

Dynamically Aggregating Diverse Information*

Annie Liang[†] Xiaosheng Mu[‡] Vasilis Syrgkanis[§]

June 8, 2021

Abstract

An agent has access to multiple information sources, each modeled as a Brownian motion whose drift provides information about a different component of an unknown Gaussian state. Information is acquired continuously—where the agent chooses both which sources to sample from, and also how to allocate attention across them—until an endogenously chosen time, at which point a decision is taken. We demonstrate conditions on the agent’s prior belief under which it is possible to exactly characterize the optimal information acquisition strategy. We then apply this characterization to derive new results regarding: (1) endogenous information acquisition for binary choice, (2) the dynamic consequences of attention manipulation, and (3) strategic information provision by biased news sources.

*We are grateful to Cuimin Ba, Yash Deshpande, Mira Frick, Drew Fudenberg, Boyan Jovanovic, Erik Madsen, George Mailath, Konrad Mierendorff, Lan Min, Jacopo Peregò, Peter Norman Sørensen, Jakub Steiner, Philipp Strack, Can Urgun and anonymous referees for helpful comments and suggestions. Xiaosheng Mu acknowledges the hospitality of Columbia University and the Cowles Foundation at Yale University, which hosted him during parts of this research. Annie Liang thanks National Science Foundation Grant SES-1851629 for financial support.

[†]Northwestern University

[‡]Princeton University

[§]Microsoft Research

1 Introduction

We study dynamic acquisition of information when a decision-maker has access to multiple sources of information, and limited resources with which to acquire that information. Our decision-maker seeks to learn a Gaussian state, and each information source is modeled as a diffusion process whose drift is an unknown “attribute” that contributes linearly to the state. Attributes are potentially correlated. This structure captures information acquisition in many economic settings, including for example:

- A governor wants to learn the number of cases of a disease outbreak, and can acquire information about the incidence rate of the disease in different cities.
- An investor wants to assess the value of an asset portfolio, and acquires information about the value of each asset included in the portfolio.
- An analyst wants to forecast a macroeconomic variable such as GDP growth, and needs to aggregate recent economic activities across different industries and locations.

At every instant of continuous time, the decision-maker allocates a fixed budget of attention/resources across the information sources, where this allocation determines the precision of information extracted from each source. For example, the governor may have a limited number of tests to allocate across testing centers each day, where more tests lead to a more precise estimate of the incidence rate. The decision-maker acquires information until an endogenously chosen stopping time, at which point he makes a decision whose payoff depends on the unknown state.

Our contribution is to demonstrate that the optimal dynamic acquisition strategy can be explicitly characterized under certain conditions on the prior belief, and to explain what those conditions are. Under the optimal strategy, the decision-maker initially exclusively acquires information from the single most informative source, where “more informative” is evaluated with respect to his prior belief over the unknown attribute values. At fixed times, the decision-maker begins learning from additional sources, and divides attention over these new sources as well as the ones he was learning from previously. Eventually, the decision-maker acquires information from all sources using a final and constant mixture. Notably, the optimal information acquisition strategy is not only history-independent but also robust across all decision problems. This implies, for example, that one does not need to solve for the

optimal stopping time and information acquisition strategy jointly in this problem—optimal information acquisition is independent of when the decision-maker stops. We make use of this implication in Section 5.1 to derive new results about the optimal stopping behavior in a binary choice problem.

To gain some intuitions for the optimal information acquisition strategy, it is useful to compare our problem with a simpler one, in which the decision-maker acquires information for a decision at a fixed end date. Since the payoff-relevant state and all sources of information are Gaussian in our setting, the Blackwell-optimal solution would then be to acquire information in any way that minimizes posterior variance of the state at the known end date (Blackwell, 1951; Hansen and Torgersen, 1974). We show that under certain conditions on the prior belief, it is possible to “string together” these solutions across different end dates using a single history-independent dynamic strategy, which thus minimizes posterior variance at every moment of time. Generalizing a result of Greenshtein (1996), we show that this strategy—which we call the “uniformly optimal” strategy—is best for every decision problem and every distribution over stopping times, including those that are endogenously chosen.

When a uniformly optimal strategy does not exist, the variance-minimizing strategies for some end dates are in conflict with one another. In these environments the decision-maker must trade off across possible end dates, where the optimal way of doing this in general depends on the stopping time distribution and details of the payoff function. Thus, the existence of a uniformly optimal strategy is key to guaranteeing the properties of history independence and robustness across decision problems that we have outlined for our solution.

The question of whether a uniformly optimal strategy exists turns out to relate to a classic problem in consumer theory regarding the normality of demand—i.e., whether a consumer’s demand for various goods is weakly increasing in income. In our setting, the decision-maker’s “utility function” is the negative of the posterior variance function, and his “income” is the budget constraint on attention. When a uniformly optimal strategy exists, this means that the decision-maker’s demand for information from each source is weakly increasing in his overall attention budget. One of our sufficient conditions for existence of a uniformly optimal strategy—“perpetual complementarity” of the different sources—directly connects to a known sufficient condition in the literature for normality of demand. We additionally utilize the specific structure of our problem to provide two new sufficient conditions, and

show that all of these conditions can be simply stated in terms of the decision-maker’s prior belief. See Section 4 for further details.

Beyond the specific statements of the results, a main contribution of this paper is demonstrating that in the present framework (i) the study of endogenous information acquisition is quite tractable, permitting explicit and complete characterizations; and (ii) there is enough richness in the setting to accommodate various economically interesting questions (e.g., about the role of primitives such as correlation across attributes). This makes the characterizations useful for deriving new substantive results in settings motivated by particular economic questions. We illustrate this with three applications, where we use our main results to generalize (Application 1) and complement (Application 2) existing results in the literature, as well as to solve for the equilibrium in a new game between competing news sources (Application 3). In all three applications, we discover new economic insights.

Application 1: Binary Choice. A large literature in economics and neuroscience (originating with [Ratcliff and McKoon \(2008\)](#)) considers a consumer’s decision process for choosing between two goods with unknown payoffs. Recently, [Fudenberg et al. \(2018\)](#) proposed a model in which a decision-maker endogenously allocates attention across learning about two normally distributed, but i.i.d., payoffs. This model is nested in our framework. We use our main result to generalize [Fudenberg et al. \(2018\)](#)’s Proposition 3 and Theorem 5 beyond i.i.d. payoffs to settings with (1) correlation in the payoffs and (2) asymmetry in the consumer’s initial knowledge about the two payoffs. These generalizations bring important realism to the setting, since correlation and asymmetry are common features of choice environments. We characterize the optimal attention allocation given an arbitrary normal prior about the payoffs, and show that [Fudenberg et al. \(2018\)](#)’s main economic insight regarding the relationship between choice speed and accuracy holds in this general setting.

Application 2: Attention Manipulation. Next, we use our framework to study the dynamic consequences of a one-time attention manipulation. Recently, [Gossner et al. \(2020\)](#) studied this question in a model where a decision-maker chooses between goods with independent payoffs. Under some assumptions, they show that a one-time manipulation of attention towards a given good leads to persistently higher cumulative attention devoted to that good, and persistently lower cumulative attention to every other good. We derive a complementary result in our setting, focusing on how correlation across the unknown at-

tributes affects the consequences of attention manipulation. We show that with two sources, [Gossner et al. \(2020\)](#)’s insights hold under flexible patterns of correlation. On the other hand, with more than two sources, the nature of correlation matters. We identify a property of the prior belief under which all sources provide substitutable information, and show that under this property (but not in general), attention manipulation leads to persistently higher attention for that source and lower attention for others.

Application 3: Biased News Sources. In our final application, we consider a stylized game between a liberal and a conservative news source that report on a common unknown (e.g., the fiscal cost of a policy proposal), where their reporting is biased in opposite directions. The sources choose the size of their bias, as well as the informativeness of their reporting, and compete over readers’ attention. Using our characterization of information acquisition, we are able to derive the complete time path of readers’ attention allocations given the sources’ choices of precision and bias levels, which allows us to characterize equilibrium news provision in this model. One particular insight that emerges from this analysis is that incentives for bias not only lead to greater polarization in equilibrium, but also lead to lower-quality news provision (i.e., larger noise choices). This analysis contributes to a literature about how competition across news sources affects the quality of news ([Gentzkow and Shapiro, 2008](#); [Galperti and Trevino, 2020](#); [Chen and Suen, 2019](#); [Perego and Yuksel, 2021](#)), where our work is distinguished in considering the role of the time path of information demand.

1.1 Related Literature

Our work builds on a large literature regarding dynamic acquisition of information. One part of this literature considers choice between unconstrained information structures at entropic (or more generally, “posterior-separable”) costs, see e.g. [Yang \(2015\)](#), [Steiner et al. \(2017\)](#), [Hébert and Woodford \(2021\)](#), [Morris and Strack \(2019\)](#), and [Zhong \(2019\)](#).¹ Under this modeling approach, the cost to acquiring information depends on the decision-maker’s current belief, and is often interpreted as a mental processing cost ([Mackowiak et al., 2021](#)). In

¹It is interesting that [Steiner et al. \(2017\)](#) also showed how the solution to their dynamic problem reduces to a series of static optimizations, similar to our multi-stage characterization. However, their argument is based on the additive property of entropy and differs from ours.

contrast, a second set of papers—to which our paper belongs—models agents as dynamically allocating a fixed budget of resources across a prescribed (and finite) set of experiments, see e.g. [Che and Mierendorff \(2019\)](#), [Mayskaya \(2020\)](#), [Fudenberg et al. \(2018\)](#), [Gossner et al. \(2020\)](#), and [Azevedo et al. \(2020\)](#). These papers, and ours, assume that the cost of information is independent of what the decision-maker currently knows. We view such information costs as a better match for applications in which the cost to information acquisition is physical, e.g., a limit on the number of available COVID tests that can be administered in a given day.

Relative to this latter strand of literature, a distinguishing feature of our work is the presence of flexible correlation. Dynamic learning about correlated unknowns is generally intractable, so there has been relatively little work done in this area. An exception is a model introduced by [Callander \(2011\)](#), where the available signals are the realizations of a single Brownian motion path at different points, and the agent (or a sequence of agents) chooses myopically. This informational setting has since been extended by [Garfagnini and Strulovici \(2016\)](#), which considers the optimal experimentation strategy for a forward-looking agent with acquisition costs, and [Bardhi \(2020\)](#), which introduces a potential conflict between an agent acquiring the information and a principal making the decision. These models differ from ours in that agents can perfectly observe any of an infinite number of attributes, and the correlation structure across the attributes is derived from a primitive notion of similarity or distance. We show that in a different model with flexible correlation across a finite number of sources (and with noisy observations), it is sometimes possible to exactly characterize the optimal forward-looking solution.²

Our work additionally connects to a large literature on sequential sampling in statistics and operations research. Since the information acquisition decisions in our model are not directly linked to flow payoffs, our model does not fall under the classic multi-armed bandit framework ([Gittins, 1979](#); [Bergemann and Välimäki, 2008](#)). This feature also distinguishes our results relative to a classic literature on “learning by experimentation” ([Easley and Kiefer, 1988](#); [Aghion et al., 1991](#); [Keller et al., 2005](#)). The “best-arm identification” problem ([Bubeck et al., 2009](#); [Russo, 2016](#)) is more closely related to us, as it considers a decision-

²Also related are [Klabjan et al. \(2014\)](#) and [Sanjurjo \(2017\)](#), which study learning about multiple attributes. Besides having noisy Gaussian signals, the main distinction of our informational setting is again that we allow for correlation across attributes and focus on what this correlation implies for the optimal learning strategy.

maker who samples for a number of periods before selecting an arm and receiving its payoff. Indeed, the special case of two arms with jointly normal payoffs is nested in our framework under the case of two attributes and equal payoff weights. Our Theorem 1 thus builds on a prior result of [Frazier et al. \(2008\)](#), which showed that myopic information acquisition—or the “knowledge gradient” policy in the language of that literature—is optimal when the two arms have independent normal payoffs. Our result generalizes [Frazier et al. \(2008\)](#)’s result by allowing for correlated payoffs and a broader class of decision problems.

The best-arm identification problem between three or more arms falls outside of our framework, since payoffs in that problem depend on a multi-dimensional unknown.³ From a number of papers including [Chick and Frazier \(2012\)](#) and [Ke and Villas-Boas \(2019\)](#), it is well-understood that characterizing the optimal strategy in those problems is quite challenging (although [Frazier et al. \(2008\)](#) and [Frazier et al. \(2009\)](#) showed that the knowledge gradient policy performs well asymptotically). Our setting also involves multi-dimensional uncertainty, but we assume that the unknowns are linearly aggregated into a one-dimensional payoff-relevant variable. We show that under this restriction, exact characterization of the optimal strategy is feasible, and in fact it is the knowledge gradient policy (suitably defined in continuous time). We also discover new properties of the knowledge gradient policy in our continuous-time setting: In each of a finite number of stages, the policy attends to a fixed set of attributes with a constant ratio of attention, until this set expands and the next stage commences.

A key technical tool behind our characterization builds on a literature about the comparison of normal experiments. Following the classic work of [Blackwell \(1951\)](#), [Hansen and Torgersen \(1974\)](#) showed that in a static decision problem, different normal signals about a one-dimensional payoff-relevant state can be Blackwell-ranked based on how much they reduce the variance of the state. [Greenshtein \(1996\)](#) subsequently derived comparisons between deterministic *sequences* of conditionally independent normal signals about an unknown state. His Theorem 3.1 implies that one sequence Blackwell-dominates another if and only if it leads to lower posterior variances about the state at every time. Our Lemma 5 shows that [Greenshtein \(1996\)](#)’s characterization is valid in a more general setting, in which time is continuous, and the sequence of signals can be chosen in a history-dependent manner (i.e.,

³With two arms, the difference in their payoffs is a sufficient statistic for choosing which arm is better. Such reduction to a one-dimensional unknown is not available in the case of many arms.

the first signal’s realization can determine which signal is chosen next).

2 Model

An agent has access to $K \geq 2$ *information sources*, each of which is a diffusion process that provides information about an unknown *attribute* $\theta_i \in \mathbb{R}$. The random vector $(\theta_1, \dots, \theta_K)$ is jointly normal with a known prior mean vector μ and prior covariance matrix Σ . We assume Σ has full rank, so the attributes are linearly independent.

As we describe in more detail below, the agent’s decision depends on a *payoff-relevant state* $\omega \in \mathbb{R}$. We assume the state is a linear combination of the attributes:

Assumption 1. $\omega = \alpha_1\theta_1 + \dots + \alpha_K\theta_K$ for known weights $\alpha_1, \dots, \alpha_K \in \mathbb{R}$.

It is equivalent (up to a constant) to assume that the vector $(\omega, \theta_1, \dots, \theta_K)$ is jointly normal, and that there is no residual uncertainty about ω conditional on the attribute values.⁴ Because any attribute value can be replaced with its negative, assuming $\alpha_i \geq 0$ is without loss. For ease of exposition, we will further assume that each weight α_i is strictly positive.

Time is continuous, and the agent has a unit budget of attention to allocate at every instant of time. Formally, at each $t \in [0, \infty)$, the agent chooses an attention vector $\beta_1(t), \dots, \beta_K(t)$ subject to the constraints $\beta_i(t) \geq 0$ (attentions are positive) and $\sum_i \beta_i(t) \leq 1$ (allocations respect the budget constraint).

Attention choices influence the diffusion processes X_1, \dots, X_K observed by the agent in the following way:

$$dX_i^t = \beta_i(t) \cdot \theta_i \cdot dt + \sqrt{\beta_i(t)} \cdot d\mathcal{B}_i^t. \quad (1)$$

Above, each \mathcal{B}_i is an independent Brownian motion, and the term $\sqrt{\beta_i(t)}$ is a standard normalizing factor to ensure constant informativeness per unit of attention devoted to each source. Thus, devoting T units of time to observation of source i is equivalent to observation of the normal signal $\theta_i + \mathcal{N}(0, 1/T)$, or T independent observations of the standard normal signal $\theta_i + \mathcal{N}(0, 1)$. This formulation treats “attention” and “time” in the same way, in the

⁴If $\omega, \theta_1, \dots, \theta_K$ are jointly normal, then ω can be rewritten as a linear combination of the θ_i plus a residual term that is independent of each θ_i . The assumption of no residual uncertainty means that the residual term is a constant, returning Assumption 1 up to an additive constant (which can be normalized to zero in our problem).

sense that devoting $1/2$ attention to source i for a unit of time provides the same amount of information about θ_i as devoting full attention to source i for a $1/2$ unit of time. We also note that since all sources are assumed to be equally informative about their corresponding attributes, it is *with loss* to further normalize the payoff weights α_i to be equal to one another.⁵

Remark 1. As these comments suggest, there is a natural discrete-time analogue to our continuous-time model: At each period $t \in \mathbb{Z}_+$, the agent has a unit budget of precision to allocate across K normal signals. Choice of attention vector $(\pi_1(t), \dots, \pi_K(t))$ results in one observation of the normal signal $\theta_i + \mathcal{N}(0, 1/\pi_i(t))$ for each source $i = 1, \dots, K$. All of our main results have an immediate corollary in that model. See Section 6 for further discussion.

Let $(\Omega, \mathbb{P}, \{\mathcal{F}_t\}_{t \in \mathbb{R}_+})$ describe the relevant probability space, where the information \mathcal{F}_t that the agent observes up to time t is the collection of paths $\{X_i^{\leq t}\}_{i=1}^K$ (with $X_i^{\leq t}$ representing the sample path of X_i from time 0 to time t). An *information acquisition strategy* S is a map from $\{X_i^{\leq t}\}_{i,t}$ into $\Delta(\{1, \dots, K\})$, representing how the agent divides attention at each instant as a function of the observed diffusion processes.⁶ In addition to allocating his attention, the agent chooses how long to acquire information for; that is, at each instant he determines (based on the history of observations) whether to continue acquiring information, or to stop and take an action. Formally, the agent chooses a *stopping time* τ , which is a map from Ω into $[0, +\infty]$ satisfying the measurability requirement $\{\tau \leq t\} \in \mathcal{F}_t$ for all t .

At the endogenously chosen end time τ , the agent will choose from a set of actions A and receive the payoff $u(\tau, a, \omega)$, where u is a known payoff function that depends on the stopping time τ , the action taken a and the payoff-relevant state ω . This formulation allows for additively separable waiting costs, $u(\tau, a, \omega) = u_1(a, \omega) - c(\tau)$, as well as geometric discounting, $u(\tau, a, \omega) = \delta^\tau \cdot u_2(a, \omega)$. The agent's posterior belief about ω at time τ determines the action that maximizes his expected flow payoff $\mathbb{E}[u(\tau, a, \omega)]$. We will only impose the following weak assumption on the payoff function:

Assumption 2. *Given any (normal) belief about ω , $\max_a \mathbb{E}[u(\tau, a, \omega)]$ is decreasing in τ .*

⁵In fact, our subsequent results indicate that the case of equal weights is special. For example, with two sources, the conclusions of Theorem 1 always hold when $\alpha_1 = \alpha_2$, but do not hold in general.

⁶We assume that given the agent's attention strategy, the stochastic differential equations in (1) have a solution. This is true for example if each $\beta_i(t)$ is a deterministic function of t (as in the optimal strategy that we describe in Theorems 1 and 2), or if $\sqrt{\beta_i(t)}$ satisfies standard Lipschitz conditions (see Section 6.1 of Yong and Zhou (1999)).

That is, we assume that holding fixed the agent’s belief at the time of decision, an earlier decision is better. In the case of $u(\tau, a, \omega) = u_1(a, \omega) - c(\tau)$, this assumption requires the waiting cost $c(\tau)$ to be non-decreasing in τ ; in the case of $u(\tau, a, \omega) = \delta^\tau \cdot u_2(a, \omega)$, the assumption is that the optimal flow payoff $\max_a \mathbb{E}[u_2(a, \omega)]$ is non-negative (which is satisfied, for example, if there is a default action that always yields zero payoff).

