}(t)\right\}$. Under the experiment induced by β^* , this is $L_t^*(x) = \exp\left\{xS^*(t) - \frac{x^2}{2}\mathcal{I}^*(t)\right\}$. Under the reference Wiener law $\mathbf{W}_0$, define $\mathcal{E}_s(x) := \exp\left\{xC_s - \frac{x^2}{2}s\right\}$ for $s \in [0, U]$. The standard Brownian motions C^* (from step 3) and $\tilde{C}$ (from step 4) are distinct random elements, defined with respect to different filtration. However, on the target space $C([0, U], \mathbb{R})$ they both admit Wiener law $\mathbf{W}_0$.

Fix $x \in \mathbb{R}$, a finite t , and $A \in \mathbb{H}_t$. The random time $\tilde{\mathcal{I}}(t)$ is bounded by the deterministic finite number $I^*(t)$. Therefore, bounded optional sampling gives

$$\mathbb{E}_{J_0^{\tilde{\beta}}} \left[\mathcal{E}_{\mathcal{I}^*(t)}(x) \mid \mathcal{G}_{\tilde{\mathcal{I}}(t)}^{\tilde{\beta}} \right] = \mathcal{E}_{\tilde{\mathcal{I}}(t)}(x) \quad (36)$$

Then, we have

$$\begin{aligned} (P_x^{\beta^*} K)(A) &= \int \bar{K}(c, A) \mathcal{E}_{\mathcal{I}^*(t)}(x) \mathbf{W}_0(\mathrm{d}c) \\ &= \mathbb{E}_{J_0^{\tilde{\beta}}} \left[\mathbf{1}_{\{\tilde{H} \in A\}} \mathcal{E}_{\mathcal{I}^*(t)}(x) \right] \\ &= \mathbb{E}_{J_0^{\tilde{\beta}}} \left[\mathbf{1}_{\{\tilde{H} \in A\}} \mathbb{E}_{J_0^{\tilde{\beta}}} [\mathcal{E}_{\mathcal{I}^*(t)}(x) \mid \mathcal{G}_{\tilde{\mathcal{I}}(t)}^{\tilde{\beta}}] \right] \\ &= \mathbb{E}_{J_0^{\tilde{\beta}}} \left[\mathbf{1}_{\{\tilde{H} \in A\}} \mathcal{E}_{\tilde{\mathcal{I}}(t)}(x) \right] \\ &= P_x^{\beta}(A) \end{aligned} \quad (37)$$

The first equality follows from $K^{\beta}(h^*, A) = \bar{K}^{\beta}(C^*(h^*), A)$ and the adaptedness supplied

by Lemma 6: for $A \in \mathbb{H}_t$, the kernel evaluation $\bar{K}^\beta(c, A)$ has a version depending only on $c_{\leq I^*(t)}$. The law of this process under state x has density $\mathcal{E}_{I^*(t)}(x)$ relative to its law under state zero. The second equality is the disintegration identity 35. Since $\{\tilde{H} \in A\} \in \mathcal{G}_{\tilde{I}(t)}^\beta$, the third equality follows by conditioning, and the fourth by bounded optional sampling in 36. Finally, $\tilde{C}_{\tilde{I}(t)} = \tilde{S}(t)$ implies $\mathcal{E}_{\tilde{I}(t)}(x) = \tilde{L}_t(x)$, so the likelihood formula above gives the last equality. Integrating this statewise identity against any Borel prior π_0^μ establishes the result for every prior, with the same kernel K .

Fix any admissible stopping-and-action plan under $\tilde{\beta}$. Composing this plan with the adapted garbling kernel specifies a randomised stopping-and-action rule under β^* : the probability of stopping by any date and choosing an action in any given set depends only on observations available by that date. The representation of randomised stopping times, together with an adapted conditional distribution of the terminal action, permits implementation using private randomisation. Since the kernel reproduces the competing history law in every state, the transferred plan preserves the joint distribution of the state, stopping time, and terminal action, and therefore yields the same expected payoff. Restricting attention to β^* consequently entails no loss in decision value. ■

C Proofs of Comparative Statics, Optimal Stopping and Optimal Attention Allocation with Gaussian Prior

C.1 Proofs of Results from Subsection 5.1

To prove Proposition 3 and Proposition 4, we prove four separate claims.

Claim 1. Fix $t \in (T_m, T_{m+1})$ and let $i \in \{1, \dots, m\}$. Then, there exists a $\hat{\sigma}_i > 0$ such that $\frac{\partial q_i^*(t)}{\partial \sigma_i} \leq 0$ if and only if $\sigma_i \geq \hat{\sigma}_i$.

Proof of Claim 1. Fix $t \in (T_m, T_{m+1})$ and let $i \in \{1, \dots, m\}$. From Proposition 1, the cumulative attention allocated to source i until time t is

$$q_i^*(t) = \frac{\sigma_i}{v_i^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} - \sigma_i \right)$$

Since $t \in (T_m, T_{m+1})$, a marginal change in σ_i (which marginally affects T_m and T_{m+1}) will not change the fact that $t \in (T_m, T_{m+1})$. Hence, the t -consumption set does not vary and

$$\begin{aligned}\frac{\partial q_i^*(t)}{\partial \sigma_i} &= \frac{1}{v_i^2} \left(1 - \frac{\sigma_i/v_i^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right) \left(\frac{t + \sum_{k=1}^m \sigma_k^2/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} - 2\sigma_i \right) \\ &= \frac{1}{v_i^2} \left(1 - \frac{\sigma_i/v_i^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right) \left(\frac{t + \sum_{k=1}^m \sigma_k(\sigma_k - 2\sigma_i)/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right) \\ &= \frac{1}{v_i^2} \left(1 - \frac{\sigma_i/v_i^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right) \left(\frac{t + \sum_{k \neq i} \sigma_k^2/v_k^2 - 2\sigma_i \sum_{k \neq i} \sigma_k/v_k^2 - \sigma_i^2/v_i^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right)\end{aligned}$$

Given the positivity of the first two terms and of the denominator of the third term, the sign of $\frac{\partial q_i^*(t)}{\partial \sigma_i}$ is the same as the sign of $\xi(\sigma_i) := t + \sum_{k \neq i} \sigma_k^2/v_k^2 - 2\sigma_i \sum_{k \neq i} \sigma_k/v_k^2 - \sigma_i^2/v_i^2$.

Notice that $\xi(0) = t + \sum_{k \neq i} \sigma_k^2/v_k^2 > 0$, $\xi(\sigma_i) \rightarrow -\infty$ as $\sigma_i \rightarrow +\infty$, and $\frac{\partial \xi(\sigma_i)}{\partial \sigma_i} < 0$. This means that there exists a unique point $\hat{\sigma}_i > 0$ such that: if $\sigma_i < \hat{\sigma}_i$, then $\frac{\partial q_i^*(t)}{\partial \sigma_i} > 0$; if $\sigma_i = \hat{\sigma}_i$, then $\frac{\partial q_i^*(t)}{\partial \sigma_i} = 0$; if $\sigma_i > \hat{\sigma}_i$, then $\frac{\partial q_i^*(t)}{\partial \sigma_i} < 0$.

It is possible to solve for $\hat{\sigma}_i$ from $\xi(\hat{\sigma}_i) = 0$. The second-degree equation has only one admissible root (since $\hat{\sigma}_i > 0$):

$$\hat{\sigma}_i = v_i^2 \left[\sqrt{\left(\sum_{k \neq i} \frac{\sigma_k}{v_k^2} \right)^2 + \frac{1}{v_i^2} \left(t + \sum_{k \neq i} \frac{\sigma_k^2}{v_k^2} \right)} - \sum_{k \neq i} \frac{\sigma_k}{v_k^2} \right]$$

■

Claim 2. Fix $t \in (T_m, T_{m+1})$ and let $i \in \{1, \dots, m\}$. Then, $\frac{\partial q_i^*(t)}{\partial v_i} \leq 0$.

Proof of Claim 2. Fix $t \in (T_m, T_{m+1})$ and let $i \in \{1, \dots, m\}$. From Proposition 1, the cumulative attention allocated to source i until time t is

$$q_i^*(t) = \frac{\sigma_i}{v_i^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} - \sigma_i \right)$$

Since $t \in (T_m, T_{m+1})$, a marginal change in v_i (which marginally affects T_m and T_{m+1}) will

not change the fact that $t \in (T_m, T_{m+1})$. Hence, the t -consumption set does not vary and

$$\frac{\partial q_i^*(t)}{\partial v_i} = -\frac{2\sigma_i}{v_i^3} \left(1 - \frac{\sigma_i/v_i^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right) \left(\frac{t + \sum_{k=1}^m \sigma_k^2/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} - \sigma_i \right)$$