To summarize, the agent chooses his information acquisition strategy and stopping time (S, τ) to maximize $\mathbb{E}[\max_a \mathbb{E}[u(\tau, a, \omega)|\mathcal{F}_\tau]]$. In this paper we primarily focus on characterizing the optimal information acquisition strategy S . In general the strategies S and τ should be determined jointly, but our results will show that in many cases the optimal S can be characterized independently from the choice of τ .

3 Main Results

In Section 3.1 we consider the case of two attributes, where we are able to derive a slightly stronger result. In Section 3.2 we characterize the optimal attention allocation strategy for any finite number of attributes. All proofs appear in the appendix, and we provide an extended explanation of these results in Section 4.

3.1 Two Attributes

Suppose there are two attributes θ_1 and θ_2 , and the payoff-relevant state is $\omega = \alpha_1\theta_1 + \alpha_2\theta_2$, with each $\alpha_i > 0$. The agent’s prior over the unknown attributes is $(\theta_1, \theta_2) \sim \mathcal{N}(\mu, \Sigma)$. Then the prior covariance between each attribute i and the payoff-relevant state ω is $cov_i := \text{Cov}(\omega, \theta_i) = \alpha_i\Sigma_{ii} + \alpha_j\Sigma_{ji}$, and we assume that these covariances satisfy the following relationship:

Assumption 3. $cov_1 + cov_2 = \alpha_1(\Sigma_{11} + \Sigma_{12}) + \alpha_2(\Sigma_{21} + \Sigma_{22}) \geq 0$.

Since both variances Σ_{11}, Σ_{22} are positive, this property holds if the covariance Σ_{12} is not too negative relative to the size of either variance. If the weights on the two attributes are equal (i.e., $\alpha_1 = \alpha_2$), this property holds for all prior beliefs over the attributes (since $2 \cdot |\Sigma_{12}| \leq 2 \cdot \sqrt{\Sigma_{11} \cdot \Sigma_{22}} \leq \Sigma_{11} + \Sigma_{22}$). If the attributes are positively correlated ($\Sigma_{12} = \Sigma_{21} \geq 0$), then this property holds for all payoff weights α_1 and α_2 .

Theorem 1. *Suppose $K = 2$ and Assumption 3 is satisfied. W.l.o.g. let $cov_i \geq cov_j$. Define*

$$t_i^* := \frac{cov_i - cov_j}{\alpha_j \det(\Sigma)}.$$

Then there exists an optimal information acquisition strategy which is history-independent and hence can be expressed as a deterministic path of attention allocations $(\beta_1(t), \beta_2(t))_{t \geq 0}$.

This path consists of two stages:

- **Stage 1:** *At all times $t < t_i^*$, the agent optimally allocates all attention to attribute i (that is, $\beta_i(t) = 1$ and $\beta_j(t) = 0$).*
- **Stage 2:** *At all times $t \geq t_i^*$, the agent optimally allocates attention in the constant proportion $(\beta_1(t), \beta_2(t)) = \left(\frac{\alpha_1}{\alpha_1 + \alpha_2}, \frac{\alpha_2}{\alpha_1 + \alpha_2}\right)$.*

There are two stages of information acquisition. In the first stage, which ends at some time t_i^* , the agent allocates all of his attention to the attribute i with higher prior covariance with the payoff-relevant state. After time t_i^* , he divides his attention across the attributes in a constant ratio, where the long-run instantaneous attention allocation is proportional to the weights α . Note that depending on the agent's stopping rule, Stage 2 of information acquisition may never be reached along some histories of the diffusion processes. But as long as the agent continues acquiring information, his attention allocations are as given above. In Online Appendix O.1, we show that under mild technical assumptions, the optimal attention strategy is in fact unique up to the stopping time τ (after which attention allocations obviously do not matter).

The characterization reveals that the optimal information acquisition strategy is completely determined from the prior covariance matrix Σ and the payoff weight vector α (the prior mean vector μ does not play a role). In particular, the strategy does not depend on details of the agent's payoff function $u(\tau, a, \omega)$, including his time preferences. When the prior belief satisfies Assumption 3, the optimal information acquisition strategy is thus constant across different objectives and also across different stopping rules. Relatedly, as we demonstrate in Section 5.1, we can solve for the optimal stopping rule in this setting as if information acquisition were *exogenously* given by Theorem 1. In Online Appendix O.2.1, we provide an example to illustrate that these properties can fail when Assumption 3 is violated. Online Appendix O.2.2 further shows that for the case of two attributes, Assumption 3 is not only sufficient but also necessary for our characterization to hold independently of the agent's payoff criterion.

Below we illustrate this optimal strategy using a few simple examples.

Example 1 (Independent Attributes). Suppose $(\theta_1, \theta_2) \sim \mathcal{N}(\mu, \begin{pmatrix} 6 & 0 \\ 0 & 1 \end{pmatrix})$ and the payoff-relevant state is $\omega = \theta_1 + \theta_2$. Then, applying Theorem 1, the agent initially puts all attention towards learning θ_1 . At time $t_1^* = \frac{5}{6}$, his posterior covariance matrix is the identity matrix. After this time he optimally splits attention equally between the two attributes, which are now symmetrically distributed.

Example 2 (Correlated Attributes). Suppose $(\theta_1, \theta_2) \sim \mathcal{N}(\mu, \begin{pmatrix} 6 & 2 \\ 2 & 1 \end{pmatrix})$, and the payoff-relevant state is still $\omega = \theta_1 + \theta_2$. Applying Theorem 1, the agent initially puts all attention towards learning θ_1 . At time $t_1^* = \frac{5}{2}$, his posterior covariance matrix becomes $\begin{pmatrix} 3/8 & 1/8 \\ 1/8 & 3/8 \end{pmatrix}$. Compared to the previous example, it takes longer for the agent's uncertainty about the two attributes to equalize, since information about θ_1 also reduces the agent's uncertainty about θ_2 . After $t_1^* = \frac{5}{2}$ he optimally splits attention equally between the two attributes.

Example 3 (Unequal Payoff Weights). Consider the prior belief given in the previous example, but suppose now that the payoff-relevant state is $\omega = \theta_1 + 2\theta_2$. As before, the agent initially puts all attention towards learning θ_1 . Stage 1 ends at time $t_1^* = \frac{3}{2}$, when the posterior covariance matrix is $\begin{pmatrix} 3/5 & 1/5 \\ 1/5 & 2/5 \end{pmatrix}$. Note that because of the asymmetry in the payoff weights, the agent's posterior uncertainty about the two attributes is not the same at this switch point. As we will discuss in Section 4, however, the *marginal values* of learning about the two attributes are equal to one another at time t_1^* . After this time, the agent optimally acquires information in the mixture $(1/3, 2/3)$.

3.2 K Attributes

We now consider the case of multiple attributes. We provide three different sufficient conditions under which the optimal information acquisition strategy can be exactly characterized.

Assumption 4 (Perpetual Substitutes). Σ^{-1} has non-positive off-diagonal entries.

Assumption 5 (Perpetual Complements). Σ has non-positive off-diagonal entries and $\Sigma \cdot \alpha$ has non-negative coordinates.

Assumption 6 (Diagonal Dominance). Σ^{-1} is diagonally-dominant. That is, $[\Sigma^{-1}]_{ii} \geq \sum_{j \neq i} |[\Sigma^{-1}]_{ij}|$ for all $1 \leq i \leq K$.

Assumption 4 generalizes a previous sufficient condition $\Sigma_{12} \geq 0$ to more than two attributes. It requires that the *partial correlation* between any two attributes—controlling for all other attributes—is positive.^{7,8} Proposition 7 in Online Appendix O.3 shows that Assumption 4 is an if and only if condition for information from any pair of sources to be “perpetually substitutable.” By this we mean that the value of acquiring information from one source is decreasing in the amount of information from the other source, and that this property holds not only at time $t = 0$ but also along any path of information acquisitions.

Assumption 5 imposes that all attributes are negatively correlated with one another, but that each attribute is initially positively correlated with the payoff-relevant state. Under this assumption, prior covariances are “mildly” negative. Proposition 8 in Online Appendix O.3 shows that Assumption 5 is an if and only if condition for information from any pair of sources to be perpetually complementary, in the sense that the value of acquiring information from one source is increasing in the amount of information from the other source, again along the entire path of information acquisitions.

The last Assumption 6 requires the inverse of the prior covariance matrix to be diagonally-dominant. Roughly speaking, this assumption allows for some pairs of attributes to be complements and other pairs to be substitutes, but puts an upper bound on the magnitude of any complementarity or substitution effects. When $K = 2$, Assumption 6 reduces to the simple condition $|\Sigma_{12}| \leq \min\{\Sigma_{11}, \Sigma_{22}\}$, which is sufficient for our previous Assumption 3. For general K , Assumption 6 is implied by a similar condition $|\Sigma_{ij}| \leq \frac{1}{2K-3}\Sigma_{ii}$ for all $i \neq j$ (see Appendix A.3) We explain the role of these assumptions in Section 4.

Theorem 2. *Suppose any of Assumption 4, 5, or 6 is satisfied.⁹ Then, there exist times*

$$0 = t_0 < t_1 < \dots < t_m = +\infty$$

⁷The partial correlation between attributes θ_i and θ_j under covariance matrix Σ is equal to $-\frac{[\Sigma^{-1}]_{ij}}{\sqrt{[\Sigma^{-1}]_{ii}[\Sigma^{-1}]_{jj}}}$.

⁸Positive-definite matrices with non-positive off-diagonal entries are known as *M*-matrices (Plemmons, 1977). It is well-known that the inverse of such matrices has positive entries everywhere. Thus Assumption 4 implies positive correlation $\Sigma_{ij} \geq 0$, but is strictly stronger when $K > 2$. We mention that assuming $\Sigma_{ij} \geq 0$ for all $i \neq j$ does not guarantee Theorem 2 to hold. A counterexample with three sources is presented in Online Appendix O.2.3.

⁹We point out that while Assumption 4, 5, and 6 are each sufficient for the theorem to hold, they are not in general necessary (unlike Assumption 3 in the $K = 2$ case).

and nested sets

$$\emptyset \subsetneq B_1 \subsetneq \cdots \subsetneq B_m = \{1, \dots, K\},$$

such that an optimal information acquisition strategy is described by a deterministic path of attention allocations $(\beta(t))_{t \geq 0}$. This attention path consists of $m \leq K$ stages: For each $1 \leq k \leq m$, $\beta(t)$ is constant at all times $t \in [t_{k-1}, t_k)$ and supported on the sources in B_k . In particular, the optimal attention allocation at any time $t \geq t_{m-1}$ is proportional to α .

The full path of attention allocations $(\beta_1(t), \dots, \beta_K(t))$ (including the times t_k , the nested sets B_k , and the constant attention allocation at each stage k) can again be determined directly from the primitives Σ and α . In Appendix D, we provide an algorithm for computing this path. Theorem 2 thus tells us that the agent can reduce the dynamic information acquisition problem to a sequence of $m \leq K$ static problems, each of which involves finding the optimal constant division of attention for a fixed period of time (from t_{k-1} to t_k). Moreover, as in the $K = 2$ case, the optimal information acquisition strategy does not depend on the agent’s payoff function, and is history-independent.

We note that each of Assumption 4, 5, or 6 is “absorbing” in the following sense: If the prior covariance matrix satisfies any of these conditions, then so does any posterior covariance matrix. Propositions 7 and 8 respectively show that this is true for Assumptions 4 and 5, whereas diagonal dominance is absorbing because any information acquisition strategy only increases the diagonal entries of the precision matrix Σ^{-1} .¹⁰ The absorbing property implies that our characterization not only applies to the prior belief, but also to *any* posterior belief even if the history involves sub-optimal attention allocations. This feature enables us to study the effect of attention manipulation, an application that we pursue in Section 5.2.

Finally, we mention that starting from any prior belief (including those that fail all three of the sufficient conditions we have provided), so long as the agent does not stop learning about any attribute, his posterior belief must eventually satisfy Assumption 6.¹¹ Theorem 2 applies at these posterior beliefs, implying in particular that the agent’s optimal attention allocation is eventually constant and proportional to the weight vector α .

¹⁰By directly computing cov_1 and cov_2 at posterior beliefs, it can be shown that our previous Assumption 3 for the case of two attributes is also absorbing.

¹¹This is because any attention devoted to attribute i simply increases the i -th diagonal entry of the precision matrix Σ^{-1} . With sufficient attention devoted to each attribute, the posterior precision matrix must be diagonally-dominant.

4 Explanation of Results

4.1 Fixed Stopping Time t

Consider the simpler problem in which the agent makes a decision at an exogenously fixed and known time t . Because normal signals can be completely Blackwell-ordered based on their precisions (Hansen and Torgersen, 1974), different mixtures over the sources can be compared based on how much they reduce the variance of the payoff-relevant state.

Formally, the agent's past attention allocations integrate to a *cumulative attention vector* $q(t) = (q_1(t), \dots, q_K(t))' \in \mathbb{R}_+^K$ at time t , describing how much attention has been paid to each attribute thus far. The agent's posterior variance of ω is

$$V(q) = \alpha'(\Sigma^{-1} + \text{diag}(q(t)))^{-1}\alpha \quad (2)$$

where Σ is the prior covariance matrix over the attribute values, and $\text{diag}(q(t))$ is the diagonal matrix with entries $q_1(t), \dots, q_K(t)$. This posterior variance depends solely on the payoff weights α , the prior covariance matrix Σ , and the cumulative acquisitions $q(t)$. It does not depend on the realizations of the diffusion processes or the order of information acquisitions. So the problem of optimizing for a fixed end date t reduces to a static problem of optimally allocating t units of attention.

Define the *t-optimal* attention vector

$$n(t) = \underset{q_1, \dots, q_K \geq 0, \sum_i q_i = t}{\text{argmin}} V(q_1, \dots, q_K).$$

to be the allocation of t units of attention that minimizes posterior variance of ω , which is unique by Lemma 4 in the appendix. Every information acquisition strategy that cumulates to the attention vector $n(t)$ at time t is optimal for any decision problem at that time.

4.2 Uniformly Optimal Strategy: Definition

In general, the family of solutions $\{n(t)\}_{t \geq 0}$ corresponding to optimal allocation of t units of attention does not determine the solution to the dynamic problem. To see this, suppose $n(1) = (0, 1)$, so that the optimal way to allocate attention for a decision at time $t = 1$ is to allocate it all to the second attribute, but $n(2) = (2, 0)$, so that the optimal way to allocate attention for a decision at time $t = 2$ is to allocate it all to the first attribute. Clearly

these t -optimal attention vectors are incompatible under a single dynamic strategy: Optimal attention allocation for a decision at time $t = 1$ precludes achieving the optimal attention allocation for time $t = 2$.¹² So when the agent faces the possibility of stopping at either time $t = 1$ or $t = 2$ (depending on the realizations of the diffusion processes), he must trade off between achieving more precise information about the payoff-relevant state at either time. The optimal resolution of this trade-off depends on the specific decision problem.

If, however, it were possible to continuously string together the t -optimal attention vectors $n(t)$ along the path of one information acquisition strategy, then such intertemporal trade-offs would not arise, and we might further conjecture such a strategy to be optimal for all stopping problems. This turns out to be true.

Definition 1. *Say that an information acquisition strategy S is uniformly optimal if it is deterministic (i.e., history-independent) and its induced cumulative attention vector at each time t is the t -optimal vector $n(t)$.*

A uniformly optimal strategy, by definition, minimizes posterior variance at every instant t . A result of [Greenshtein \(1996\)](#) implies that such a strategy is best among all history-independent information acquisition strategies. In [Lemma 5](#), we extend [Greenshtein \(1996\)](#)'s result to show that a uniformly optimal strategy is best among *all* strategies, including those that condition future attention allocations on past signal realizations. In brief, we first observe that compared to any alternative information acquisition strategy, the uniformly optimal strategy achieves the same precision of beliefs about the state at earlier times. We then use this observation to show that any decision rule (i.e., any stopping time and final action) achievable under the alternative strategy can be replicated under the uniformly optimal strategy in a way that makes the agent stop earlier, but maintains his belief at the time of stopping. This “replicating” decision rule, together with the uniformly optimal attention strategy, thus yields a higher expected payoff. We note that this proof relies crucially on the normal environment, which allows us to capture the agent’s uncertainty through the single statistic of posterior variance.

Thus, whenever a uniformly optimal strategy exists, it must be an optimal strategy in

¹²That a decision-maker cannot choose to forget or undo past information acquisitions is a standard property satisfied by any learning or information acquisition strategy. In [Nieuwerburgh and Veldkamp \(2010\)](#), this is referred to as the *no-forgetting constraint*.

our problem.¹³ It remains to show that under the assumptions we provided previously, a uniformly optimal strategy does exist, and has the structure described in Theorems 1 and 2.

4.3 Uniformly Optimal Strategy: Existence

Lemma 1 (Monotonicity). *A uniformly optimal strategy exists if and only if the t -optimal attention vector $n(t)$ weakly increases (in each coordinate) over time.*

In words, a uniformly optimal strategy exists if and only if for every $t' > t$, the optimal allocation of t' units of attention devotes a (weakly) higher amount of attention to *each* source compared to the optimal allocation of t units. Thus, monotonicity of $n(t)$ and existence of a uniformly optimal strategy are equivalent.

Whether $n(t)$ is monotone turns out to be closely related to a classic problem in consumer theory: Suppose a consumer has a utility function $U(q_1, \dots, q_K)$ over consumption of q_k units of each of K goods, and let $D(\mathbf{p}, w)$ denote his Marshallian demand subject to the budget constraint $\mathbf{p} \cdot \mathbf{q} \leq w$. Then, the consumer's demand is *normal* if each coordinate of $D(\mathbf{p}, w)$ increases with income w . In our setting, we can set $U = -V(q_1, \dots, q_K)$ to be the negative of the posterior variance, so that minimizing V is the same as maximizing U . Our t -optimal attention vector $n(t)$ is then precisely the Marshallian demand, when prices are identically equal to 1 and income is equal to t . Thus, normality of the consumer's demand under utility function $U = -V$ is equivalent to existence of a uniformly optimal strategy.

The literature on consumer demand provides conditions on U that imply normality of demand. When there are just two goods, [Alarie and Bronsard \(1990\)](#) and [Bilancini and Boncinelli \(2010\)](#) showed that demand is normal if and only if $\partial_{ij}U \cdot \partial_jU \geq \partial_iU \cdot \partial_{jj}U$ for $i, j \in \{1, 2\}$. For more than two goods, a sufficient condition for normality of demand given by [Chipman \(1977\)](#), and more recently generalized by [Quah \(2007\)](#), is concavity and super-modularity of U .¹⁴

The key difference between these conditions and ours is that the known conditions apply to the utility function U . In our environment, $U = -V = -\alpha'(\Sigma^{-1} + \text{diag}(q))^{-1}\alpha$ is a composite function of the primitive objects Σ (prior belief) and α (payoff weights). Our

¹³Our argument based on Blackwell comparisons gets to the optimal policy (i.e., attention allocation) without going through the HJB equation and value function, which may be difficult to solve for explicitly.

¹⁴The super-modularity property is also called ‘‘ALEP-complementarity’’ to distinguish from Marshallian and Hicksian complementarity.

Assumptions 3-6 apply to these primitives, which have no natural analogue in the consumer demand problem. Moreover, the conditions in the literature are stated for the utility function U evaluated at *all* demand levels q . In contrast, we show it is possible to reduce such conditions to conditions on Σ and α only, which are much easier to verify.