It is immediate to see that the first term is negative and the second one is positive. To see that the third term is non-negative, suppose by contradiction that it is negative. Notice that the negativity of the third term can be written as follows:

$$\frac{t + \sum_{k=1}^m \sigma_k^2/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} - \sigma_i < 0 \iff t < \sum_{k=1}^m \frac{\sigma_k(\sigma_i - \sigma_k)}{v_k^2}$$

Since $i \in \{1, \dots, m\}$, $\sigma_i \leq \sigma_m$. Hence,

$$\sum_{k=1}^m \frac{\sigma_k(\sigma_i - \sigma_k)}{v_k^2} \leq \sum_{k=1}^m \frac{\sigma_k(\sigma_m - \sigma_k)}{v_k^2} = \sum_{k=1}^{m-1} \frac{\sigma_k(\sigma_m - \sigma_k)}{v_k^2} = T_m$$

Hence, we would have $t < T_m$, a contradiction with the fact that $t \in (T_m, T_{m+1})$. We can then conclude that $\frac{\partial q_i^*(t)}{\partial v_i}$ is always non-positive. ■

Claim 3. Fix $t \in (T_m, T_{m+1})$ and let $i, j \in \{1, \dots, m\}$ with $i \neq j$. Then, there exists a $\hat{\sigma}_i > 0$ such that $\frac{\partial q_j^*(t)}{\partial \sigma_i} \geq 0$ if and only if $\sigma_i \geq \hat{\sigma}_i$.

Proof of Claim 3. Fix $t \in (T_m, T_{m+1})$ and let $i, j \in \{1, \dots, m\}$ with $i \neq j$. From Proposition 1, the cumulative attention allocated to source j until time t is

$$q_j^*(t) = \frac{\sigma_j}{v_j^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} - \sigma_j \right)$$

Since $t \in (T_m, T_{m+1})$, a marginal change in σ_i (which marginally affects T_m and T_{m+1}) will not change the fact that $t \in (T_m, T_{m+1})$. Hence, the t -consumption set does not vary and

$$\frac{\partial q_j^*(t)}{\partial \sigma_i} = \frac{\sigma_j}{v_j^2 v_i^2} \frac{1}{\sum_{k=1}^m \sigma_k/v_k^2} \left(2\sigma_i - \frac{t + \sum_{k=1}^m \sigma_k^2/v_k^2}{\sum_{k=1}^m \sigma_k/v_k^2} \right)$$

It is immediate to see that the first two terms are positive, and hence the sign of $\frac{\partial q_j^*(t)}{\partial \sigma_i}$ is

determined by the sign of the third term. We can also notice that the third term is the negative of the third term studied in Claim 1. We can therefore immediately conclude that there exists a unique point $\hat{\sigma}_i > 0$ (the same as in Claim 1) such that: if $\sigma_i < \hat{\sigma}_i$, then $\frac{\partial q_j^*(t)}{\partial \sigma_i} < 0$; if $\sigma_i = \hat{\sigma}_i$, then $\frac{\partial q_j^*(t)}{\partial \sigma_i} = 0$; if $\sigma_i > \hat{\sigma}_i$, then $\frac{\partial q_j^*(t)}{\partial \sigma_i} > 0$. ■

Claim 4. Fix $t \in (T_m, T_{m+1})$ and let $i, j \in \{1, \dots, m\}$ with $i \neq j$. Then, $\frac{\partial q_j^*(t)}{\partial v_i} \geq 0$.

Proof of Claim 4. Fix $t \in (T_m, T_{m+1})$ and let $i, j \in \{1, \dots, m\}$ with $i \neq j$. From Proposition 1, the cumulative attention allocated to source j until time t is

$$q_j^*(t) = \frac{\sigma_j}{v_j^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} - \sigma_j \right)$$

Since $t \in (T_m, T_{m+1})$, a marginal change in v_i (which marginally affects T_m and T_{m+1}) will not change the fact that $t \in (T_m, T_{m+1})$. Hence, the t -consumption set does not vary and

$$\frac{\partial q_j^*(t)}{\partial v_i} = \frac{2\sigma_i\sigma_j}{v_i^3 v_j^2} \frac{1}{\sum_{k=1}^m \sigma_k / v_k^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} - \sigma_i \right)$$

It is immediate to see that the first two terms are positive. To see that the third term is non-negative, we can apply the same reasoning of Claim 2. We can then conclude that $\frac{\partial q_j^*(t)}{\partial v_i}$ is always non-negative. ■

Putting together the above claims yields Proposition 3 and 4.