Specifically, our Assumption 3 for the $K = 2$ case is necessary and sufficient for the resulting function $-V$ to satisfy the condition given in Alarie and Bronsard (1990) and Bilancini and Boncinelli (2010), and our Assumption 5 (“Perpetual Complementarity”) can be shown to imply that $U = -V$ is not only concave but also super-modular, thus coinciding with the condition given by Chipman (1977) and Quah (2007). Generalizing beyond Perpetual Complementarity, we use the special form of V to develop alternative sufficient conditions. Perhaps surprisingly, we show that if the sources are substitutes at every posterior belief (i.e., $-V$ is sub-modular), then $n(t)$ is also monotone. This “Perpetual Substitution” property, too, can be stated as a simple condition on the prior covariance matrix (our Assumption 4). Finally, if correlation is not too strong—as implied by Assumption 6, which bounds the size of the covariances between the attributes relative to the size of their variances—then again we obtain monotonically increasing t -optimal attention vectors.

In our proof of Theorem 2, we show that these three different economic conditions are each sufficient to imply the following technical property: At every moment of time, those attributes that covary most strongly with the payoff-relevant state ω all have *positive* covariance with ω (Lemma 8). As demonstrated in Lemma 9, this key technical property delivers monotonicity of $n(t)$. In fact, we show in the Supplementary Materials that a weaker version of this technical property is also necessary for our characterization to hold. Further exploration of our sufficient conditions, and whether they imply results for other utility functions U besides our special case of $U = -V$, is an interesting topic for subsequent work.

4.4 Uniformly Optimal Strategy: Structure

When $n(t)$ is indeed monotone in t , the uniformly optimal strategy that achieves these vectors is simply the one that sets each attention allocation $\beta(t)$ to be the time-derivative of $n(t)$. Under this strategy, the agent divides attention at every moment across those attributes that maximize the *instantaneous* marginal reduction of posterior variance V . While the set of attributes is pinned down by first-order conditions, the precise mixture over those attributes is pinned down by second-order conditions, which ensure that this set of attributes continue

to have equal and marginal values at future instants. Specifically, for each set of attributes, there is a unique linear combination that corresponds to the “learnable component” of ω given those attributes (formally, it is a projection of ω). It turns out that at each stage, it is optimal to acquire information in a mixture proportional to the weights of this linear combination, thus producing an unbiased signal about the learnable component of ω . In the final stage, when the agent pays attention to every attribute—so that the learnable component of ω is ω itself—the optimal mixture is proportional to the payoff weights α .

As beliefs about a set of attributes become more precise, their shared marginal value decreases continuously relative to the marginal value of learning about other attributes. Eventually the marginal value of learning about some other attribute “catches up” and joins the set of maximizers. This yields the nested-set property in Theorems 1 and 2.

5 Applications

The characterizations in Theorems 1 and 2 suggest that the study of dynamic information acquisition in our setting is quite tractable. We now apply this characterization to derive new results in a diverse set of applications.

In Section 5.1, we consider optimal attention allocation for choice between two goods and generalize recent results from Fudenberg et al. (2018). In Section 5.2, we consider the dynamic implications of a one-time attention manipulation, complementing a recent exercise in Gossner et al. (2020). In Section 5.3 we develop a game between biased news sources, and characterize the degree of polarization and the quality of information in equilibrium. The applications in Sections 5.1 and 5.2 show that we can use our main results to tractably introduce correlation in settings that have been previously studied under strong assumptions of independence. Our applications in Sections 5.1 and 5.3 show how our main results can be used as an intermediate step to derive results about other economic behaviors.

5.1 Application 1: Binary Choice

Building on a large literature regarding “binary choice” problems, Fudenberg et al. (2018) (henceforth FSS) recently proposed the *uncertain drift-diffusion* model: An agent has a choice between two goods with unknown payoffs v_1 and v_2 , and can learn about those payoffs

before making a choice. The two payoffs are jointly normal and i.i.d.:

$$(v_1, v_2)' \sim \mathcal{N}\left(\mu, \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}\right) \quad (3)$$

The agent observes two Brownian processes whose drifts are the unknown payoffs. He then chooses a stopping time τ to maximize the objective

$$\mathbb{E}[\mathbb{E}[\max\{v_1, v_2\} \mid \mathcal{F}_\tau] - c\tau],$$

where $c\tau$ is a linear waiting cost.

FSS's main economic insight is that earlier decision times are associated with more accurate decisions. Formally, let $p(t)$ be the probability of choosing the higher-value good conditional on stopping at time t . FSS's Proposition 3 shows that $p(t)$ is monotonically (weakly) decreasing over time. This comparative static is not obvious because two forces push in opposite directions: On the one hand, the agent has more information at later times, suggesting that later decisions may be more accurate. On the other hand, because the stopping time is endogenously chosen, the agent is more likely to stop earlier when the decision is easy (i.e., when one good's value is much higher than the other). FSS's result implies that this second force dominates.

FSS show moreover that this result is robust to endogenous attention allocation under a budget constraint. Specifically, suppose that at each moment of time, the agent has one unit of attention to allocate across learning either v_1 or v_2 . Then, FSS's Theorem 5 shows that the agent optimally divides attention equally between learning about the two payoffs at every moment of time, similar to the exogenous process specified in their main model.

FSS's model of endogenous attention and binary choice is nested within our framework. To map this setting back into our main model, define $\theta_1 = v_1$ and $\theta_2 = -v_2$. Then, since the payoff difference $v_1 - v_2$ is a sufficient statistic for the agent's decision, this problem corresponds to payoff-relevant state $\omega = v_1 - v_2 = \theta_1 + \theta_2$ in our framework.

We now show that we can use our results to go beyond the case of independent and identically distributed payoffs (as imposed in the prior in (3)). Different from FSS, suppose that the agent's prior over (θ_1, θ_2) is normal with an arbitrary covariance matrix $\Sigma := \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$. Since the payoff weights are $\alpha_1 = \alpha_2 = 1$, our Theorem 1 applies and characterizes the agent's optimal attention allocations over time given any Σ . The following corollary is an immediate generalization of FSS's Theorem 5:

Corollary 1. *Suppose $\Sigma_{11} \geq \Sigma_{22}$. The agent's optimal information acquisition strategy $(\beta_1(t), \beta_2(t))_{t \geq 0}$ in this binary choice problem consists of two stages:*

- **Stage 1:** *At all times*

$$t < t_1^* = \frac{\Sigma_{11} - \Sigma_{22}}{\det(\Sigma)},$$

the agent optimally allocates all attention to θ_1 .

- **Stage 2:** *At times $t \geq t_1^*$, the agent optimally allocates equal attention to θ_1 and θ_2 .*

Thus, when the agent is initially more uncertain about one of the two payoffs, he spends a period of time exclusively learning about that payoff. Starting at time t_1^* , the agent divides attention equally across learning about the two goods, as in FSS's i.i.d. case. From the closed-form expression for t_1^* , it is straightforward to show that the length of the first stage is increasing in the asymmetry of initial uncertainty, and also increasing in the degree of correlation between the two payoffs.

We now use this characterization of optimal attention allocation to further generalize FSS's main economic insight regarding the relationship between choice speed and accuracy.

Proposition 1. *For any Σ , $p(t)$ is weakly decreasing in t . Thus, choice accuracy is weakly higher at earlier stopping times.*

The logic of this result is as follows (see Online Appendix O.4 for detailed analysis). First, Corollary 1 implies that we can separate the problem of optimal stopping from the problem of optimal information acquisition. That is, we can take information as exogenously given by the process described in Corollary 1, and characterize properties of optimal stopping within this problem. In particular, Corollary 1 pins down the evolution of the agent's posterior covariance matrix Σ_t , which will be important for the subsequent arguments.

While Σ_t is a deterministic process, the agent's posterior expectation for the payoff difference $\theta_1 + \theta_2$ evolves according to a random process Y_t . As in FSS, the symmetric stopping boundary at time t is given by a function $k^*(\Sigma_t)$ of the agent's posterior covariance matrix Σ_t ; that is, the agent optimally stops at time t if and only if $|Y_t| \geq k^*(\Sigma_t)$. Given these stopping boundaries, the choice accuracy $p(t)$ conditional on stopping at time t has the following simple form:

$$p(t) = \Phi\left(\frac{k^*(\Sigma_t)}{\sigma_t}\right), \quad (4)$$

where σ_t^2 is the agent’s posterior variance of $\theta_1 + \theta_2$ at time t , and Φ is the normal c.d.f. function.

So it remains to understand how (4) evolves. There are two forces, which turn out to go in the same direction. First, uncertainty about the payoff difference $\theta_1 + \theta_2$ decreases over time. As FSS already showed in the i.i.d. case, this effect weakly decreases the ratio $k^*(\Sigma_t)/\sigma_t$. Roughly speaking, stopping at an earlier time when there is more residual uncertainty requires the agent to have received disproportionately stronger signals to forgo the option value. In our Lemma 16, we generalize this insight to arbitrary prior beliefs.

Second, our characterization in Corollary 1 reveals that the optimal attention strategy continuously reduces the asymmetry in uncertainty between the two attributes. We show in Lemma 15 that holding fixed uncertainty about the sum $\theta_1 + \theta_2$, asymmetry in uncertainty about the two attributes θ_1 and θ_2 allows the agent to learn faster. That is, comparing an agent with an asymmetric prior to another agent with a symmetric prior, if the two agents have the same prior variance of $\theta_1 + \theta_2$, then the asymmetric prior leads to lower posterior variance of $\theta_1 + \theta_2$ at every future time (when both agents adopt the optimal attention strategy).¹⁵ So asymmetric uncertainty increases the option value to waiting, and thus also the stopping boundary relative to the symmetric baseline. This effect, too, causes the ratio $k^*(\Sigma_t)/\sigma_t$ to decrease over time. Combining both effects yields the result that $p(t)$ decreases over time.

In fact, using similar arguments, we can further generalize Proposition 1 to asymmetric learning speeds about the two unknown payoffs. See Online Appendix O.4.5.

¹⁵A simple informal argument is as follows. Given any prior, we can upper-bound the posterior variance under the optimal strategy by the posterior variance under a strategy that devotes equal attention to θ_1 and θ_2 at all times, which is equivalent to receiving the pair of signals $\theta_1 + \mathcal{N}(0, 2)$ and $\theta_2 + \mathcal{N}(0, 2)$ over every unit of time. This pair must be weakly more informative than any summary statistic of the two signals—for example, their sum. Thus the agent learns about $\theta_1 + \theta_2$ at least as fast as if he received the signal $\theta_1 + \theta_2 + \mathcal{N}(0, 4)$ over every unit of time. When the prior is symmetric, then the equal-attention strategy considered above is optimal, and this lower bound on the speed of learning is tight. But when the prior is asymmetric, then the agent can improve upon this bound. This suggests that the agent can learn more quickly under an asymmetric prior than a symmetric one, holding fixed his prior uncertainty about $\theta_1 + \theta_2$.

5.2 Application 2: Attention Manipulation

So far, we have assumed that the agent has complete control over how to allocate his attention. In practice, businesses expend substantial effort to divert attention towards their products. Such “attention grabbing” often takes the form of a one-time intervention (e.g., an ad) rather than a continual shift in exposure, so the value of the attention diversion depends on how it shapes subsequent allocation of attention. Two questions thus naturally arise: 1) Does a one-time manipulation of attention towards a given source lead to a persistently higher amount of attention devoted to that source, or will the decision-maker quickly “compensate” for the manipulation? 2) What are the externalities on other sources—in particular, is it the case that manipulating attention towards one source decreases the amount of attention devoted to others?

Gossner et al. (2020) (henceforth GSS) recently studied this question in a model in which an agent sequentially learns about the quality of a number of goods by allocating attention to one good each period. One of their main results (Theorem 1) resolves the two questions posed above in the following way:

- (1) the cumulative amount of attention paid to that good remains persistently higher following the attention manipulation, and
- (2) the cumulative amount of attention paid to any other good remains persistently lower following the attention manipulation.

A key assumption in GSS is that the attention strategy used by the agent satisfies a version of Independence of Irrelevant Alternatives (IIA): Conditional on *not* focusing on the good to which attention is diverted, the agent’s belief about that good does not affect the relative probabilities of focusing on the remaining goods. Proposition 5 in Gossner et al. (2020) shows that when the agent adopts a class of “satisficing” stopping rules, the optimal attention strategy satisfies IIA for independent goods.¹⁶

We can use our framework and main characterization to study a related but different problem, where the agent learns about multiple attributes of an unknown (one-dimensional) payoff-relevant state, and—importantly—those attribute values can be correlated. Additionally, we differ from GSS by focusing on the optimal attention allocation strategy and how it

¹⁶This additional assumption on the stopping rule is not required in our setting, since we focus on the uniformly optimal attention strategy which is independent of stopping behavior.

is affected by attention manipulation. Outside of the special case of independent attributes, the optimal strategy in our setting fails IIA when there are more than two attributes. Nevertheless, we show GSS’s finding in (1) holds for flexible patterns of correlation, and the finding in (2) holds under an additional condition, which we make precise.

Formally, suppose attention is manipulated such that the agent only attends to source 1 from time zero to time T , where $T > 0$ is fixed in this section. After time T , the agent adopts the optimal attention strategy given his posterior belief at T . The dynamic effect of the one-time attention manipulation is then understood by comparing the cumulative attention vectors under the optimal strategy and under the manipulated strategy. We assume throughout that our previous conditions on the prior covariance matrix apply (i.e., Assumption 3 if $K = 2$ and Assumption 4, 5 or 6 if $K > 2$).

Proposition 2. *Let $T^* \geq T$ be the earliest time at which cumulative attention towards source 1 exceeds T under the baseline (unmanipulated optimal) strategy. Then, cumulative attention towards source 1 is strictly larger under the manipulated attention strategy at every moment of time $t \in (T, T^*)$, and equal to the baseline at all later times $t \geq T^*$.*

Thus, the finding in (1) holds under arbitrary correlation (so long as our characterizations apply): Attention manipulation towards source 1 has a persistent positive effect on the total amount of attention source 1 receives up to every future time. On the other hand, we show that this increase in cumulative attention vanishes in the long run, with the cumulative attention paid to source 1 under the baseline strategy “catching up” to the manipulated strategy by time T^* .

The proof of this proposition is simple given our previous analysis. The cumulative attention vector under the optimal strategy is the t -optimal vector defined as

$$n(t) = (n_1(t), \dots, n_K(t)) = \underset{q_1, \dots, q_K \geq 0: \sum_i q_i = t}{\operatorname{argmin}} V(q), \quad (5)$$

where V is the posterior variance function given by (2). On the other hand, the manipulated strategy induces the following *constrained* t -optimal vector $\hat{n}(t)$ at any time $t > T$:¹⁷

$$\hat{n}(t) = (\hat{n}_1(t), \dots, \hat{n}_K(t)) = \underset{q_1, \dots, q_K \geq 0: \sum_i q_i = t \text{ and } q_1 \geq T}{\operatorname{argmin}} V(q). \quad (6)$$

¹⁷As discussed, the “absorbing property” of our sufficient conditions implies that our characterizations apply to the posterior belief at time T after the agent has paid T units of attention to source 1. Thus, the cumulative attention vector at time $t \geq T$ must minimize the posterior variance V among feasible attention vectors following the manipulation. This leads to the constrained t -optimal vector $\hat{n}(t)$.

If $n_1(t) \geq T$, then the unconstrained t -optimal vector $n(t)$ satisfies the constraint in (6), so it coincides with the constrained t -optimal vector $\hat{n}(t)$. Moreover, $n_1(t) \rightarrow \infty$ as $t \rightarrow \infty$ because our characterization says that source 1 receives positive and constant attention at every instant in the final stage. Thus while initially source 1 must receive higher cumulative attention under the manipulated strategy, eventually the cumulative attention devoted to source 1 must be the same under the baseline and manipulated strategies. To show that T^* is this switch point, note that $t \geq T^*$ implies $n_1(t) \geq T$ (by definition of T^* and monotonicity of $n_1(t)$), in which case $n_1(t) = \hat{n}_1(t)$. And if $T < t < T^*$, we have $n_1(t) < T \leq \hat{n}_1(t)$, so the manipulated amount of attention devoted to source 1 strictly exceeds that of the baseline strategy. This yields the result.

When there are only two attributes, Proposition 2 also delivers GSS's second finding, namely that diversion of attention towards learning about one attribute weakly reduces cumulative attention towards learning about the remaining attribute at every moment of time. With more than two attributes, however, correlation between the attributes can overturn this result.

Example 4. Suppose there are 3 attributes, the payoff-relevant state is $\omega = \theta_1 + \theta_2 + \theta_3$, and the prior covariance matrix is

$$\Sigma = \begin{pmatrix} 3 & -2 & 0 \\ -2 & 3 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

Since Σ^{-1} is diagonally-dominant, Theorem 2 applies.

Without attention manipulation, the optimal strategy devotes the first 0.5 units of attention towards θ_3 . At $t = 0.5$, the three sources have exactly equal marginal values (and equal payoff weights), so equal attention is optimal afterwards. Thus $n(t) = (0, 0, t)$ for $t < 0.5$ and $n(t) = (\frac{t-0.5}{3}, \frac{t-0.5}{3}, \frac{t+1}{3})$ for $t \geq 0.5$.

Now suppose instead that the agent is forced to attend to source 1 for 0.1 unit of time. Then after the first 0.1 units of attention devoted towards θ_1 , the agent optimally still begins by learning about θ_3 . This lasts until $t^* = \frac{7}{15} < 0.5$, at which time source 2 has the same marginal value as source 3. The second stage then involves learning about θ_2 and θ_3 using the constant attention ratio 3 : 7. The third stage begins at $t^{**} = 0.8$, after which equal attention across the three sources is optimal.

It can be checked that the manipulation of attention towards source 1 weakly *increases* the cumulative attention towards source 2 at all times, and strictly so during the period $t \in (\frac{7}{15}, 0.8)$.

Intuitively, in this example sources 1 and 2 provide complementary information (since $\Sigma_{12} < 0$). Manipulating attention towards source 1 thus increases the marginal value of source 2, and the agent begins observing source 2 earlier than he would have otherwise. In contrast, we might expect that when all sources are substitutes with one another, attention manipulation to source 1 must decrease the amount of attention devoted to every other source. The challenge is understanding what the appropriate notion of “substitutes” is. This turns out to be the property given earlier in Assumption 4, which guarantees that each pair of attributes has a positive partial correlation coefficient.

Proposition 3. *Suppose all pairs of sources are substitutes (i.e., Assumption 4 is satisfied). Then cumulative attention towards every source $i > 1$ is weakly smaller under the manipulated strategy than under the baseline strategy, at every moment of time.*

This result complements GSS by demonstrating a class of correlated attributes for which manipulation of attention towards one reduces cumulative attention towards all others. Together with our previous Proposition 2, it shows that GSS’s Attention Theorem extends beyond their IIA assumption. Since our environment and GSS’s are non-nested, these results collectively point to the possibility of a more general set of sufficient conditions, which we leave to future work.

5.3 Application 3: Biased News Sources

Our previous two applications demonstrate that we can use our results to build on existing results from the literature. We now propose a new model—specifically, a game between strategic information providers—and show that our characterization of optimal information acquisition can be used to derive qualitative insights in this new setting.

In our model, a representative news reader seeks to learn an unknown state $\omega \sim \mathcal{N}(\mu_\omega, \sigma_\omega^2)$, and can learn about ω from either of two news sources labelled $i \in \{1, 2\}$, which are associated with opposite political parties. Source 1’s political party would like for the reader to perceive the state to be $\omega + b$, while source 2’s political party would like for the reader to

perceive the state to be $\omega - b$. The bias b is unknown to the reader, and the reader's prior is that $b \sim \mathcal{N}(\mu_b, \sigma_b^2)$.

Each news source decides how much to bias their reporting in the direction favorable to their party, and how accurate to make their reporting. Formally, each source chooses a bias intensity $\phi_i > 0$ and noise parameter $\zeta_i > 0$. A unit of time spent on source 1 is informationally equivalent to a realization of the random variable

$$Z_1 \sim \mathcal{N}(\omega + \phi_1 b, \zeta_1^2),$$

while a unit of time spent on source 2 is informationally equivalent to a realization of the random variable

$$Z_2 \sim \mathcal{N}(\omega - \phi_2 b, \zeta_2^2).$$

Both choices ϕ_i and ζ_i are fixed across time and observable to the reader.¹⁸

For any given choices of precisions and bias intensities $(\phi, \zeta) = (\phi_1, \phi_2; \zeta_1, \zeta_2)$, the reader's optimal attention allocations are denoted $(\beta_1^{(\phi, \zeta)}(t), \beta_2^{(\phi, \zeta)}(t))_{t \geq 0}$ and can be derived from the characterization in Theorem 1.¹⁹ Each source i 's payoff is a combination of a preference for viewership and a preference for bias:

$$U_i(\phi, \zeta) = \int_0^\infty r e^{-rt} \beta_i^{(\phi, \zeta)}(t) dt - \lambda(1 - \phi_i)^2.$$

The first part, $\int_0^\infty r e^{-rt} \beta_i^{(\phi, \zeta)}(t) dt$, is the discounted average attention paid to source i given the common discount rate r ,²⁰ which we interpret as a reduced form for advertising revenue. The second part, $-\lambda(1 - \phi_i)^2$, penalizes the news source for distance from the ideal bias for its party. This part of the payoff function is maximized by choosing $\phi_i = 1$, in which case

¹⁸For example, a source may have a reputation for providing very biased information but having high-quality reporting.

¹⁹To apply this theorem, we can transform the current setting to our main model: Define $\theta_1 = \frac{1}{\zeta_1}(\omega + \phi_1 b)$ and $\theta_2 = \frac{1}{\zeta_2}(\omega - \phi_2 b)$, so that a unit of time spent on each source i produces an equally informative (standard normal) signal about θ_i . The payoff-relevant state can be rewritten as $\omega = \alpha_1 \theta_1 + \alpha_2 \theta_2$ with payoff weights $\alpha_1 = \zeta_1 \cdot \frac{\phi_2}{\phi_1 + \phi_2}$ and $\alpha_2 = \zeta_2 \cdot \frac{\phi_1}{\phi_1 + \phi_2}$. It can be checked that $\text{Cov}(\omega, \theta_i) = \sigma_\omega^2 / \zeta_i > 0$, so our Assumption 3 is satisfied and Theorem 1 holds.

²⁰Note that in this formulation, we implicitly assume that the reader never stops information acquisition, which simplifies our subsequent analysis. However, $\int_0^\infty r e^{-rt} \beta_i(t) dt$ can be interpreted as the limiting discounted average attention that source i receives when the reader chooses an endogenous stopping time under vanishingly small information acquisition costs. Never stopping can also be justified in an extension of our main model where the agent faces multiple decisions across time (see Section 6).

source 1's signal is centered at $\omega + b$ and source 2's signal is centered at $\omega - b$. The parameter $\lambda \in \mathbb{R}_+$ determines the strength of incentive for bias, relative to the incentive for viewership.

The following proposition reports equilibrium in this game between the two sources. For technical reasons, we require an assumption that the incentive for bias is not too weak.

Proposition 4. *Suppose $\lambda \geq 1.6$.²¹ The unique pure strategy equilibrium is $(\phi_1^*, \zeta_1^*; \phi_2^*, \zeta_2^*)$ where*

$$\phi_1^* = \phi_2^* = \frac{1}{2} \left(1 + \sqrt{1 - \frac{1}{2\lambda}} \right)$$

and

$$\zeta_1^* = \zeta_2^* = \frac{\sigma_b}{2\sqrt{r}} \cdot \left(1 + \sqrt{1 - \frac{1}{2\lambda}} \right).$$

Given these equilibrium choices, the reader optimally devotes equal attention to the two sources at every moment.

The subsequent corollary regarding the informativeness of news in equilibrium follows immediately.²²

Corollary 2 (Informativeness of News). *The equilibrium noise level, $\zeta^* = \frac{\sigma_b}{2\sqrt{r}} \left(1 + \sqrt{1 - \frac{1}{2\lambda}} \right)$,*

- (a) *is increasing in the incentive for bias, λ ;*
- (b) *is increasing in the prior uncertainty about partisan implications, σ_b ;*
- (c) *is decreasing in the discount rate, r .*

Part (a) says that incentives for greater bias not only increase polarization, which is expected, but also lead to a reduction in the quality of news. To understand this result, consider the incentives for source i 's choice of precision. Applying our characterization in Theorem 1, there are up to two stages of information acquisition: In Stage 1, if there is a strictly more informative source, then that source receives all viewership; in Stage 2, both sources receive a constant proportion of viewership. We show that, for any equal bias intensity choices $\phi_1 = \phi_2$, source i 's long-run share is $\frac{\zeta_i}{\zeta_1 + \zeta_2}$, while source i is chosen in Stage 1 if

²¹This assumption is imposed to guarantee the existence of a pure strategy equilibrium. Our analysis shows that $\lambda \geq 1.6$ is sufficient for existence, whereas a weaker condition $\lambda \geq \frac{9}{16}$ is necessary.

²²It can be computed that when the sources choose $\zeta_1 = \zeta_2 = \zeta^*$ and $\phi_1 = \phi_2$, the news reader's posterior variance of the payoff-relevant state ω at time t is $\left(\frac{1}{\sigma_\omega^2} + \frac{t}{(\zeta^*)^2} \right)^{-1}$. This confirms why ζ^* is a sufficient statistic for equilibrium informativeness.

and only if its noise term is smaller ($\zeta_i < \zeta_{3-i}$). Thus, sources face a trade-off between optimizing for greater long-run viewership—where a larger noise choice ζ_i increases the long-run share—versus competing to be chosen in the short-run—which encourages smaller ζ_i . Intuitively, more precise information improves the competitive value of a source at the beginning of time, but reduces the value of continual engagement with that source. In equilibrium, sources choose the same ζ_i , thus washing out the first stage of information acquisition.

The size of this common noise level, however, depends on the incentives for bias. When sources provide biased news, the reader must attend to both sources to learn the truth. Polarized news sources thus live in symbiosis, where the extremity of bias on one side increases the value of information from the other. In the language of our paper, two sufficiently polarized news sources on opposite sides provide complementary information (while in contrast, two unbiased sources about ω provide fully substitutable information). The strength of complementarity increases monotonically with the degree of polarization.

Since the reader has stronger preferences for mixing over the two sources when they are complements, this means that the length of Stage 1 (when it exists) is decreasing in the degree of polarization. Thus, the more polarized the news sources are, the more emphasis these sources place on the long run, which in turn leads to lower quality news provision as we have discussed. This gives the conclusion in Part (a) of Corollary 2. In addition, larger prior uncertainty about b implies higher value of de-biasing and thus also a shorter Stage 1, leading to Part (b).

Part (c) of Corollary 2 holds by very similar reasoning. Less patient news sources compete over short-run profits (i.e., being chosen in Stage 1), and thus prefer precise signals, while patient sources compete for long-run profits (i.e., long-run proportion), and thus prefer imprecise signals. So the less patient the sources are (larger r), the more precise their signals will be in equilibrium (smaller ζ^*).

We believe that some of these qualitative insights may extend to a richer model, for example if sources simultaneously report on multiple states, the reader endogenously stops acquiring news once his belief is sufficiently precise, or if the sources have direct preferences over the reader’s beliefs. We leave exploration of these interesting extensions to future work.

6 Discussion

Information acquisition is a classic problem within economics, but there are relatively few dynamic models that are simultaneously rich and tractable. In this paper we present a class of dynamic information acquisition problems whose solution can be explicitly characterized in closed-form. We show that a complete analysis is feasible if we assume: (1) Gaussian uncertainty, (2) a one-dimensional payoff-relevant state, and (3) correlation across the unknowns that satisfies certain assumptions (for example if correlation is not too strong). Given these restrictions, a great deal of generality can be accommodated in other aspects of the problem, such as the payoff function. The tractability of the solution and the flexibility of the environment open the door to interesting applications, a number of which we have illustrated here.

We conclude by briefly mentioning a few other potential extensions and variations.

Discrete Time. Although our main model is in continuous time, our results have direct analogues in a related discrete-time model. Specifically, for the model previously described in Remark 1, we have the following result: *Suppose any of Assumption 4, 5, or 6 holds. Then at each period $t \in \mathbb{Z}_+$, the optimal allocation of precision is $(\pi_1(t), \dots, \pi_K(t))$ where $\pi_i(t) = \int_t^{t+1} \beta_i(s) ds$ for each i , with $\beta_i(s)$ being the optimal attention allocation for the continuous-time model as described in Theorem 2.*²³

Exogenous Stopping. Although we have assumed that the agent endogenously chooses when to stop acquiring information, our results hold without modification if instead the end date arrives according to an arbitrary exogenous distribution. Additionally, in that alternative model, the optimal path of attention allocations is uniquely characterized by our results so long as the exogenous distribution of end date has full support.

Intertemporal Decision Problems. Our main model assumes that the agent takes only one action, which simplifies the exposition. But since our analysis based on the notion of uniform optimality is independent of details of the payoff function, it can be easily generalized

²³In a companion piece, [Liang et al. \(2019\)](#), we discretize not only time but also information acquisitions: At each period t , the agent has to choose one of K standard normal signals, without the ability to allocate fractional precisions. The necessity of integer approximation complicates the characterization of the full sequence of signal choices. In that paper we instead provide conditions under which myopic acquisition is (eventually) optimal.

to a setting where the agent takes N actions a_1, \dots, a_N at times $\tau_1 \leq \dots \leq \tau_N$. Our characterization of the optimal attention strategy extends for any (intertemporal) payoff function $u(\tau_1, \dots, \tau_N, a_1, \dots, a_N, \omega)$ that is decreasing in the decision times τ_1, \dots, τ_N .

A Preliminaries

A.1 Posterior Variance Function

Given q_i units of attention devoted to learning about each attribute i , the posterior variance of ω can be written in two ways:

Lemma 2. *It holds that*

$$V(q_1, \dots, q_K) = \alpha' [(\Sigma^{-1} + \text{diag}(q))^{-1}] \alpha = \alpha' [\Sigma - \Sigma(\Sigma + \text{diag}(1/q))^{-1}\Sigma] \alpha$$

where $\text{diag}(1/q)$ is the diagonal matrix with entries $1/q_1, \dots, 1/q_K$.

This function V extends to a rational function (quotient of polynomials) over all of \mathbb{R}^K , i.e., even if some q_i are negative.

Proof. The equality $(\Sigma^{-1} + \text{diag}(q))^{-1} = \Sigma - \Sigma(\Sigma + \text{diag}(1/q))^{-1}\Sigma$ is well-known. To see that V is a rational function, simply note that $(\Sigma^{-1} + \text{diag}(q))^{-1}$ can be written as the *adjugate matrix* of $\Sigma^{-1} + \text{diag}(q)$ divided by its determinant. Thus each entry of the posterior covariance matrix is a rational function in q . \square

Below we calculate the first and second derivatives of the posterior variance function V :

Lemma 3. *Given a cumulative attention vector $q \geq 0$, define*

$$\gamma := \gamma(q) = (\Sigma^{-1} + \text{diag}(q))^{-1}\alpha \tag{7}$$

which is a vector in \mathbb{R}^K . Then the first and second derivatives of V are given by

$$\partial_i V = -\gamma_i^2, \quad \partial_{ij} V = 2\gamma_i \gamma_j \cdot [(\Sigma^{-1} + \text{diag}(q))^{-1}]_{ij}.$$

Proof. From Lemma 2 and the formula for matrix derivatives, we have

$$\partial_i V = -\alpha'(\Sigma^{-1} + \text{diag}(q))^{-1}\Delta_{ii}(\Sigma^{-1} + \text{diag}(q))^{-1}\alpha = -[e'_i(\Sigma^{-1} + \text{diag}(q))^{-1}\alpha]^2 = -\gamma_i^2$$

where e_i is the i -th coordinate vector in \mathbb{R}^K , and $\Delta_{ii} = e_i \cdot e_i'$ is the matrix with “1” in the (i, i) -th entry and “0” elsewhere. For the second derivative, we compute that

$$\partial_{ij}V = -2\gamma_i \cdot \frac{\partial \gamma_i}{\partial q_j} = 2\gamma_i \cdot e_i' (\Sigma^{-1} + \text{diag}(q))^{-1} \Delta_{jj} (\Sigma^{-1} + \text{diag}(q))^{-1} \alpha = 2\gamma_i \cdot [(\Sigma^{-1} + \text{diag}(q))^{-1}]_{ij} \cdot \gamma_j$$

as we desire to show. The last equality follows by writing $\Delta_{jj} = e_j \cdot e_j'$, and using $e_i' (\Sigma^{-1} + \text{diag}(q))^{-1} e_j = [(\Sigma^{-1} + \text{diag}(q))^{-1}]_{ij}$ as well as $e_j' (\Sigma^{-1} + \text{diag}(q))^{-1} \alpha = e_j' \gamma = \gamma_j$. \square

Corollary 3. *V is decreasing and convex in q_1, \dots, q_K whenever $q_i \geq 0$.*

Proof. By Lemma 3, the partial derivatives of V are non-positive, so V is decreasing. Additionally, its Hessian matrix is

$$2 \text{diag}(\gamma) \cdot (\Sigma^{-1} + \text{diag}(q))^{-1} \cdot \text{diag}(\gamma),$$

which is positive semi-definite whenever $q \geq 0$. So V is convex. \square

We use these properties to show that for each t , the t -optimal vector $n(t)$ is unique:

Lemma 4. *For each $t \geq 0$, there is a unique t -optimal vector $n(t)$.*

Proof. Suppose for contradiction that two vectors (r_1, \dots, r_K) and (s_1, \dots, s_K) both minimize the posterior variance at time t . Relabeling the sources if necessary, we can assume $r_i - s_i$ is positive for $1 \leq i \leq k$, negative for $k+1 \leq i \leq l$ and zero for $l+1 \leq i \leq K$. Since $\sum_i r_i = \sum_i s_i = t$, the cutoff indices k, l satisfy $1 \leq k < l \leq K$.

For $\lambda \in [0, 1]$, consider the vector $q^\lambda = \lambda \cdot r + (1 - \lambda) \cdot s$ which lies on the line segment between r and s . Then by assumption we have $V(r) = V(s) \leq V(q^\lambda)$. Since V is convex, equality must hold. This means $V(q^\lambda)$ is a constant for $\lambda \in [0, 1]$. But $V(q^\lambda)$ is a rational function in λ , so its value remains the same constant even for $\lambda > 1$ or $\lambda < 0$. In particular, consider the limit as $\lambda \rightarrow +\infty$. Then the i -th coordinate of q^λ approaches $+\infty$ for $1 \leq i \leq k$, approaches $-\infty$ for $k+1 \leq i \leq l$ and equals r_i for $i > l$.

For each q^λ , let us also consider the vector $|q^\lambda|$ which takes the absolute value of each coordinate in q^λ . Note that as $\lambda \rightarrow +\infty$, $\text{diag}(1/|q^\lambda|)$ has the same limit as $\text{diag}(1/q^\lambda)$. Thus by the second expression for V (see Lemma 2), $\lim_{\lambda \rightarrow \infty} V(|q^\lambda|) = \lim_{\lambda \rightarrow \infty} V(q^\lambda) = V(r)$. For large λ , the first l coordinates of $|q^\lambda|$ are strictly larger than the corresponding coordinates of r , and the remaining coordinates coincide. So the fact that V is decreasing and $V(|q^\lambda|) = V(r)$ implies $\partial_i V(r) = 0$ for $1 \leq i \leq l$.

Consider the vector $\gamma = (\Sigma^{-1} + \text{diag}(r))^{-1}\alpha$. By Lemma 3, $\partial_i V(r) = -\gamma_i^2$ for $1 \leq i \leq K$. Thus $\gamma_1 = \dots = \gamma_l = 0$. Since α and thus γ is not the zero vector, there exists $j > l$ s.t. $\gamma_j \neq 0$. It follows that $\partial_1 V(r) = 0 > \partial_j V(r)$. But then the posterior variance V would be reduced if we slightly decreased the first coordinate of r (which is strictly positive since $r_1 > s_1$) and increased the j -th coordinate by the same amount. This contradicts the assumption that r is a t -optimal vector. Hence the lemma holds. \square

A.2 Optimality and Uniform Optimality

The following result ensures that a strategy that minimizes the posterior variance uniformly at all times is an optimal strategy in any decision problem.

Lemma 5. *Suppose the payoff function $u(\tau, a, \omega)$ satisfies Assumption 2, then a uniformly optimal attention strategy is dynamically optimal.*

Proof. Without loss of generality we may assume the prior mean of ω is zero; otherwise shift ω by a constant and modify the utility function accordingly. Let S^* be the uniformly optimal attention strategy, and $\{\mathcal{F}_t^*\}$ be the induced filtration. Given S^* , the optimal stopping rule τ is a solution to

$$\sup_{\tau} \mathbb{E} \left[\max_a \mathbb{E}[u(\tau, a, \omega) \mid \mathcal{F}_{\tau}^*] \right].$$

Note that the stochastic process of posterior means $M_t^* = \mathbb{E}[\omega \mid \mathcal{F}_t^*]$ is a continuous martingale adapted to the filtration $\{\mathcal{F}_t^*\}$, with $M_0^* = 0$. Moreover, since information is Gaussian, the quadratic variation $\langle M^* \rangle_t$ is simply $v_0 - v_t^*$, where v_t^* is the posterior variance of ω at time t under the strategy S^* , and v_0 is the prior variance. By definition of uniform optimality, for each t the random variable v_t^* is deterministic and moreover smallest among possible posterior variances at time t .

Thus, by the Dambis–Dubins–Schwartz Theorem (see Theorem 1.7 in Chapter V of Revuz and Yor (1999)), there exists a Brownian motion $(B_{\nu}^*)_{\nu \in [0, v_0]}$ such that

$$B_{v_0 - v_t^*}^* = \mathbb{E}[\omega \mid \mathcal{F}_t^*].$$

This allows us to change variable from the time t to the cumulative precision $v_0 - v_t^*$.

To formulate the resulting optimization problem, for each $\nu \in [0, v_0]$ we denote by $T^*(\nu)$ the time t such that $v_t^* = v_0 - \nu$; T^* is a deterministic and increasing function of ν . Then,

under the attention strategy S^* , the agent's optimal payoff can be rewritten as

$$\sup_{\tau} \mathbb{E} \left[\max_a \mathbb{E}[u(\tau, a, \omega) \mid \mathcal{F}_{\tau}^*] \right] = \sup_{\nu} \mathbb{E} \left[\max_a \mathbb{E}[u(T^*(\nu), a, \omega) \mid B_{\nu}^*] \right]. \quad (8)$$

In other words, instead of optimizing over stopping times τ adapted to $\{\mathcal{F}_t^*\}$, we can think of the agent choosing an optimal $\nu = v_0 - v_t^*$ adapted to the Brownian motion B^* .

We will show this payoff is greater than the optimal payoff under any other attention strategy S . To do this, let $\{\mathcal{F}_t\}$ be the induced filtration under S . Similar to the above, we consider the stochastic process $M_t = \mathbb{E}[\omega \mid \mathcal{F}_t]$, adapted to $\{\mathcal{F}_t\}$. Applying the Dambis-Dubins-Schwartz Theorem again, there exists a Brownian motion $(B_{\nu})_{\nu \in [0, v_0]}$ such that

$$B_{v_0 - v_t} = \mathbb{E}[\omega \mid \mathcal{F}_t].$$

Here v_t is the posterior variance under strategy S , which is in general random but always satisfies $v_t \geq v_t^*$. Note also that B may not be the same process as B^* .