Proof of Proposition 3 & 4. The result follows from Claims 1-4. ■

C.2 Proofs of Results in Subsection 5.3

Proof of Corollary 2. Suppose $\beta_{i+1}^*(t) \leq \beta_i^*(t)$. Then, $\sigma_{i+1}/v_{i+1}^2 \leq \sigma_i/v_i^2$. Since $\sigma_i \leq \sigma_{i+1}$, we conclude that $v_i^2 \leq v_{i+1}^2$, which contradicts the fact that the set is mutually undominated. ■

Proof of Corollary 3. Suppose $i < j$, so that $\sigma_i \leq \sigma_j$ and $T_i \leq T_j$. There are three relevant cases. First, if $t < T_i \leq T_j$, then $q_i^*(t) = q_j^*(t) = 0$. Second, if $T_i \leq t < T_j$, then $q_i^*(t) > 0$

and $q_j^*(t) = 0$, so $q_i^*(t) > q_j^*(t)$. Third, we have the case where $T_i \leq T_j \leq t$. In this third case, for any $m > j$, we have

$$q_i^*(t) = \frac{\sigma_i}{v_i^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} - \sigma_i \right) \quad \text{and} \quad q_j^*(t) = \frac{\sigma_j}{v_j^2} \left(\frac{t + \sum_{k=1}^m \sigma_k^2 / v_k^2}{\sum_{k=1}^m \sigma_k / v_k^2} - \sigma_j \right)$$

for $t \in [T_m, T_{m+1})$. When $t \in [T_m, T_{m+1})$, we can compute the difference as follows:

$$\begin{aligned} q_j^*(t) - q_i^*(t) &= \frac{\frac{\sigma_j}{v_j^2} \left(t + \sum_{k=1}^m \frac{\sigma_k(\sigma_k - \sigma_j)}{v_k^2} \right) - \frac{\sigma_i}{v_i^2} \left(t + \sum_{k=1}^m \frac{\sigma_k(\sigma_k - \sigma_i)}{v_k^2} \right)}{\sum_{k=1}^m \frac{\sigma_k}{v_k^2}} \\ &= \left(\frac{\frac{\sigma_j}{v_j^2} - \frac{\sigma_i}{v_i^2}}{\sum_{k=1}^m \frac{\sigma_k}{v_k^2}} \right) t + \left[\frac{\frac{\sigma_j}{v_j^2} \left(\sum_{k=1}^m \frac{\sigma_k(\sigma_k - \sigma_j)}{v_k^2} \right) - \frac{\sigma_i}{v_i^2} \left(\sum_{k=1}^m \frac{\sigma_k(\sigma_k - \sigma_i)}{v_k^2} \right)}{\sum_{k=1}^m \frac{\sigma_k}{v_k^2}} \right] \end{aligned}$$

Within $[T_m, T_{m+1})$, the component in square brackets is constant. Since the set is undominated and $\sigma_i \leq \sigma_j$, we have $v_i > v_j$ and hence $\frac{\sigma_j}{v_j^2} > \frac{\sigma_i}{v_i^2}$. Hence, the coefficient of t is strictly positive. Therefore, $q_j^*(t) - q_i^*(t)$ is strictly increasing in t for $t \in [T_m, T_{m+1})$. Moreover, both $q_i^*(t)$ and $q_j^*(t)$ are continuous in t , hence there are no jumps at the time thresholds. We can then conclude that $q_j^*(t) - q_i^*(t)$ is strictly increasing in t for $t \geq T_j$. Since $q_j^*(T_j) - q_i^*(T_j) < 0$, and $q_j^*(t) - q_i^*(t) \rightarrow +\infty$ as $t \rightarrow +\infty$, there exists a unique (in $[T_j, +\infty)$) point $\tilde{t}_{ij}$ such that $q_j^*(\tilde{t}_{ij}) - q_i^*(\tilde{t}_{ij}) = 0$.

To conclude, notice that the first and the second case verify that, for $t < T_j$, we have $q_i^*(t) \geq q_j^*(t)$. On the other hand, the third case verifies that, for $t \geq T_j$, $q_i^*(t) < q_j^*(t) \iff t > \tilde{t}_{ij}$. Putting the three cases together, we conclude that $q_i^*(t) < q_j^*(t) \iff t > \tilde{t}_{ij}$. ■

C.3 Proof of Results in Section 6

Let $u(a, \mu) = f(a - \mu) + \zeta$, where ζ is some constant and $f : \mathbb{R} \rightarrow \mathbb{R}$ is defined as in Section 6. To establish Proposition 5 we first prove a lemma about the gross value of the decision problem.¹⁴

¹⁴To simplify notation we drop this constant in proofs.