Observe that for any $t \geq 0$ we have $t = T^*(v_0 - v_t^*) \geq T^*(v_0 - v_t)$. Thus the agent's payoff under strategy S is bounded above by

$$\sup_{\tau} \mathbb{E} \left[\max_a \mathbb{E}[u(\tau, a, \omega) \mid \mathcal{F}_{\tau}] \right] \leq \sup_{\tau} \mathbb{E} \left[\max_a \mathbb{E}[u(T^*(v_0 - v_{\tau}), a, \omega) \mid \mathcal{F}_{\tau}] \right],$$

where we used Assumption 2. Now we can change variable again from τ to $\nu = v_0 - v_{\tau}$, and rewrite the payoff as

$$\sup_{\nu} \mathbb{E} \left[\max_a \mathbb{E}[u(T^*(\nu), a, \omega) \mid B_{\nu}] \right] \quad (9)$$

This is the same as the RHS of (8), since B and B^* are both Brownian motions. Hence the payoff under S does not exceed the payoff under S^* , completing the proof. \square

We also have a simple converse result:

Lemma 6. *Fixing Σ and α . Suppose an information acquisition strategy is optimal for all payoff functions $u(\tau, a, \omega)$ that satisfy Assumption 2, then it is uniformly optimal.*

Proof. Take an arbitrary time t and consider the payoff function with $u(\tau, a, \omega) = -(a - \omega)^2 - c(\tau)$, where $c(\tau) = 0$ for $\tau \leq t$ and $c(\tau)$ very large for $\tau > t$. Then the agent's optimal stopping rule is to stop exactly at time t . Since his information acquisition strategy is optimal for this payoff function, the induced cumulative attention vector must achieve t -optimality. Varying t yields the result. \square

A.3 Sufficient Condition for Assumption 6

Lemma 7. *Suppose the prior covariance matrix Σ satisfies $\Sigma_{ii} \geq (2K - 3) \cdot |\Sigma_{ij}|$ for all $i \neq j$. Then its inverse matrix satisfies $[\Sigma^{-1}]_{ii} \geq (K - 1) \cdot |[\Sigma^{-1}]_{ij}|$ for all $i \neq j$, and is thus diagonally-dominant.*

Proof. By symmetry, we can focus on $i = 1$. Let $s_j = [\Sigma^{-1}]_{1j}$ for $1 \leq j \leq K$, and without loss assume s_2 has the greatest absolute value among s_2, \dots, s_K . It suffices to show

$$s_1 \geq (K - 1)|s_2|.$$

From $\Sigma^{-1} \cdot \Sigma = I$ we have $\sum_{j=1}^K [\Sigma^{-1}]_{1j} \cdot \Sigma_{j2} = 0$. Thus $\sum_{j=1}^K s_j \cdot \Sigma_{2j} = 0$ because $\Sigma_{j2} = \Sigma_{2j}$. Rearranging yields

$$|s_1 \cdot \Sigma_{21}| = |s_2 \cdot \Sigma_{22} + \sum_{j>2} s_j \cdot \Sigma_{2j}| \geq |s_2 \cdot \Sigma_{22}| - \sum_{j>2} |s_j \cdot \Sigma_{2j}| \geq |s_2 \cdot \Sigma_{22}| - \sum_{i>2} \frac{|s_2 \cdot \Sigma_{22}|}{2K - 3},$$

where the last inequality uses $|s_j| \leq |s_2|$ and $|\Sigma_{2j}| \leq \frac{1}{2K-3}|\Sigma_{22}|$ for $j > 2$. The above inequality simplifies to

$$|s_1 \cdot \Sigma_{21}| \geq \frac{K - 1}{2K - 3} \cdot |s_2 \cdot \Sigma_{22}|.$$

And since $\Sigma_{21} \leq \frac{1}{2K-3}|\Sigma_{22}|$, we conclude that $|s_1| \geq (K - 1)|s_2|$ as desired. Note that $s_1 = [\Sigma^{-1}]_{11}$ is necessarily positive, thus $s_1 \geq (K - 1)|s_2|$. \square

B Proof of Theorem 1

Define cov_1, cov_2 as in the statement of Theorem 1, and define $x_i = \alpha_i \det(\Sigma)$ to ease notation.

Given a cumulative attention vector q , let Q be a shorthand for the diagonal matrix $\text{diag}(q)$. Then by direct computation, we have

$$\begin{aligned} \gamma &:= (\Sigma^{-1} + Q)^{-1} \cdot \alpha = (\Sigma^{-1} \cdot (I + \Sigma Q))^{-1} \cdot \alpha \\ &= (I + \Sigma Q)^{-1} \cdot \Sigma \cdot \alpha = (I + \Sigma Q)^{-1} \cdot \begin{pmatrix} cov_1 \\ cov_2 \end{pmatrix} \\ &= \frac{1}{\det(I + \Sigma Q)} \begin{pmatrix} 1 + q_2 \Sigma_{22} & -q_2 \Sigma_{12} \\ -q_1 \Sigma_{21} & 1 + q_1 \Sigma_{11} \end{pmatrix} \cdot \begin{pmatrix} cov_1 \\ cov_2 \end{pmatrix} = \frac{1}{\det(I + \Sigma Q)} \begin{pmatrix} x_1 q_2 + cov_1 \\ x_2 q_1 + cov_2 \end{pmatrix}. \end{aligned}$$

By Lemma 3, this implies the marginal values of the two sources are given by:

$$\partial_1 V(q_1, q_2) = \frac{-(x_1 q_2 + cov_1)^2}{\det^2(I + \Sigma Q)}; \quad \partial_2 V(q_1, q_2) = \frac{-(x_2 q_1 + cov_2)^2}{\det^2(I + \Sigma Q)}. \quad (10)$$

Note that Assumption 3 translates into $cov_1 + cov_2 \geq 0$. Under this assumption, we will characterize the t -optimal vector $(n_1(t), n_2(t))$ and show it is increasing over time. Without loss assume $cov_1 \geq cov_2$, then cov_1 is non-negative. Let $t_1^* = \frac{cov_1 - cov_2}{x_2}$. Then when $q_1 + q_2 \leq t_1^*$ we always have

$$x_1 q_2 + cov_1 \geq cov_1 \geq x_2 q_1 + cov_2,$$

since $x_1 q_2 \geq 0$ and $x_2 q_1 \leq x_2(q_1 + q_2) \leq x_2 t_1^* = cov_1 - cov_2$. We also have

$$x_1 q_2 + cov_1 \geq -(x_2 q_1 + cov_2),$$

since $x_1 q_2, x_2 q_1 \geq 0$ and by assumption $cov_1 + cov_2 \geq 0$. Thus, (10) implies that $\partial_1 V(q_1, q_2) \leq \partial_2 V(q_1, q_2)$ at such attention vectors q . So for any budget of attention $t \leq t_1^*$, putting all attention to source 1 minimizes the posterior variance V . That is, $n(t) = (t, 0)$ for $t \leq t_1^*$.

For $t > t_1^*$, observe that (10) implies $\partial_1 V(0, t) < \partial_2 V(0, t)$ as well as $\partial_1 V(t, 0) > \partial_2 V(t, 0)$. Thus the t -optimal vector $n(t)$ is interior (i.e., $n_1(t)$ and $n_2(t)$ are both strictly positive). The first-order condition $\partial_1 V = \partial_2 V$, together with (10) and the budget constraint $n_1(t) + n_2(t) = t$, yields the solution

$$n(t) = \left(\frac{x_1 t + cov_1 - cov_2}{x_1 + x_2}, \frac{x_2 t - cov_1 + cov_2}{x_1 + x_2} \right).$$

Hence $n(t)$ is indeed increasing in t . The instantaneous attention allocations $\beta(t)$ are the time-derivatives of $n(t)$, and they are easily seen to be described by Theorem 1. In particular, the long-run attention allocation to source i is $\frac{x_i}{x_1 + x_2}$, which simplifies to $\frac{\alpha_i}{\alpha_1 + \alpha_2}$. This completes the proof.

C Proof of Theorem 2

We will first prove the result under Assumption 6. The proof is similar under the alternative Assumption 4 or 5, and is presented at the end.

Given Lemma 5, it is sufficient to show that the t -optimal vector $n(t)$ is weakly increasing in t , that its time-derivative is locally constant, and that the time-derivative has an expanding support set (as described in the theorem). The proof is divided into several sections below.

C.1 Technical Property of γ

We will use the following lemma regarding the marginal values of different sources:

Lemma 8. *Suppose Σ^{-1} is diagonally-dominant. Given an arbitrary attention vector q , define γ as in Lemma 3 and denote by B the set of indices i such that $|\gamma_i|$ is maximized. Then γ_i is the same positive number for every $i \in B$.*

Proof. We use Q to denote $\text{diag}(q)$. Since $(\Sigma^{-1} + Q)^{-1}\alpha = \gamma$, we equivalently have

$$\alpha = (\Sigma^{-1} + Q)\gamma.$$

Suppose for contradiction that $\gamma_i \leq 0$ for some $i \in B$. Using the above vector equality for the i -th coordinate, we have

$$0 < \alpha_i = \sum_{j=1}^K [\Sigma^{-1} + Q]_{ij} \cdot \gamma_j.$$

Rearranging, we then have

$$[\Sigma^{-1} + Q]_{ii} \cdot (-\gamma_i) < \sum_{j \neq i} [\Sigma^{-1} + Q]_{ij} \cdot \gamma_j \leq \sum_{j \neq i} |[\Sigma^{-1} + Q]_{ij}| \cdot |\gamma_j|,$$

which is impossible because $-\gamma_i \geq |\gamma_j|$ for each $j \neq i$ and $[\Sigma^{-1} + Q]_{ii} \geq \sum_{j \neq i} |[\Sigma^{-1} + Q]_{ij}|$. Thus γ_i is positive for any $i \in B$. The result that these γ_i are the same follows from the definition that their absolute values are maximal. \square

C.2 The Last Stage

To prove Theorem 2, we first consider those times t when each of the K sources has been sampled. The following lemma shows that after any such time, it is optimal to maintain a constant attention allocation proportional to α .

Lemma 9. *Suppose Σ^{-1} is diagonally-dominant. If at some time \underline{t} , the \underline{t} -optimal vector satisfies $\partial_1 V(n(\underline{t})) = \dots = \partial_K V(n(\underline{t}))$, then the t -optimal vector at each time $t \geq \underline{t}$ is given by*

$$n(t) = n(\underline{t}) + \frac{t - \underline{t}}{\alpha_1 + \dots + \alpha_K} \cdot \alpha.^{24}$$

²⁴That is, $n_i(t) = n_i(\underline{t}) + \frac{t - \underline{t}}{\alpha_1 + \dots + \alpha_K} \cdot \alpha_i$ for each i .

Proof. Consider increasing $n(\underline{t})$ by a vector proportional to α . If we can show the equalities $\partial_1 V = \dots = \partial_K V$ are preserved, then the resulting cumulative attention vector must be t -optimal. This is because for the convex function V , a vector q minimizes $V(q)$ subject to $q_i \geq 0$ and $\sum_i q_i = t$ if and only if it satisfies the KKT first-order conditions.

We check the equalities $\partial_1 V = \dots = \partial_K V$ by computing the marginal changes of each $\partial_i V$ when the attention vector $q = n(\underline{t})$ increases in the direction of α . Denoting $\text{diag}(q)$ by Q to save notation, this marginal change equals

$$\delta_i := \sum_{j=1}^K \partial_{ij} V \cdot \alpha_j = 2 \sum_{j=1}^K \gamma_i \gamma_j [(\Sigma^{-1} + Q)^{-1}]_{ij} \cdot \alpha_j$$

by Lemma 3. Applying Lemma 8, we have $\gamma_1 = \dots = \gamma_K$. Thus the above simplifies to

$$\delta_i = 2\gamma_1^2 \sum_{j=1}^K [(\Sigma^{-1} + Q)^{-1}]_{ij} \cdot \alpha_j = 2\gamma_1^2 \gamma_i = 2\gamma_1^3.$$

Hence $\partial_1 V = \dots = \partial_K V$ continues to hold, completing the proof. \square

C.3 Earlier Stages

In general, we need to show that even when the agent is choosing from a subset of the sources, the t -optimal vector $n(t)$ is still increasing over time. This is guaranteed by the following lemma, which says that the agent optimally attends to those sources that maximize the marginal reduction of V , until a new source becomes another maximizer. For ease of exposition we work under the stronger assumption that Σ^{-1} is *strictly* diagonally-dominant, in the sense that $[\Sigma^{-1}]_{ii} > \sum_{j \neq i} |[\Sigma^{-1}]_{ij}|$ for all $1 \leq i \leq K$. Later we discuss how the lemma should be modified without this strictness.

Lemma 10. *Suppose Σ^{-1} is strictly diagonally-dominant. Choose any time \underline{t} and denote*

$$B = \text{argmin}_i \partial_i V(n(\underline{t})) = \text{argmax}_i |\gamma_i|.$$

Then there exists $\beta \in \Delta^{K-1}$ supported on B and $\bar{t} > \underline{t}$ such that $n(t) = n(\underline{t}) + (t - \underline{t}) \cdot \beta$ at times $t \in [\underline{t}, \bar{t}]$.

The vector β depends only on Σ, α and B . The time \bar{t} is the earliest time after \underline{t} at which $\text{argmin}_i \partial_i V(n(\bar{t}))$ is a strict superset of B . Moreover, when $|B| < K$ it holds that $\bar{t} < \infty$. Whereas when $|B| = K$, it holds that $\bar{t} = \infty$ and β is proportional to α .

Proof. The case when $|B| = K$ has been proved in Lemma 9, so we only consider $|B| < K$. Without loss we assume $B = \{1, \dots, k\}$ with $1 \leq k < K$. Let $q = n(\underline{t})$ and define γ as before. By Lemma 8, γ_i is the same positive number for $i \leq k$. Moreover, t -optimality implies that $q_j = 0$ whenever $j > k$. Otherwise the posterior variance could be reduced by decreasing q_j and increasing q_1 , as source 1 has strictly higher marginal value than source j .

We now use a trick to deduce the current lemma from the previous Lemma 9. Specifically, given the prior covariance matrix Σ , we can choose another basis of the attributes $\theta_1, \dots, \theta_k, \theta_{k+1}^*, \dots, \theta_K^*$ with two properties:

1. each θ_j^* ($j > k$) is a linear combination of the original attributes $\theta_1, \theta_2, \dots, \theta_K$;
2. $\text{Cov}[\theta_i, \theta_j^*] = 0$ for all $i \leq k < j$, where the covariance is computed according to the prior belief Σ .

Denote by $\tilde{\theta}$ the vector $(\theta_1, \dots, \theta_k)'$, and by θ^* the vector $(\theta_{k+1}^*, \dots, \theta_K^*)'$. The payoff-relevant state $\omega = \alpha' \cdot \theta$ can thus be rewritten as $\tilde{\alpha}' \cdot \tilde{\theta} + \alpha^{*'} \cdot \theta^*$ for some constant coefficient vectors $\tilde{\alpha} \in \mathbb{R}^k$ and $\alpha^* \in \mathbb{R}^{K-k}$. Using property 2 above, we can solve for $\tilde{\alpha}$ from Σ , α and B :

$$\tilde{\alpha} = (\Sigma_{TL})^{-1} \cdot (\Sigma_{TL}, \Sigma_{TR}) \cdot \alpha \quad (11)$$

where Σ_{TL} is the $k \times k$ top-left sub-matrix of Σ and Σ_{TR} is the $k \times (K - k)$ top-right block.

With this transformation, we have reduced the original problem with K sources to a smaller problem with only the first k sources. To see why this reduction is valid, recall that sampling sources $1 \sim k$ only provides information about $\tilde{\theta}$, which is orthogonal to θ^* according to the prior. So as long as the agent has only looked at the first k sources, the transformed attributes continue to satisfy property 2 above (zero covariances) under any posterior belief. It follows that the posterior variance of ω is simply the variance of $\tilde{\alpha}' \cdot \tilde{\theta}$ plus the variance of $\alpha^{*'} \cdot \theta^*$. Since the latter uncertainty cannot be reduced, the agent's objective (at those times when only the first k sources are attended to) is equivalent to minimizing the posterior variance of $\tilde{\omega} = \tilde{\alpha}' \cdot \tilde{\theta}$.

Thus, in this smaller problem, the prior covariance matrix is Σ_{TL} and the payoff weights are $\tilde{\alpha}$. Assuming that $\tilde{\alpha}$ has strictly positive coordinates, we can then apply Lemma 9: As long as the agent attends to the first k sources proportional to $\tilde{\alpha}$, $\partial_1 V = \dots = \partial_k V$ continues to hold.²⁵ Moreover, at $q = n(\underline{t})$, the definition of the set B implies that these k partial

²⁵Lemma 9 implies $\partial_1 \tilde{V} = \dots = \partial_k \tilde{V}$, where $\tilde{V}(q_1, \dots, q_k)$ is the posterior variance of $\tilde{\alpha}' \cdot \tilde{\theta}$ in the smaller problem. But as discussed, \tilde{V} differs from V by a constant, so its derivatives are the same as those of V .

derivatives have greater magnitude (i.e., more negative) than the rest. By continuity, the same comparison holds until some time $\bar{t} > \underline{t}$. Thus, when $t \in [\underline{t}, \bar{t})$, the cumulative attention vector (under this strategy) still satisfies the first-order condition $B = \operatorname{argmin}_{1 \leq i \leq K} \partial_i V$ and $q_j = 0$ for $j \notin B$. Since V is convex, this must be the t -optimal vector as we desire to show. It also follows that $\bar{t} < \infty$, because at $t = \infty$ the minimum possible posterior variance is zero, which cannot be achieved by attending only to a subset B of sources.

It remains to prove that $\tilde{\alpha}_i$ is positive for $1 \leq i \leq k$. To this end, define $\tilde{Q} = \operatorname{diag}(q_1, \dots, q_k)$ to be the $k \times k$ top-left sub-matrix of Q , and let

$$\tilde{\gamma} = ((\Sigma_{TL})^{-1} + \tilde{Q})^{-1} \cdot \tilde{\alpha}. \quad (12)$$

We will show that $\tilde{\gamma}$ is just the first k coordinates of γ . Indeed, for $1 \leq i \leq k$, $\tilde{\gamma}_i$ is by definition the covariance between θ_i and $\tilde{\omega} = \tilde{\alpha}' \cdot \tilde{\theta}$ under the posterior belief at time \underline{t} . Since $\omega = \tilde{\omega} + \alpha^{*'} \cdot \theta^*$, and the vector θ^* is by construction independent of θ_i , we deduce that $\operatorname{Cov}(\theta_i, \tilde{\omega}) = \operatorname{Cov}(\theta_i, \omega)$. Thus $\tilde{\gamma}_i = \gamma_i$ as desired.

Given this, Lemma 8 tells us that $\tilde{\gamma}_i$ is the same positive number for $1 \leq i \leq k$. Rewriting (12) as $\tilde{\alpha} = ((\Sigma_{TL})^{-1} + \tilde{Q}) \cdot \tilde{\gamma}$, we see that $\tilde{\alpha}_i$ is proportional to the i -th row sum of the matrix $(\Sigma_{TL})^{-1} + \tilde{Q}$, which is just the row sum of $(\Sigma_{TL})^{-1}$ plus q_i . By Carlson and Markham (1979), if Σ^{-1} is (strictly) diagonally-dominant, then so is $(\Sigma_{TL})^{-1}$ for any principal sub-matrix Σ_{TL} (because $(\Sigma_{TL})^{-1}$ is the Schur complement of Σ^{-1} with respect to its bottom-right block). So the row sums of $(\Sigma_{TL})^{-1}$ are all strictly positive, implying $\tilde{\alpha}_i > 0$. \square

C.4 Piecing Together Different Stages

We now apply Lemma 10 repeatedly to prove Theorem 2. Continuing to assume strict diagonal dominance, we can apply Lemma 10 with $\underline{t} = 0$ and deduce that up to some time $t_1 = \bar{t} > 0$, t -optimality can be achieved by a constant attention strategy supported on $B_1 = \operatorname{argmin}_{1 \leq i \leq K} \partial_i V(0)$. Applying Lemma 10 again with $\underline{t} = t_1$, we know that the agent can maintain t -optimality from time t_1 to some time t_2 with a constant attention strategy supported on $B_2 = \operatorname{argmin}_{1 \leq i \leq K} \partial_i V(n(t_1))$. So on and so forth. The sets $\emptyset = B_0, B_1, B_2, \dots$ are nested by construction, so eventually $B_m = \{1, \dots, K\}$. This delivers the result.

C.5 The Case of Weak Diagonal Dominance

Here we demonstrate how to prove Theorem 2 assuming only that Σ^{-1} is weakly diagonally-dominant. The new difficulty is that in the proof of Lemma 10, we cannot conclude the optimal attention allocation (which is proportional to $\tilde{\alpha}$) has *strictly* positive coordinates on B . Thus the agent does not necessarily mix over *all* of the sources that maximize marginal reduction of variance. This might lead to the failure of Theorem 2 for two reasons. First, it is possible that the agent optimally divides attention across a *subset* of the sources that he has paid attention to in the past, which would violate the requirement of nested observation sets. Second, when a new source achieves maximal marginal value, the agent might (not attend to it and) use a different mixture over the sources previously sampled, which would violate the requirement of constant attention allocation for a given observation set.

We now show that neither occurs in our setting. In response to the first concern above, note that we can still follow the proof of Lemma 10 to deduce that the optimal instantaneous attention $\tilde{\alpha}_i$ given to a source $i \in \operatorname{argmin}_j \partial_j V(t)$ is proportional to the i -th row sum of $(\Sigma_{TL})^{-1}$ plus q_i . Since $(\Sigma_{TL})^{-1}$ is weakly diagonally-dominant, its row sums are weakly positive. Thus $\tilde{\alpha}_i > 0$ whenever $q_i > 0$. In words, any source that has received attention in the past will be allocated strictly positive attention at every future instant.

To address the second concern, consider two times $\tilde{t} < \hat{t}$ with $\operatorname{argmin}_j \partial_j V(n(\tilde{t})) \subsetneq \operatorname{argmin}_j \partial_j V(n(\hat{t}))$. Reordering the attributes, we may assume that at time \tilde{t} the first \tilde{k} sources have the highest marginal value, whereas at time \hat{t} this set expands to the first $\hat{k} > \tilde{k}$ sources. Let $\tilde{\alpha} \in \mathbb{R}^{\tilde{k}}$ and $\hat{\alpha} \in \mathbb{R}^{\hat{k}}$ be the optimal attentions associated with these subsets, as given by (11). We want to show that if $\hat{\alpha}$ is supported on the same set of sources as $\tilde{\alpha}$, then $\hat{\alpha}$ coincides with $\tilde{\alpha}$ on their support. Indeed, by definition of $\hat{\alpha}$ (see the proof of Lemma 10),

$$\omega = \sum_{i \leq \tilde{k}} \hat{\alpha}_i \theta_i + \text{residual term orthogonal to } \theta_1, \dots, \theta_{\tilde{k}}.$$

If $\hat{\alpha}$ has the same support as $\tilde{\alpha}$, then the above implies

$$\omega = \sum_{i \leq \tilde{k}} \hat{\alpha}_i \theta_i + \text{residual term orthogonal to } \theta_1, \dots, \theta_{\tilde{k}},$$

where we use the fact that any term orthogonal to the first \hat{k} attributes is clearly orthogonal to the first \tilde{k} attributes. This last representation of ω reduces to the definition of $\tilde{\alpha}$. Hence $\hat{\alpha}_i = \tilde{\alpha}_i$ for $1 \leq i \leq \tilde{k}$, as we desire to prove.

C.6 The Case of Perpetual Substitutes or Perpetual Complements

We now prove Theorem 2 under Assumption 4 or 5. Our proof above uses the diagonal dominance assumption at two places. It is crucial for proving Lemma 8 (that is, the coordinates of γ with greatest magnitude are all positive), as well as for showing that the transformed weight vector $\tilde{\alpha}$ is positive in the proof of Lemma 10. Thus we just need to verify these two steps under the alternative assumptions.

Lemma 8 continues to hold because, as we show in the proof of Propositions 7 and 8 in Online Appendix O.3, these alternative assumptions imply that $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$ has non-negative coordinates for any $q \geq 0$. It trivially follows that those coordinates with maximal absolute value must be strictly positive.

As for $\tilde{\alpha}$ in the proof of Lemma 10, first consider Assumption 4 which imposes that Σ^{-1} is an M -matrix. We use the matrix identity $(\Sigma_{TL})^{-1} \cdot \Sigma_{TR} = -(\Sigma^{-1})_{TR} \cdot [(\Sigma^{-1})_{BR}]^{-1}$, which can be proved using $\Sigma \cdot \Sigma^{-1} = I_K$ and comparing the top-right block. By assumption, $(\Sigma^{-1})_{TR}$ has non-positive entries, since it only consists of off-diagonal entries of Σ^{-1} . Moreover, $[(\Sigma^{-1})_{BR}]^{-1}$ has non-negative entries, since $(\Sigma^{-1})_{BR}$ is an M -matrix. We thus conclude that $(\Sigma_{TL})^{-1} \cdot \Sigma_{TR}$ is a matrix with non-negative entries. From (11) we have $\tilde{\alpha} = (\Sigma_{TL})^{-1} \cdot (\Sigma_{TL}, \Sigma_{TR}) \cdot \alpha = (I_k, (\Sigma_{TL})^{-1} \Sigma_{TR}) \cdot \alpha$. So each coordinate of $\tilde{\alpha}$ is larger than the corresponding coordinate of α , and is thus strictly positive.

If instead Assumption 5 is satisfied, then Σ itself is an M -matrix, and so is the principal sub-matrix Σ_{TL} . Thus $(\Sigma_{TL})^{-1}$ has non-negative entries off the diagonal and strictly positive entries on the diagonal. From (12) we have $\tilde{\alpha} = ((\Sigma_{TL})^{-1} + \tilde{Q}) \cdot \tilde{\gamma}$. Since $\tilde{\gamma}$ is a positive vector (with equal coordinates), we deduce $\tilde{\alpha} \gg 0$. This completes the proof of Theorem 2.

D Algorithm for Computing the Optimal Strategy

Here we provide an algorithm for recursively finding the times t_k and sets B_k in Theorem 2. Set Q_0 to be the $K \times K$ matrix of zeros, and $t_0 = 0$. For each stage $k \geq 1$:

1. Computation of the observation set B_k . Define the $K \times 1$ vector $\gamma^k = (\Sigma^{-1} + Q_{k-1})^{-1} \cdot \alpha$ where Σ is the prior covariance matrix, and α is the weight vector. The set of attributes that the agent attends to in stage k is

$$B_k = \operatorname{argmax}_i |\gamma_i^k|.$$

These sources have highest marginal reduction of posterior variance (see Lemma 3).

2. Computation of the constant attention allocation in stage k . If B_k is the set of all sources, then we are already in the last stage and the algorithm ends. Otherwise let $\ell = |B_k| < K$. We can re-order the attributes so that the ℓ attributes in B_k are the first ℓ attributes. In an abuse of notation, let Σ be the covariance matrix for the re-ordered attribute vector θ . Define Σ_{TL} to be the $\ell \times \ell$ top-left submatrix of Σ and Σ_{TR} to be the $\ell \times (K - \ell)$ top-right block. Finally let

$$\alpha^k = (\Sigma_{TL})^{-1} \cdot (\Sigma_{TL}, \Sigma_{TR}) \cdot \alpha$$

be an $\ell \times 1$ vector. The agent's optimal attention allocation in stage k is proportional to α^k :

$$\beta_i^k = \begin{cases} \alpha_i^k / \sum_i \alpha_i^k & \text{if } i \leq \ell \\ 0 & \text{otherwise} \end{cases}$$

As the agent acquires information in this mixture during stage k , the marginal values of learning about different attributes in B_k remain the same, and strictly higher than learning about any attribute outside of the set.

3. Computation of the next time t_k . For arbitrary t , define

$$Q^k(t) := Q_{k-1} + (t - t_{k-1}) \cdot \text{diag}(\beta^k).$$

Let t_k be the smallest $t > t_{k-1}$ such that the coordinates maximizing $(\Sigma^{-1} + Q^k(t))^{-1} \cdot \alpha$ are a strict superset of B_k . At this time, the marginal value of some attribute(s) outside of B_k equalizes the attributes in B_k , and stage $k + 1$ commences, with $Q_k = Q^k(t_k)$.

The time t_k can be computed as follows. For each source $j > \ell$, consider the following (polynomial) equation in t :

$$e'_j \cdot (\Sigma^{-1} + Q^k(t))^{-1} \cdot \alpha = \pm e'_1 \cdot (\Sigma^{-1} + Q^k(t))^{-1} \cdot \alpha.$$

Any solution $t > t_{k-1}$ is a time at which source j has the same marginal value as sources $1, \dots, \ell$. So t_k is the smallest such solution t across all $j > \ell$.

References

- AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.
- ALARIE, Y. AND C. BRONSARD (1990): “Preferences and Normal Goods: A Necessary and Sufficient Condition,” *Journal of Economic Theory*, 51, 423–430.
- AZEVEDO, E., A. DENG, J. L. M. OLEA, J. RAO, AND G. WEYL (2020): “A/B Testing with Fat Tails,” *Journal of Political Economy*, 128, 4614–4672.
- BARDHI, A. (2020): “Attributes: Selective Learning and Influence,” Working Paper.
- BERGEMANN, D. AND J. VÄLIMÄKI (2008): “Bandit Problems,” in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Basingstoke: Palgrave Macmillan.
- BILANCINI, E. AND L. BONCINELLI (2010): “Preferences and normal goods: An easy-to-check necessary and sufficient condition,” *Economics Letters*, 108, 13–15.
- BLACKWELL, D. (1951): “Comparison of Experiments,” in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley and Los Angeles: University of California Press.
- BUBECK, S., R. MUNOS, AND G. STOLTZ (2009): “Pure Exploration in Multi-armed Bandits Problems,” in *Algorithmic Learning Theory. ALT 2009. Lecture Notes in Computer Science, vol 5809*, Springer, Berlin, Heidelberg.
- CALLANDER, S. (2011): “Searching and Learning by Trial and Error,” *American Economic Review*, 101, 2277–2308.
- CARLSON, D. AND T. L. MARKHAM (1979): “Schur Complements of Diagonally Dominant Matrices,” *Czechoslovak Mathematical Journal*, 29, 246–251.
- CHE, Y.-K. AND K. MIERENDORFF (2019): “Optimal Dynamic Allocation of Attention,” *American Economic Review*, 108, 2993–3029.
- CHEN, H. AND W. SUEN (2019): “Competition for Attention and News Quality,” Working Paper.
- CHICK, S. E. AND P. FRAZIER (2012): “Sequential Sampling with Economics of Selection Procedures,” *Management Science*, 58, 550–569.
- CHIPMAN, J. S. (1977): “An Empirical Implication of Auspitz-Lieben-Edgeworth-Pareto Complementarity,” *Journal of Economic Theory*, 14, 228–231.

- EASLEY, D. AND N. M. KIEFER (1988): “Controlling a Stochastic Process with Unknown Parameters,” *Econometrica*, 56, 1045–1064.
- FRAZIER, P., W. POWELL, AND S. DAYANIK (2008): “A Knowledge-Gradient Policy for Sequential Information Collection,” *SIAM Journal of Control and Optimization*, 47, 2410–2439.
- (2009): “The Knowledge-Gradient Policy for Correlated Normal Beliefs,” *Informations Journal on Computing*, 21, 599–613.
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2018): “Speed, Accuracy, and the Optimal Timing of Choices,” *American Economic Review*, 108, 3651–3684.
- GALPERTI, S. AND I. TREVINO (2020): “Coordination Motives and Competition for Attention in Information Markets,” *Journal of Economic Theory*, 188.
- GARFAGNINI, U. AND B. STRULOVICI (2016): “Social Experimentation with Interdependent and Expanding Technologies,” *Review of Economic Studies*, 83, 1579–1613.
- GENTZKOW, M. AND J. SHAPIRO (2008): “Competition and Truth in the Market for News,” *Journal of Economic Perspectives*, 22, 133–154.
- GITTINS, J. C. (1979): “Bandit Processes and Dynamic Allocation Indices,” *Journal of the Royal Statistical Society, Series B*, 148–177.
- GOSSNER, O., J. STEINER, AND C. STEWART (2020): “Attention Please!” *Econometrica*, forthcoming.
- GREENSHTEIN, E. (1996): “Comparison of Sequential Experiments,” *The Annals of Statistics*, 24, 436–448.
- HANSEN, O. H. AND E. N. TORGERSEN (1974): “Comparison of Linear Normal Experiments,” *The Annals of Statistics*, 2, 367–373.
- HÉBERT, B. AND M. WOODFORD (2021): “Rational Inattention When Decisions Take Time,” Working Paper.
- JOHNSON, C. R. (1982): “Inverse M-Matrices,” *Linear Algebra and Its Applications*, 47, 195–216.
- KE, T. T. AND J. M. VILLAS-BOAS (2019): “Optimal Learning Before Choice,” *Journal of Economic Theory*, 180, 383–437.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73, 2787–2802.
- KLABJAN, D., W. OLSZWSKI, AND A. WOLINSKY (2014): “Attributes,” *Games and Eco-*

- nomic Behavior*, 88, 190–206.
- LIANG, A., X. MU, AND V. SYRGKANIS (2019): “Optimal and Myopic Information Acquisition,” Working Paper.
- MACKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2021): “Rational Inattention: A Review,” *Journal of Economic Literature*, forthcoming.
- MAYSKAYA, T. (2020): “Dynamic Choice of Information Sources,” Working Paper.
- MORRIS, S. AND P. STRACK (2019): “The Wald Problem and the Relation of Sequential Sampling and Ex-ante Information Costs,” Working Paper.
- NIEUWERBURGH, S. V. AND L. VELDKAMP (2010): “Information Acquisition and Under-Diversification,” *Review of Economic Studies*, 77, 779–805.
- PEREGO, J. AND S. YUKSEL (2021): “Media Competition and Social Disagreement,” *Econometrica*, forthcoming.
- PLEMMONS, R. (1977): “M-Matrix Characterizations.I—Nonsingular M-Matrices,” *Linear Algebra and Its Applications*, 18, 175–188.
- QUAH, J. K.-H. (2007): “The Comparative Statics of Constrained Optimization Problems,” *Econometrica*, 75, 401–431.
- RATCLIFF, R. AND G. MCKOON (2008): “The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks,” *Neural Computation*, 20, 873–922.
- REVUZ, D. AND M. YOR (1999): *Continuous Martingales and Brownian Motion*, New York: Springer.
- RUSSO, D. (2016): “Simple Bayesian Algorithms for Best-Arm Identification,” *Operations Research*, 68, 1625–1647.
- SANJURJO, A. (2017): “Search with Multiple Attributes: Theory and Empirics,” *Games and Economic Behavior*, 104, 535–562.
- STEINER, J., C. STEWART, AND F. MATĚJKA (2017): “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” *Econometrica*, 85, 521–553.
- YANG, M. (2015): “Coordination with Flexible Information Acquisition,” *Journal of Economic Theory*, 158, 721–738.
- YONG, J. AND X. Y. ZHOU (1999): *Stochastic Controls: Hamiltonian Systems and HJB Equations*, New York: Springer.
- ZHONG, W. (2019): “Optimal Dynamic Information Acquisition,” Working Paper.

O For Online Publication

O.1 Uniqueness of Optimal Information Acquisition

By Lemma 5, whenever a uniformly optimal strategy exists, it is the optimal attention strategy regardless of the form of $u(\tau, a, \omega)$. Without further assumptions on u , there could exist other optimal attention strategies. For example, consider the payoff function used in the proof of Lemma 6. Under this payoff function, the agent always stops at some fixed time t . Hence any strategy that achieves the t -optimal vector $n(t)$ would be optimal in this problem.

Nonetheless, such examples can be ruled out by an assumption on the agent's stopping rule:

Assumption 7. *Given any attention allocation strategy S , any history of signal realizations up to time t such that the agent has not stopped, and any $t' > t$, there exists a positive probability of continuation histories such that the agent optimally stops in the interval $(t, t']$.*

Proposition 5. *Suppose Assumption 2 holds strictly, and Assumption 7 is satisfied. Then, any optimal attention strategy S induces the same posterior variance as the uniformly optimal strategy S^* at every history where the agent has not stopped. Consequently, the two strategies induce the same cumulative attention vectors, and coincide at almost every time before stopping.*

Proof. Suppose S induces larger posterior variance than S^* at some time t , then by continuity the same holds from time t to some later time $t' > t$ along this history. By assumption, the agent stops between these times with positive probability. Thus there is positive probability that the agent stops with posterior variance strictly larger than the minimal variance. From the proof of Lemma 5, we see that payoff under S is strictly below S^* , contradicting the optimality of S . The second part of the result follows from the uniqueness of $n(t)$, and the fact that $\beta(t)$ integrates to $n(t)$. \square

Although Assumption 7 is stated in terms of the endogenous stopping rule, it is satisfied in any problem where the agent always stops to take some action when he has an extremely high (or low) expectation about ω . This is guaranteed if extreme values of ω agree on the optimal action and the marginal cost of delay is bounded away from zero. These conditions on the primitives are rather weak, and are satisfied in many applications such as binary choice with linear waiting cost.

O.2 Non-existence of Uniformly Optimal Strategy

O.2.1 Counterexample for $K = 2$

The example below illustrates how and why Theorem 1 might fail without Assumption 3:

Example 5. There are two unknown attributes with prior distribution

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} 10 & -3 \\ -3 & 1 \end{pmatrix} \right).$$

The agent wants to learn $\theta_1 + 4\theta_2$.

Given q_1 units of attention devoted to learning θ_1 , and q_2 devoted to θ_2 , the agent's posterior variance about ω is given by (2). Simplifying, we have $V(q_1, q_2) = \frac{2+16q_1+q_2}{(1+q_1)(10+q_2)-9}$. The t -optimal cumulative attention vectors $n(t)$ (see Section 4) are defined to minimize $V(q_1, q_2)$ subject to $q_1, q_2 \geq 0$ and the budget constraint $q_1 + q_2 \leq t$.

These vectors do not evolve monotonically: Initially, the marginal value of learning θ_1 exceeds that of learning θ_2 , since the agent has greater prior uncertainty about θ_1 (even accounting for the difference in payoff weights). Thus at all times $t < 1/4$, the t -optimal vector is $(t, 0)$, and the agent learns only about attribute 1.

After a quarter-unit of time devoted to learning θ_1 , the agent's posterior covariance matrix becomes $\begin{pmatrix} 20/7 & -6/7 \\ -6/7 & 5/14 \end{pmatrix}$. Note that the two sources have equal marginal values at $t = 1/4$, since $\gamma_1 = \frac{-4}{7}$ and $\gamma_2 = \frac{4}{7}$ have the same absolute value. However, to maintain equal marginal values at future instants, it would be optimal to take attention away from attribute 1 and redistribute it to attribute 2. Specifically, at all times $t \in [1/4, 1]$ the t -optimal vector is given by $n(t) = (\frac{-t+1}{3}, \frac{4t-1}{3})$, and the optimal cumulative attention toward θ_1 is decreasing in this interval.