Lemma 7. Let $u \in \mathcal{U}$. The gross value $\pi^\mu \mapsto g(\pi^\mu) = \max_a \mathbb{E}_{\pi^\mu}[u(a, \mu)]$ is a decreasing function of the variance σ_π^2 and it is independent of the mean.

Proof of Lemma 7. Regarding the gross value $g(\pi^\mu)$ we have

$$\begin{aligned} g(\pi^\mu) &= \max_a \mathbb{E}_{\pi^\mu}[u(a, \mu)] \\ &= \sup_a \int_{\mathbb{R}} f(a - \mu) k_\pi(\mu) d\mu \\ &= \sup_a (f * k_\pi)(a) \end{aligned}$$

where k_π is the density of π^μ and $(f * k_\pi)$ is the convolution of f and k_π .

Now consider two normal distributions $\pi_1 = N(\eta_{\pi_1}, \sigma_{\pi_1}^2)$ and $\pi_2 = N(\eta_{\pi_2}, \sigma_{\pi_2}^2)$ such that $\pi_1, \pi_2 \in \Gamma$ and $\sigma_{\pi_1}^2 < \sigma_{\pi_2}^2$. Since π_1 has a smaller variance, we can define $\sigma_{\pi_3} := \sqrt{\sigma_{\pi_2}^2 - \sigma_{\pi_1}^2}$. We know that the sum of two independent normal random variables is still normal, with variance equal to the sum of variances. Thus there exist a normal distribution π_3 of which the variance is $\sigma_{\pi_3}^2$ satisfying $k_{\pi_2} = k_{\pi_1} * k_{\pi_3}$. Since convolutions form a semigroup, namely the operator $*$ respects associativity, we have

$$\begin{aligned} g(\pi_2^\mu) &= \max_a \mathbb{E}_{\pi_2^\mu}[u(a, \mu)] \\ &= \sup_a \int_{\mathbb{R}} f(a - \mu) k_{\pi_2}(\mu) d\mu \\ &= \sup_a (f * k_{\pi_2})(a) \\ &= \sup_a (f * (k_{\pi_1} * k_{\pi_3}))(a) \\ &= \sup_a ((f * k_{\pi_1}) * k_{\pi_3})(a) \\ &= \sup_a \int_{\mathbb{R}} (f * k_{\pi_1})(a - \mu) k_{\pi_3}(\mu) d\mu \end{aligned}$$

Now since k_{π_3} is a probability density, one can see that

$$\int_{\mathbb{R}} (f * k_{\pi_1})(a - \mu) k_{\pi_3}(\mu) d\mu \leq \int_{\mathbb{R}} \sup_x \{f * k_{\pi_1}\}(x) k_{\pi_3}(\mu) d\mu = \sup_x \{f * k_{\pi_1}\}(x) = g(\pi_1^\mu)$$

where $x := a - \mu$. Notice the inequality is strict if f is not almost surely constant. Therefore

$g(\pi^\mu)$ is a decreasing function in the variance.

Next we show that $g(\pi^\mu)$ is independent of the mean. Notice we can rewrite $g(\pi^\mu)$ as follow:

$$g(\pi^\mu) = \max_a \mathbb{E}_{\pi^\mu}[u(a, \mu)] = \max_a \mathbb{E}_{\pi^\mu}[f(a - \mu)] = \max_a \mathbb{E}_{\pi^\mu}[f(a - \eta_\pi + \eta_\pi - \mu)]$$

where $\pi^\mu = N(\eta_\pi, \sigma_\pi^2)$. It is clear that $\mu - \eta_\pi \sim N(0, \sigma_\pi^2)$. Now let $y := a - \eta_\pi$ and $v := \mu - \eta_\pi$, we have

$$g(\pi^\mu) = \max_y \mathbb{E}_{\pi^v}[f(y - v)]$$

Since π^v is a normal distribution with mean zero and normal distributions can be parametrised by their means and variances, $g(\pi^\mu)$ is a function of the variance σ_π^2 only. ■

Given the result in Lemma 7, we are now ready to prove Proposition 5.