Consequently, there does not exist a uniformly optimal strategy in this example (Lemma 1). Hence the optimal information acquisition strategy varies according to when the agent expects to stop, and Theorem 1 cannot hold independently of the payoff criterion.

O.2.2 Necessity of Assumption 3 for Theorem 1

We show here that when $K = 2$, the assumption $cov_1 + cov_2 \geq 0$ is also *necessary* for the existence of a uniformly optimal strategy. The result generalizes Example 5 above.

Proposition 6. *Suppose $K = 2$ and Assumption 3 is violated. Then a uniformly optimal strategy does not exist.*

Proof. Suppose that $cov_1 + cov_2 < 0$. First note that one of cov_1, cov_2 is positive, because $\alpha_1 cov_1 + \alpha_2 cov_2 = \alpha' \Sigma \alpha > 0$. So without loss we can assume $cov_2 > 0 > -cov_2 > cov_1$. Moreover, from $\alpha_1 cov_1 + \alpha_2 cov_2 > 0$ we obtain $\alpha_2 > \alpha_1$ and hence $x_2 > x_1$. Below we show the t -optimal attention vector $n(t)$ is non-monotonic.

Suppose $\frac{-(cov_1+cov_2)}{x_2} < t < \frac{-(cov_1+cov_2)}{x_1}$, then by (10) we have $\partial_1 V(0, t) < \partial_2 V(0, t)$ and $\partial_1 V(t, 0) > \partial_2 V(t, 0)$. These imply that $n(t)$ is interior, and the first-order condition yields

$$x_1 n_2(t) + cov_1 = -(x_2 n_1(t) + cov_2),$$

where we use the fact that for t in this range $x_1 q_2 + cov_1$ is always negative. Together with $n_1(t) + n_2(t) = t$, we can solve that $n(t) = (\frac{-x_1 t - cov_1 - cov_2}{x_2 - x_1}, \frac{x_2 t + cov_1 + cov_2}{x_2 - x_1})$. Thus as t increases in this range, $n_1(t)$ actually decreases. So a uniformly optimal strategy does not exist. \square

O.2.3 Counterexample for $K = 3$

We present another example where a uniformly optimal strategy does not exist. In the following example, the three attributes have positive correlation with each other, but not positive *partial correlation*. Thus the example illustrates the subtlety of Assumption 4.

Let the primitives be $K = 3$, $\alpha_1 = \alpha_2 = 1$, $\alpha_3 = 20$ and

$$\Sigma = \begin{pmatrix} 19 & 3 & 0 \\ 3 & 5 & 3 \\ 0 & 3 & 2 \end{pmatrix} \quad \Sigma^{-1} = \begin{pmatrix} 1 & -6 & 9 \\ -6 & 38 & -57 \\ 9 & -57 & 86 \end{pmatrix}$$

The choice of $\Sigma_{11} = 19$ is such that Σ has determinant exactly 1. This makes it easier to calculate its inverse matrix, while still ensuring that Σ is positive-definite. We highlight the fact that $[\Sigma^{-1}]_{13} = 9$ is positive, so this example does not satisfy Assumption 4.

Consider the cumulative attention vector $q = (1, 14, 0)'$. Simple calculation gives $(\Sigma^{-1} + Q) \cdot (-1, 1, 1)' = (1, 1, 20)' = \alpha$. Thus $\gamma(q) = (\Sigma^{-1} + Q)^{-1} \alpha = (-1, 1, 1)'$. Since the three coordinates of γ have equal absolute value, the sources have equal marginal reduction of V at the attention vector q . This means q is the t -optimal vector for $t = 15$.

Next consider a different attention vector $\hat{q} = (0, 15, 20)$. We can similarly calculate that $(\Sigma^{-1} + \hat{Q}) \cdot (-1, 1, 1)' = (2, 2, 40)' = 2\alpha$. So $\gamma(\hat{q}) = (-1/2, 1/2, 1/2)$. By the same reasoning, \hat{q} is the t -optimal vector for $t = 35$. Hence we see that when t increases from 15 to 35, the optimal amount of attention devoted to θ_1 decreases from 1 to 0. This implies that a uniformly optimal strategy does not exist in this example.

O.3 When are Sources Substitutes/Complements?

Since more information *reduces* the posterior variance V , we define two sources i and j to be substitutes if the cross-partial $\partial_{ij} V(q)$ is *non-negative* at any cumulative attention vector $q \geq 0$. The following result shows that Assumption 4 precisely characterizes when two sources are substitutes.

Proposition 7. *Given any positive payoff weight vector α . The following conditions on the prior covariance matrix Σ are equivalent:*

1. Σ^{-1} has non-positive off-diagonal entries;
2. Every pair of sources $i \neq j \in \{1, \dots, K\}$ are substitutes in the sense that $\partial_{ij}V(q) \geq 0$ at every cumulative attention vector $q \in \mathbb{R}_+^K$.

Proof. In one direction, suppose Σ^{-1} has non-positive off-diagonal entries. Then for any vector $q \geq 0$ and corresponding diagonal matrix $Q = \text{diag}(q)$, $\Sigma^{-1} + Q$ satisfies the same property. Thus the positive-definite matrix $\Sigma^{-1} + Q$ is an M -matrix whose inverse is known to have non-negative entries off the diagonal and strictly positive entries on the diagonal; see e.g. Plemmons (1977). It follows that $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$ is a vector with positive entries. Thus by Lemma 3,

$$\partial_{ij}V(q) = 2\gamma_i\gamma_j \cdot [(\Sigma^{-1} + Q)^{-1}]_{ij} \geq 0.$$

Conversely, suppose $\partial_{ij}V(q) \geq 0$ for every $q \geq 0$. Let us choose any vector q whose coordinates are all sufficiently large. Then

$$(\Sigma^{-1} + Q)^{-1} = [(\Sigma^{-1} \cdot Q^{-1} + I_K) \cdot Q]^{-1} = Q^{-1}(\Sigma^{-1} \cdot Q^{-1} + I_K)^{-1} = Q^{-1}(I_K + o(1)),$$

where $o(1)$ is a matrix that goes to zero as q_1, \dots, q_K all go to infinity (at any rate). Note that when the $o(1)$ term is sufficiently small, $(I_K + o(1)) \cdot \alpha$ is a vector with positive coordinates. Thus, $\gamma = (\Sigma^{-1} + Q)^{-1}\alpha = Q^{-1}(I_K + o(1))\alpha$ has positive coordinates. Together with $\partial_{12}V(q) \geq 0$ and Lemma 3, this implies $[(\Sigma^{-1} + Q)^{-1}]_{12} \geq 0$ for any such q . Using the matrix inverse formula, we further obtain that the determinant of the cofactor matrix $[\Sigma^{-1} + Q]_{-21}$ (i.e., the sub-matrix of $\Sigma^{-1} + Q$ with the second row and first column removed) must be non-positive. Expanding this determinant using permutations, it is easy to see that it contains the term $[\Sigma^{-1}]_{12} \cdot \prod_{k=3}^K q_k$, which is in fact the dominant term when each q_k is sufficiently large. Hence we deduce $[\Sigma^{-1}]_{12} \leq 0$, and similarly $[\Sigma^{-1}]_{ij} \leq 0$ for every pair $i \neq j$. \square

We next show that Assumption 5 characterizes when two sources are complements.

Proposition 8. *Given any positive payoff weight vector α . The following conditions on the prior covariance matrix Σ are equivalent:*

1. Σ has non-positive off-diagonal entries and $\Sigma \cdot \alpha$ has non-negative coordinates;
2. Every pair of sources $i \neq j \in \{1, \dots, K\}$ are complements in the sense that $\partial_{ij}V(q) \leq 0$ at every cumulative attention vector $q \in \mathbb{R}_+^K$.

Proof. To show the first condition implies the second, we note that since Σ is assumed to be an M -matrix, Σ^{-1} is an inverse M -matrix. By Theorem 3 in [Johnson \(1982\)](#), $\Sigma^{-1} + Q$ is also an inverse M -matrix. This implies that for any attention vector q , the posterior covariance matrix $(\Sigma^{-1} + Q)^{-1}$ must be an M -matrix with non-positive off-diagonal entries. We claim that the vector $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$ has non-negative coordinates. Once this is shown, Lemma 3 will imply that $\partial_{ij}V = 2\gamma_i\gamma_j[(\Sigma^{-1} + Q)^{-1}]_{ij} \leq 0$ for every pair $i \neq j$.

By assumption, if q is the zero vector, then $\gamma = \Sigma \cdot \alpha$ indeed has non-negative coordinates. Our goal below is to show this also holds at any $q \geq 0$. To do this, we first work under the stronger assumption that $\Sigma \cdot \alpha$ has strictly positive coordinates, so that $\gamma(0)$ is a positive vector. Recall from the proof of Lemma 3 that when γ_i is viewed as a function of q , its partial derivatives are given by

$$\frac{\partial \gamma_i}{\partial q_j} = -\gamma_j \cdot [(\Sigma^{-1} + Q)^{-1}]_{ij} \quad \text{for each } j \in \{1, \dots, K\}.$$

Now since Σ^{-1} is positive-definite, we can choose $\epsilon > 0$ such that $\Sigma^{-1} \geq \epsilon I_K$ in the matrix order. Fixing this ϵ , we consider the K functions

$$f_i(q_1, \dots, q_K) = (q_i + \epsilon) \cdot \gamma_i(q_1, \dots, q_K).$$

For every $j \neq i$, $\partial f_i / \partial q_j = (q_i + \epsilon) \cdot -\gamma_j \cdot [(\Sigma^{-1} + Q)^{-1}]_{ij}$. This product is non-negative whenever $\gamma_j \geq 0$, since $[(\Sigma^{-1} + Q)^{-1}]_{ij} \leq 0$ as shown above. On the other hand, by the product rule,

$$\frac{\partial f_i}{\partial q_i} = \gamma_i \cdot [1 - (q_i + \epsilon)[(\Sigma^{-1} + Q)^{-1}]_{ii}.$$

This is non-negative whenever $\gamma_i \geq 0$, since $[(\Sigma^{-1} + Q)^{-1}]_{ii} \leq [(\epsilon I_K + Q)^{-1}]_{ii} = (q_i + \epsilon)^{-1}$ where the inequality uses standard properties of the matrix order.

Hence, we have shown that whenever $f_1(q), \dots, f_K(q)$ are all non-negative, their derivatives with respect to each q_i are also non-negative. Moreover, we know from our stronger assumption that f_1, \dots, f_K are strictly positive at $q = 0$. These together imply $f_1(q), \dots, f_K(q)$ are always strictly positive, by the following argument. Suppose for contradiction that there exists some i and some $q \geq 0$ such that $f_i(q) \leq 0$. By continuity of f and a limit argument, we may assume q is minimal in the sense that at every $q' < q$, $f_j(q')$ is positive for every j . Thus, if we let $g(t) = f_i(t \cdot q)$ be the one-variable function defined for $t \in [0, 1]$, we have $g(0) > 0$, $g'(t) \geq 0$ for $t \in [0, 1)$ and $g(1) \leq 0$. This contradicts the Mean Value Theorem.

It follows that $\gamma(0)$ being positive implies $\gamma(q)$ is positive. We now extend this result to the case of weak inequalities. Suppose $\gamma(0) = \Sigma \cdot \alpha$ has non-negative coordinates. Then by considering $\Sigma + \delta I_K$ instead of Σ , we see that the corresponding γ vector is strictly positive at $q = 0$. The above analysis applied to the M -matrix $\Sigma + \delta I_K$ thus implies $[(\Sigma + \delta I_K)^{-1} + Q]^{-1} \cdot \alpha$ is positive for

any attention vector $q \geq 0$. Letting $\delta \rightarrow 0$ yields $[\Sigma^{-1} + Q]^{-1} \cdot \alpha \geq 0$, as we desire to show. This completes the proof that the first condition in the proposition guarantees complementarity.

Turning to the converse, we assume at every $q \geq 0$, $\partial_{ij}V = 2\gamma_i(q) \cdot \gamma_j(q) \cdot [(\Sigma^{-1} + Q)^{-1}]_{ij} \leq 0$. Let us choose q such that $\gamma_i(q)$ is nonzero for each i ; its existence will be verified later. For this q , we claim that $\gamma_i(q)$ must be positive for each i . Indeed, since $\gamma = (\Sigma^{-1} + Q)^{-1}\alpha$, we have $\alpha'\gamma = \alpha'(\Sigma^{-1} + Q)^{-1}\alpha > 0$. Thus γ must have at least one positive coordinate. Suppose for contradiction that γ has a negative coordinate, then we can without loss assume γ_i is positive for $i \leq k$ and negative for $i > k$, where k satisfies $1 \leq k < K$. Using $\partial_{ij}V \leq 0$, we deduce $[(\Sigma^{-1} + Q)^{-1}]_{ij} \geq 0$ for every pair $i \leq k$ and $j > k$. Now let us decompose $(\Sigma^{-1} + Q)^{-1}$ into four block sub-matrices: $[(\Sigma^{-1} + Q)^{-1}]_{TL}$, $[(\Sigma^{-1} + Q)^{-1}]_{TR}$, $[(\Sigma^{-1} + Q)^{-1}]_{BL}$ and $[(\Sigma^{-1} + Q)^{-1}]_{BR}$ are the top-left $k \times k$, top-right $k \times (K - k)$, bottom-left $(K - k) \times k$ and bottom-right $(K - k) \times (K - k)$ sub-matrices of $(\Sigma^{-1} + Q)^{-1}$ respectively. Recall that $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$. By looking at the last $K - k$ coordinates, we obtain

$$\begin{aligned} (\gamma_{k+1}, \dots, \gamma_K)' &= ([(\Sigma^{-1} + Q)^{-1}]_{BL}, [(\Sigma^{-1} + Q)^{-1}]_{BR}) \cdot \alpha \\ &= [(\Sigma^{-1} + Q)^{-1}]_{BL} \cdot (\alpha_1, \dots, \alpha_k)' + [(\Sigma^{-1} + Q)^{-1}]_{BR} \cdot (\alpha_{k+1}, \dots, \alpha_K)'. \end{aligned}$$

The preceding analysis tells us that $[(\Sigma^{-1} + Q)^{-1}]_{BL}$ has non-negative entries, so the vector $[(\Sigma^{-1} + Q)^{-1}]_{BL} \cdot (\alpha_1, \dots, \alpha_k)'$ is non-negative. In addition, $[(\Sigma^{-1} + Q)^{-1}]_{BR}$ is positive-definite, so the vector $[(\Sigma^{-1} + Q)^{-1}]_{BR} \cdot (\alpha_{k+1}, \dots, \alpha_K)'$ has at least one positive coordinate. Thus, the above displayed equation contradicts the assumption that $\gamma_{k+1}, \dots, \gamma_K$ are all negative.

We thus know that if $\gamma(0)$ has nonzero coordinates, then it is in fact a positive vector. Complementarity further requires $\Sigma_{ij} \leq 0$ for all $i \neq j$, which would complete the proof. In the general case, $\gamma(0)$ may have some coordinates equal to zero, so we instead look for q close to the zero vector such that $\gamma(q)$ has nonzero coordinates. To see why such q exists, note that we can calculate the matrix inverse $(\Sigma^{-1} + Q)^{-1}$ using the determinants of cofactor matrices. From this we see that modulo a multiplicative factor of $[\det(\Sigma^{-1} + Q)]^{-1}$, each $\gamma_i(q) = (\Sigma^{-1} + Q)^{-1}\alpha$ is a nonzero multi-linear polynomial in the $K - 1$ variables $\{q_j\}_{j \neq i}$ (with leading term $\alpha_i \cdot \prod_{j \neq i} q_j$). Thus, the set of vectors q that make $\gamma_i(q)$ equal to zero has measure zero. It follows that we can choose q with arbitrarily small coordinates, such that $\gamma_i(q)$ is nonzero for all i . By the earlier analysis, $\gamma(q)$ is a positive vector, and $[(\Sigma^{-1} + Q)^{-1}]_{ij} \leq 0$ for all $i \neq j$. Letting $q \rightarrow 0$, we conclude by continuity that $\gamma(0) \geq 0$ and $\Sigma_{ij} \leq 0$. This is what we desire to show. \square

O.4 Proof of Proposition 1

O.4.1 Preliminary Analysis

Suppose the agent's prior about the two unknown payoffs is normal with covariance matrix $\begin{pmatrix} \Sigma_{11} & -\Sigma_{12} \\ -\Sigma_{21} & \Sigma_{22} \end{pmatrix}$. Throughout we assume $\Sigma_{11} \geq \Sigma_{22}$. The objective is to maximize

$$\mathbb{E}[\mathbb{E}[\max\{v_1, v_2\} \mid \mathcal{F}_\tau] - c\tau],$$

To reduce this problem to our main model, we use $\max\{v_1, v_2\} = \max\{v_1 - v_2, 0\} + v_2$ to rewrite the objective as

$$\mathbb{E}[\mathbb{E}[\max\{v_1 - v_2, 0\} \mid \mathcal{F}_\tau] - c\tau] + \mathbb{E}[\mathbb{E}[v_2 \mid \mathcal{F}_\tau]]$$

The posterior expectations of v_2 , $M_t = \mathbb{E}[v_2 \mid \mathcal{F}_t]$, form a continuous martingale with continuous paths. Moreover, the family $\{M_t\}$ are uniformly integrable because they are conditional expectations of an integrable random variable v_2 . Thus we can apply Doob's Optional Sampling Theorem to deduce $\mathbb{E}[\mathbb{E}[v_2 \mid \mathcal{F}_\tau]] = \mathbb{E}[v_2]$, which is just the prior expectation of v_2 (and does not depend on the agent's strategy). It follows that the agent simply maximizes

$$\mathbb{E}[\mathbb{E}[\max\{v_1 - v_2, 0\} \mid \mathcal{F}_\tau] - c\tau].$$

As a corollary, the payoff difference $v_1 - v_2$ is a sufficient statistic for the agent's decision. Now if we let $\theta_1 = v_1, \theta_2 = -v_2$, then the prior covariance matrix about θ is simply $\Sigma := \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$. This returns our main model with prior covariance matrix Σ and payoff-relevant state $\omega = v_1 - v_2 = \theta_1 + \theta_2$. Since $\alpha_1 = \alpha_2 = 1$, our Theorem 1 applies and yields Corollary 1.

For the subsequent analysis, we need to keep track of how fast the posterior variance of ω evolves over time. These (minimal) posterior variances are given below:

Lemma 11. *Suppose $\Sigma_{11} \geq \Sigma_{22}$. When adopting the optimal information acquisition strategy, the agent's posterior variance of $\omega = \theta_1 + \theta_2$ at time t is given by*

$$\sigma_t^2 = \begin{cases} \frac{\Sigma_{11} + \Sigma_{22} + 2\Sigma_{12} + \det(\Sigma)t}{1 + \Sigma_{11}t} & \text{if } t \leq t_1^*; \\ \frac{4 \det(\Sigma)}{\Sigma_{11} + \Sigma_{22} - 2\Sigma_{12} + \det(\Sigma)t} & \text{if } t \geq t_1^*. \end{cases}$$

Proof. At time $t \leq t_1^*$, the posterior covariance matrix is

$$\left[\begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}^{-1} + \begin{pmatrix} t & 0 \\ 0 & 0 \end{pmatrix} \right]^{-1} = \frac{1}{1 + \Sigma_{11}t} \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} + \det(\Sigma)t \end{pmatrix}.$$

This gives the first part of the lemma.