Proof of Proposition 5. Given Proposition 2, we know that β^* solves the attention allocation problem. Hence, we now consider the optimal stopping time given this attention allocation scheme. Fix a flow cost $c > 0$. By Lemma 7, $g(\pi_t^\mu)$ is a decreasing function in the posterior variance $\sigma_{\mu,t}^2$ of π_t^μ . Moreover, $\sigma_{\mu,t}^2$ is a decreasing function in t due to Gaussian assumptions. Thus, given the optimal attention allocation scheme β^* the value $g(\pi_t^\mu)$ is an increasing function in t , denoted by $G(t)$. Then we have

$$\begin{aligned} V(\pi_0) &= \max_\tau \mathbb{E}_{\pi_0, \beta^*}[g(\pi_\tau) - c\tau] \\ &= \max_\tau \mathbb{E}_{\pi_0}[G(\tau) - c\tau] \end{aligned}$$

Since $G(t)$ is bounded, the earliest stopping time is well defined. Also, one can see that τ^* must be deterministic as the optimisation problem given β^* is essentially deterministic.

For the monotonicity of τ^* with respect to the flow cost, consider $0 < c_1 < c_2$ and let τ_1^* and τ_2^* be the corresponding earliest stopping times. In addition, we denote the value of the decision maker when the flow cost is c and they stop at τ by $V_c(\tau)$. To build a contradiction, suppose $\tau_2^* > \tau_1^*$. If the decision maker stops at τ_2^* when the flow cost is c_1 ,

we have

$$\begin{aligned}
V_{c_1}(\tau_2^*) &= G(\tau_2^*) - c_1 \tau_2^* \\
&= G(\tau_2^*) - G(\tau_1^*) + G(\tau_1^*) - c_1(\tau_2^* - \tau_1^*) - c_1 \tau_1^* \\
&= G(\tau_2^*) - G(\tau_1^*) - c_1(\tau_2^* - \tau_1^*) + G(\tau_1^*) - c \tau_1^* \\
&= G(\tau_2^*) - G(\tau_1^*) - c_1(\tau_2^* - \tau_1^*) + V_{c_1}(\tau_1^*) \\
&> G(\tau_2^*) - G(\tau_1^*) - c_2(\tau_2^* - \tau_1^*) + V_{c_1}(\tau_1^*) \\
&= V_{c_2}(\tau_2^*) - V_{c_2}(\tau_1^*) + V_{c_1}(\tau_1^*)
\end{aligned}$$

Since by definition τ_2^* is the earliest stopping time when the flow cost is c_2 , $V_{c_2}(\tau_2^*) - V_{c_2}(\tau_1^*) > 0$. Therefore $V_{c_1}(\tau_2^*) > V_{c_1}(\tau_1^*)$, which is a contradiction. Hence the earliest stopping time is decreasing with respect to the flow cost. ■

C.4 Proofs of Results from Subsection 7.1

Throughout this section, let

$$\tilde{\mu} := \mu - \mu_0 \quad \tilde{b}_i := b_i - b_{i0} \quad \theta_i := \frac{\tilde{\mu} + \tilde{b}_i}{\sigma_i} \quad (38)$$

Proof of Lemma 3. We want to show that ω is the conditional expectation when θ is known. To simplify notation in this proof let $Z_i := \sigma_i \theta_i$. Conditional on the vector $z = (z_1, \dots, z_N)$, which is normally distributed, x follows a normal distribution as well and Bayes' rule gives the following:

$$p(\tilde{\mu} \mid z) \propto \exp \left\{ -\frac{(\tilde{\mu})^2}{2\sigma_\mu^2} - \frac{1}{2} \sum_{i=1}^N \frac{(Z_i - \tilde{\mu})^2}{v_i^2} \right\}.$$

Expand the exponent and collect its quadratic and linear terms in x :

$$-\frac{1}{2} \left(\frac{1}{\sigma_\mu^2} + \sum_{i=1}^N \frac{1}{v_i^2} \right) (\tilde{\mu})^2 + \tilde{\mu} \sum_{i=1}^N \frac{Z_i}{v_i^2} + C(z),$$

where $C(z)$ is independent of $\tilde{\mu}$. Completing the square and using the definition of φ_∞ yields

$$\tilde{\mu} \mid z \sim N\left(\varphi_\infty \sum_{i=1}^N \frac{Z_i}{v_i^2}, \varphi_\infty\right).$$

Since $Z_i = \sigma_i \theta_i$,

$$\varphi_\infty \sum_{i=1}^N \frac{Z_i}{v_i^2} = \sum_{i=1}^N \left(\varphi_\infty \frac{\sigma_i}{v_i^2}\right) \theta_i = \alpha^\top \theta = \omega.$$

This proves $\mu \mid \theta \sim N(\mu_0 + \alpha^\top \theta, \varphi_\infty)$.

Let $\varepsilon := \mu - \mu_0 - \omega$. Substituting in the definition of ω and θ_i ,

$$\varepsilon = \varphi_\infty \left(\frac{\tilde{\mu}}{\sigma_\mu^2} - \sum_{i=1}^N \frac{\tilde{b}_i}{v_i^2} \right).$$

which implies that $\text{Var}(\varepsilon) = \varphi_\infty$. For every $j \in \{1, \dots, N\}$,

$$\begin{aligned} \text{Cov}(\varepsilon, \theta_j) &= \frac{\varphi_\infty}{\sigma_j} \text{Cov}\left(\frac{\tilde{\mu}}{\sigma_\mu^2} - \sum_i \frac{\tilde{b}_i}{v_i^2}, \tilde{\mu} + \tilde{b}_j\right) \\ &= \frac{\varphi_\infty}{\sigma_j} \left(\frac{\text{Var}(\tilde{\mu})}{\sigma_\mu^2} - \frac{\text{Var}(\tilde{b}_j)}{v_j^2} \right) = 0. \end{aligned}$$

The vector (ε, θ) is jointly Gaussian; for jointly Gaussian variables, zero covariance implies independence. This proves (17). ■

Proof of Lemma 4. We first establish the equivalence in the decision maker's objective.

$$\begin{aligned}
\mathbb{E}[u(a_\tau, \mu) - c\tau] &= \mathbb{E}[\mathbb{E}[u(a_\tau, \mu_0 + \omega + \varepsilon) | \mathcal{G}_\tau^\beta] - c\mathbb{E}[\tau]] \\
&= \mathbb{E}[\mathbb{E}[\mathbb{E}[u(a_\tau, \mu_0 + \omega + \varepsilon) | \mathcal{G}_\tau^\beta, \omega] | \mathcal{G}_\tau^\beta] - c\mathbb{E}[\tau]] \\
&= \mathbb{E}[\mathbb{E}[\mathbb{E}_\varepsilon[u(a_\tau, \mu_0 + \omega + \varepsilon)] | \mathcal{G}_\tau^\beta] - c\mathbb{E}[\tau]] \\
&= \mathbb{E}[\mathbb{E}[\bar{u}(a_\tau, \omega) | \mathcal{G}_\tau^\beta] - c\mathbb{E}[\tau]] \\
&= \mathbb{E}[\bar{u}(a_\tau, \omega) - c\tau] \\
&= \mathbb{E}[\tilde{u}(\tau, a_\tau, \omega)]
\end{aligned}$$

The third equality is due to $\varepsilon \perp\!\!\!\perp \omega, \mathcal{G}_\tau^\beta$ by Lemma 3. The fourth equality is by construction in Equation (18).

Since Y_{it}^β is a deterministic transformation of X_{it}^β , they induce the identical filtration, and thus any strategy (β, τ, a_τ) admissible in the original problem is also admissible in the transformed LMS problem. The value function $V(\pi_0) = \max_{\tau, \beta(t), a_\tau} \mathbb{E}_{\pi_0}[u(a_\tau, \mu) - c\tau] = V^{LMS}(\pi_0) := \max_{\tau, \beta(t), a_\tau} \mathbb{E}_{\pi_0}[\tilde{u}(\tau, a_\tau, \omega)]$, and an optimal strategy $(\beta^*, \tau^*, a_\tau^*)$ for the original problem also achieves the supremum for the transformed LMS problem.

We now verify that the transformed attribute covariance has full-rank and satisfies Assumption 4 in the LMS model, required for their results to hold for $N > 2$.

Let

$$D_\sigma := \text{diag}(\sigma_1, \dots, \sigma_N), \quad D_v := \text{diag}(v_1^2, \dots, v_N^2),$$

and let $\mathbf{1}$ denote the N -vector of ones. From (38),

$$\theta = D_\sigma^{-1}(\tilde{\mu} \mathbf{1} + \tilde{b}).$$

where $\tilde{b} = (\tilde{b}_1, \dots, \tilde{b}_i, \dots, \tilde{b}_N)$. The transformed prior over θ has mean zero. Independence of μ and b_i , where $i \in \{1, \dots, N\}$ implies

$$\text{Cov}(\theta_i, \theta_j) = \frac{\sigma_\mu^2 + \mathbf{1}_{\{i=j\}} v_i^2}{\sigma_i \sigma_j}.$$

Consequently,

$$\Sigma_\theta := \text{Var}(\theta) = D_\sigma^{-1}(D_v + \sigma_\mu^2 \mathbf{1}\mathbf{1}^\top)D_\sigma^{-1}. \quad (39)$$

For every nonzero $z \in \mathbb{R}^N$,

$$z^\top \Sigma_\theta z = \sum_{i=1}^N \frac{v_i^2}{\sigma_i^2} Z_i^2 + \sigma_\mu^2 \left(\sum_{i=1}^N \frac{Z_i}{\sigma_i} \right)^2 > 0.$$

The strict inequality uses $v_i > 0$ for every i . Hence Σ_θ has full rank.

The Sherman–Morrison–Woodbury identity applied to (39) gives

$$\Sigma_\theta^{-1} = D_\sigma \left[D_v^{-1} - \frac{D_v^{-1} \mathbf{1}\mathbf{1}^\top D_v^{-1}}{\sigma_\mu^{-2} + \sum_{k=1}^N v_k^{-2}} \right] D_\sigma. \quad (40)$$

Indeed, this follows from

$$\begin{aligned} & (D_v + \sigma_\mu^2 \mathbf{1}\mathbf{1}^\top)^{-1} \\ &= D_v^{-1} - D_v^{-1} \mathbf{1} \left(\sigma_\mu^{-2} + \mathbf{1}^\top D_v^{-1} \mathbf{1} \right)^{-1} \mathbf{1}^\top D_v^{-1}. \end{aligned}$$

For $i \neq j$, equation (40) yields

$$[\Sigma_\theta^{-1}]_{ij} = -\frac{\sigma_i \sigma_j}{v_i^2 v_j^2 (\sigma_\mu^{-2} + \sum_{k=1}^N v_k^{-2})} < 0.$$

Thus Σ_θ^{-1} has nonpositive off-diagonal entries, which is exactly LMS Assumption 4 (“Perpetual Substitutes”). Moreover, $\alpha_i = \varphi_\infty \sigma_i / v_i^2 > 0$, so the LMS condition for positive payoff weights also holds. ■

C.5 Proofs of Results from Subsection 7.2

Proof of Lemma 5. Conditional on $\mu = x$, with $x \in \{l, h\}$, the θ_i ’s are independent Gaussians:

$$\theta_i \mid (\mu = x) \sim N \left(\frac{x - \frac{h+l}{2}}{\sigma_i}, \frac{v_i^2}{\sigma_i^2} \right).$$

Let $Z_i := \frac{h-l}{2\sigma_i}$. Then $\mathbb{E}[\theta_i \mid \mu = l] = -Z_i$ and $\mathbb{E}[\theta_i \mid \mu = h] = Z_i$. The likelihood ratio for the two states is

$$\frac{f(\theta \mid \mu = h)}{f(\theta \mid \mu = l)} = \prod_{i=1}^N \frac{\exp \left\{ -\frac{(\theta_i - Z_i)^2}{2v_i^2/\sigma_i^2} \right\}}{\exp \left\{ -\frac{(\theta_i + Z_i)^2}{2v_i^2/\sigma_i^2} \right\}} = \exp \left\{ (h-l) \sum_{i=1}^N \frac{\sigma_i}{v_i^2} \theta_i \right\},$$

Bayes' rule yields

$$\begin{aligned} \mathbb{P}(\mu = h \mid \theta) &= \frac{p_0 \cdot f(\theta \mid \mu = h)}{p_0 f(\theta \mid \mu = h) + (1 - p_0) f(\theta \mid \mu = l)} \\ &= \frac{\frac{p_0}{1-p_0} \exp \left\{ (h-l) \sum_i \frac{\sigma_i}{v_i^2} \theta_i \right\}}{1 + \frac{p_0}{1-p_0} \exp \left\{ (h-l) \sum_i \frac{\sigma_i}{v_i^2} \theta_i \right\}} \\ &= \text{logit}^{-1} \left(l_0 + (h-l) \sum_{i=1}^N \frac{\sigma_i}{v_i^2} \theta_i \right) \end{aligned}$$

Since $\mathbb{E}[\mu \mid \theta] = l + (h-l)\mathbb{P}(\mu = h \mid \theta)$ and $\mu_0 = l + (h-l)p_0$, we obtain (19). ■