At time $t = t_1^*$, the levels of uncertainty about θ_1 and θ_2 have equalized. From this time on, each unit of time produces two normal signals $\theta_1 + \epsilon_1$ and $\theta_2 + \epsilon_2$ with ϵ_1 and ϵ_2 independently and identically distributed according to $\mathcal{N}(0, 2)$. These two signals are informationally equivalent to their sum and difference $\theta_1 + \theta_2 + \epsilon_1 + \epsilon_2$ and $\theta_1 - \theta_2 + \epsilon_1 - \epsilon_2$, with i.i.d. noise terms $\epsilon_1 + \epsilon_2, \epsilon_1 - \epsilon_2 \sim \mathcal{N}(0, 4)$. Note that as long as the agent's uncertainty about θ_1 and θ_2 are the same, $\theta_1 + \theta_2$ and $\theta_1 - \theta_2$ are independent. Thus, in terms of learning about $\theta_1 + \theta_2$, it is as if the agent receives only the signal $\theta_1 + \theta_2 + \mathcal{N}(0, 4)$ over each unit of time. This observation enables us to calculate the posterior variance of $\theta_1 + \theta_2$ as follows: For any $t \geq t_1^*$,

$$\sigma_t^2 = \left(\frac{1}{\sigma_{t_1^*}^2} + \frac{t - t_1^*}{4} \right)^{-1},$$

where $\frac{1}{\sigma_{t_1^*}^2}$ is the belief precision of $\theta_1 + \theta_2$ at time t_1^* and $\frac{t - t_1^*}{4}$ is the signal precision from those $\theta_1 + \theta_2 + \mathcal{N}(0, 4)$ signals between time t_1^* and time t . Plugging in $t_1^* = \frac{\Sigma_{11} - \Sigma_{22}}{\det(\Sigma)}$ and $\sigma_{t_1^*}^2 = \frac{2 \det(\Sigma)}{\Sigma_{11} - \Sigma_{12}}$ then yields the second part of the lemma. \square

O.4.2 Stopping Boundaries and Choice Accuracy

Using the posterior variances characterized above, we can write down the process for the posterior expectation of ω , which we denote by $Y_t = \mathbb{E}[\theta_1 + \theta_2 \mid \mathcal{F}_t]$:

$$Y_t = Y_0 + \int_0^t \sqrt{\frac{\partial \sigma_s^2}{\partial s}} \cdot dB_s, \quad (13)$$

where B_s is a standard Brownian motion with respect to the filtration of the agent's information. The volatility term $\sqrt{\frac{\partial \sigma_s^2}{\partial s}}$ is such that the variance of posterior expectation Y_t matches the reduction in posterior variance $\sigma_0^2 - \sigma_t^2$. This representation is a direct generalization of Lemma 1 in [Fudenberg et al. \(2018\)](#) and follows from standard results.

Therefore, given any prior covariance matrix Σ and any prior expectation $Y_0 = y$, the agent's problem can be rewritten as

$$\max_{\tau} \mathbb{E}[\max\{Y_{\tau}, 0\} - c\tau], \quad (14)$$

where τ can be any stopping time adapted to the Y process given above. We now define the value function $U(y, c, \Sigma)$ to be the agent's maximal payoff in this problem. It is easy to see that U is non-negative, increasing in y and decreasing in c (we refer to weak monotonicity, unless otherwise specified). In addition, just as in [Fudenberg et al. \(2018\)](#), the *stopping boundary* at time $t = 0$ is symmetric and given by

$$k^*(c, \Sigma) = \min\{x > 0 : U(-x, c, \Sigma) = 0\}.$$

What this means is that if the prior expectation satisfies $|Y_0| \geq k^*(c, \Sigma)$, then the agent optimally stops immediately and chooses good 1 or good 2 depending on whether Y_0 is positive or negative. Whereas if $|Y_0| < k^*(c, \Sigma)$, then the optimal stopping time τ is strictly positive.

To study how the agent's choice accuracy changes over time, we need to also consider the stopping boundaries at later times $t > 0$. For this we let Σ_t be the posterior covariance matrix of θ at time t , given the prior covariance matrix Σ and given the optimal attention strategy. Then the stopping boundary at time t is simply $k^*(c, \Sigma_t)$ (this implicitly uses the fact that starting from the prior Σ_t , the posterior at time s would be Σ_{t+s}).

The next lemma (essentially Theorem 2 in [Fudenberg et al. \(2018\)](#)) characterizes a necessary and sufficient condition for choice accuracy to decrease over time, in terms of these stopping boundaries:

Lemma 12. *Given c and Σ , and suppose $|Y_0| < k^*(c, \Sigma)$ (so that the agent does not immediately stop). Let p_t be the conditional probability that good 1 is better than good 2 when the agent stops at time t and chooses good 1 (and vice versa, by symmetry). Then p_t decreases in t if and only if $\frac{k^*(c, \Sigma_t)}{\sigma_t}$ decreases in t .*

Proof. If the agent stops at time t and chooses good 1, then $Y_t = k^*(c, \Sigma_t)$ (and t is the earliest time this happens). So by definition, the posterior belief of ω is normal with mean $k^*(c, \Sigma_t)$ and standard deviation σ_t . Thus, the conditional probability that $\omega > 0$ (i.e., good 1 is better) is the normal c.d.f. evaluated at $\frac{k^*(c, \Sigma_t)}{\sigma_t}$. This yields the result. \square

Note that σ_t^2 , being the posterior variance of $\theta_1 + \theta_2$, is just the sum of all the entries in the matrix Σ_t . Thus the condition in Lemma 12 is ultimately about how $k^*(c, \Sigma)$ varies with Σ . The next section studies this change in detail.

O.4.3 Effect of Σ on Stopping Boundary

We will say two covariance matrices $\tilde{\Sigma}$ and $\hat{\Sigma}$ induce the same prior uncertainty, if the prior variances $\tilde{\sigma}_0^2 = \tilde{\Sigma}_{11} + 2\tilde{\Sigma}_{12} + \tilde{\Sigma}_{22}$ and $\hat{\sigma}_0^2 = \hat{\Sigma}_{11} + 2\hat{\Sigma}_{12} + \hat{\Sigma}_{22}$ of $\theta_1 + \theta_2$ are equal.

The result below provides a sufficient condition for stopping boundaries under two different prior covariance matrices to be comparable:

Lemma 13. *Let $\tilde{\Sigma}$ and $\hat{\Sigma}$ be two covariance matrices that induce the same prior uncertainty. Suppose further that the following two conditions holds:*

- (a) $\tilde{\Sigma}_{11} - \tilde{\Sigma}_{22} \geq \hat{\Sigma}_{11} - \hat{\Sigma}_{22} \geq 0$;
- (b) $(\tilde{\sigma}_{t_1^*})^2 = \frac{2 \det(\tilde{\Sigma})}{\tilde{\Sigma}_{11} - \tilde{\Sigma}_{12}} \leq \frac{2 \det(\hat{\Sigma})}{\hat{\Sigma}_{11} - \hat{\Sigma}_{12}} = (\hat{\sigma}_{t_1^*})^2$.

Then the posterior variances satisfy $\tilde{\sigma}_t^2 \leq \hat{\sigma}_t^2$ for all $t \geq 0$. Consequently $k^*(c, \tilde{\Sigma}) \geq k^*(c, \hat{\Sigma})$.

Part (a) says that there is greater asymmetry in the agent's uncertainty about the two attributes in the prior covariance matrix $\tilde{\Sigma}$ compared to $\hat{\Sigma}$. Part (b) says that the agent's uncertainty about $\theta_1 + \theta_2$ at the optimal switchpoint \tilde{t}_1^* given prior $\tilde{\Sigma}$ is lower than the agent's uncertainty about $\theta_1 + \theta_2$ at the optimal switchpoint \hat{t}_1^* given prior $\hat{\Sigma}$, i.e., the agent has learned more about $\theta_1 + \theta_2$ in (the endogenous) Stage 1 starting from $\tilde{\Sigma}$. The lemma says that these conditions imply that the agent's uncertainty about $\theta_1 + \theta_2$ is lower at every moment of time starting from prior $\tilde{\Sigma}$, i.e., the agent learns faster.

Proof. Given any prior covariance matrix Σ and resulting path of posterior variances σ_t , we can define for each $v \in [0, \sigma_0^2]$ the hitting time $T(v)$ such that $\sigma_{T(v)}^2 = \sigma_0^2 - v$. $T(v)$ is well-defined because σ_t^2 decreases strictly and continuously in t . This monotonicity also implies that the comparison $\tilde{\sigma}_t^2 \leq \hat{\sigma}_t^2$ for each t is equivalent to $\tilde{T}(v) \leq \hat{T}(v)$ for each $v < \tilde{\sigma}_0^2 = \hat{\sigma}_0^2$. Thus in what follows we study the properties of $T(v)$.

From Lemma 11, it is not difficult to derive the following formula for $T(v)$:

$$T(v) = \begin{cases} \frac{v}{(\Sigma_{11} + \Sigma_{12})^2 - \Sigma_{11}v} & \text{if } v \in [0, v^*]; \\ \frac{4}{\sigma_0^2 - v} - \frac{\Sigma_{11} + \Sigma_{22} - 2\Sigma_{12}}{\det(\Sigma)} & \text{if } v \in [v^*, \sigma_0^2]. \end{cases}$$

Above, the switchpoint v^* is given by

$$v^* = \sigma_0^2 - \sigma_{t_1^*}^2 = (\Sigma_{11} + \Sigma_{22} + 2\Sigma_{12}) - \frac{2 \det(\Sigma)}{\Sigma_{11} - \Sigma_{12}} = \frac{(\Sigma_{11} - \Sigma_{22})(\Sigma_{11} + \Sigma_{12})}{\Sigma_{11} - \Sigma_{12}}.$$

Not surprisingly, at $v = v^*$ either formula for $T(v)$ yields the time $t_1^* = \frac{\Sigma_{11} - \Sigma_{22}}{\det(\Sigma)}$.

We now compute the (right) derivative of $T(v)$:

$$T'(v) = \begin{cases} \frac{1}{\left(\Sigma_{11} + \Sigma_{12} - \frac{\Sigma_{11}}{\Sigma_{11} + \Sigma_{12}}v\right)^2} & \text{if } v \in [0, v^*]; \\ \frac{4}{(\sigma_0^2 - v)^2} & \text{if } v \in [v^*, \sigma_0^2]. \end{cases} \quad (15)$$

This time perhaps more surprisingly, both formulae for $T'(v)$ yield the same value at $v = v^*$. Moreover, it can be checked that $T'(v) \leq \frac{4}{(\sigma_0^2 - v)^2}$ for all v .

We claim that under the two stated conditions on $\tilde{\Sigma}$ and $\hat{\Sigma}$, it holds that $\tilde{T}'(v) \leq \hat{T}'(v)$ for every v . Since $\tilde{T}(0) = \hat{T}(0) = 0$, this would imply the desired comparison $\tilde{T}(v) \leq \hat{T}(v)$. To compare those derivatives, note that given $\tilde{\Sigma}_{11} + \tilde{\Sigma}_{22} + 2\tilde{\Sigma}_{12} = \hat{\Sigma}_{11} + \hat{\Sigma}_{22} + 2\hat{\Sigma}_{12}$, the assumption $\tilde{\Sigma}_{11} - \tilde{\Sigma}_{22} \geq \hat{\Sigma}_{11} - \hat{\Sigma}_{22}$ is equivalent to $\tilde{\Sigma}_{11} + \tilde{\Sigma}_{12} \geq \hat{\Sigma}_{11} + \hat{\Sigma}_{12}$. Thus $\tilde{T}'(0) \leq \hat{T}'(0)$ holds. Moreover, the second assumption in the lemma translates into $\tilde{v}^* \geq \hat{v}^*$. Below we show these are sufficient to imply $\tilde{T}'(v) \leq \hat{T}'(v)$.

Indeed, for $v \geq \widehat{v}^*$, we deduce from (15) that

$$\widetilde{T}'(v) \leq \frac{4}{(\sigma_0^2 - v)^2} = \widehat{T}'(v).$$

On the other hand, for $v \leq \widehat{v}^*$, $T'(v)$ is given by the first term in (15) for both $\widetilde{\Sigma}$ and $\widehat{\Sigma}$. Thus the comparison between $\widetilde{T}'(v)$ and $\widehat{T}'(v)$ reduces to a comparison between two linear functions of v : We want to show $\Sigma_{11} + \Sigma_{12} - \frac{\Sigma_{11}}{\Sigma_{11} + \Sigma_{12}}v$ is larger when $\Sigma = \widetilde{\Sigma}$ than when $\Sigma = \widehat{\Sigma}$. We already know this holds at $v = 0$ and $v = \widetilde{v}^*$, so by linearity, it also holds at any v in between, completing the proof.

Hence we have shown that $\widetilde{\sigma}_t^2 \leq \widehat{\sigma}_t^2$ for every t . It remains to show that the comparison of posterior variances implies the comparison of stopping boundaries. For this we observe that lower posterior variances under $\widetilde{\Sigma}$ imply that the value function $U(y, c, \widetilde{\Sigma})$ is weakly larger than $U(y, c, \widehat{\Sigma})$ for any cost c and any prior expectation y . This follows from the same time-change argument as in the proof of Lemma 5, and the idea is simply that any stopping time under prior $\widehat{\Sigma}$ can be replicated under prior $\widetilde{\Sigma}$ at an earlier stopping time. Thus

$$0 = U(-k^*(c, \widetilde{\Sigma}), c, \widetilde{\Sigma}) \geq U(-k^*(c, \widetilde{\Sigma}), c, \widehat{\Sigma}).$$

In fact we have equality since U is non-negative. Hence $k^*(c, \widehat{\Sigma})$, being the smallest x such that $U(-x, c, \widehat{\Sigma}) = 0$, must be smaller than $k^*(c, \widetilde{\Sigma})$. \square

Two useful corollaries follow from Lemma 13 (the proofs are immediate and thus omitted):

Lemma 14 (Effect of correlation). *Let $\widetilde{\Sigma}$ and $\widehat{\Sigma}$ be two covariance matrices that induce the same prior uncertainty. Suppose further that $\widetilde{\Sigma}_{11} = \widetilde{\Sigma}_{22}$ and $\widehat{\Sigma}_{11} = \widehat{\Sigma}_{22}$ (symmetric priors). Then $k^*(c, \widetilde{\Sigma}) = k^*(c, \widehat{\Sigma})$.*

Lemma 15 (Effect of asymmetry). *Let $\widetilde{\Sigma}$ and $\widehat{\Sigma}$ be two covariance matrices that induce the same prior uncertainty. Suppose further that $\widetilde{\Sigma}_{11} > \widetilde{\Sigma}_{22}$ while $\widehat{\Sigma}_{11} = \widehat{\Sigma}_{22}$ (asymmetric versus symmetric). Then $k^*(c, \widetilde{\Sigma}) \geq k^*(c, \widehat{\Sigma})$.*

While the above results hold fixed the prior variance σ_0^2 , we also need a result that considers a change in overall prior uncertainty, and characterizes its effect on the stopping boundary.

Lemma 16 (Effect of scaling Σ). *For any c, Σ and any $\lambda \in (0, 1)$, it holds that $\lambda k^*(c, \Sigma) \geq k^*(c, \lambda^2 \Sigma)$.*

Proof. We will show that for any $\lambda > 0$,

$$k^*(c, \lambda^2 \Sigma) = \lambda k^*(c \lambda^{-3}, \Sigma). \tag{16}$$

This implies the lemma because the cost $c\lambda^{-3}$ is higher than c whenever $\lambda < 1$, which decreases the value function and thus also decreases the stopping boundary, resulting in $k^*(c\lambda^{-3}, \Sigma) \leq k^*(c, \Sigma)$.

We note that the identity (16) is a direct generalization of Equation (A6) in Fudenberg et al. (2018). Nonetheless, we provide a proof below for completeness. The key insight is that we can identify the belief processes under prior Σ and under $\lambda^2\Sigma$ via a scaling of time and space. Specifically, let Y_t denote the belief process under Σ as given by (13), with $\sigma_t^2(\Sigma)$ denoting the posterior variance at time t (under the optimal path). Similarly define Z_t for the belief process starting from the prior $\lambda\Sigma$, with $\sigma_t^2(\lambda^2\Sigma)$ denoting the posterior variance at time t . From Lemma 11, we have the relation

$$\sigma_t^2(\lambda^2\Sigma) = \lambda^2 \cdot \sigma_{\lambda^2 t}^2(\Sigma).$$

Assuming $Z_0 = \lambda Y_0$, then the process Z_t has the same distribution as the process $\lambda \cdot Y_{\lambda^2 t}$. Intuitively, receiving standard normal signals about $\lambda\omega$ for one unit of time is equivalent to receiving standard normal signals about ω for λ^2 units of time.

Now that we identify Z_t with $\lambda \cdot Y_{\lambda^2 t}$, we can rewrite the stopping problem with respect to Z_t in terms of the Y process instead. Specifically, the agent's problem is

$$\max_{\tau} \mathbb{E}[\max\{\lambda Y_{\lambda^2 \tau}, 0\} - c\tau] = \lambda \max_{\tau'} \mathbb{E}[\max\{Y_{\tau'}, 0\} - \frac{c}{\lambda^3} \tau'],$$

with $\tau' = \lambda^2 \tau$. Thus it is as if the agent chooses an optimal stopping time with respect to Y_t , but with transformed marginal cost $\frac{c}{\lambda^3}$. The agent should stop at time 0 in this problem if and only if $|Y_0| \geq k^*(\frac{c}{\lambda^3}, \Sigma)$, which is equivalent to $|Z_0| \geq \lambda k^*(\frac{c}{\lambda^3}, \Sigma)$. Thus (16) holds. \square

O.4.4 Main Proof of Proposition 1

Given Lemma 12, we just need to show that for any times $t < t'$, $\frac{k^*(c, \Sigma_t)}{\sigma_t} \geq \frac{k^*(c, \Sigma_{t'})}{\sigma_{t'}}$. Define $\lambda \in (0, 1)$ such that $\sigma_{t'} = \lambda \sigma_t$. Then the inequality becomes $\lambda k^*(c, \Sigma_t) \geq k^*(c, \Sigma_{t'})$. From Lemma 16 we have $\lambda k^*(c, \Sigma_t) \geq k^*(c, \lambda^2 \Sigma_t)$, so it is sufficient to show

$$k^*(c, \lambda^2 \Sigma_t) \geq k^*(c, \Sigma_{t'}). \quad (\text{to be shown})$$

Note that, by the definition of λ , the matrices $\lambda^2 \Sigma_t$ and $\Sigma_{t'}$ induce the same uncertainty about $\theta_1 + \theta_2$. There are three cases to consider:

Case 1: $t_1^* \leq t < t'$. In this case Σ_t and $\Sigma_{t'}$ are posterior covariance matrices in Stage 2, so they induce symmetric uncertainty about θ_1 and θ_2 . Lemma 14 thus applies to the symmetric priors $\lambda^2 \Sigma_t$ and $\Sigma_{t'}$, and yields $k^*(c, \lambda^2 \Sigma_t) = k^*(c, \Sigma_{t'})$. Intuitively, the belief process in Stage 2 is the same as in Fudenberg et al. (2018) since correlation does not matter for symmetric priors. Thus the result in this case follows from the result in that paper.

Case 2: $t < t_1^* \leq t'$. Here Lemma 14 no longer applies. We instead apply Lemma 15 to deduce $k^*(c, \lambda^2 \Sigma_t) \geq k^*(c, \Sigma_{t'})$, since $\lambda^2 \Sigma_t$ is an asymmetric prior while $\Sigma_{t'}$ is a symmetric prior. This proves the result, and as discussed, the intuition is that asymmetry increases the stopping boundary relative to symmetric priors.

Case 3: $t < t' < t_1^*$. To prove the key comparison $k^*(c, \lambda^2 \Sigma_t) \geq k^*(c, \Sigma_{t'})$, we now need to invoke the more general Lemma 13 since $\lambda^2 \Sigma_t$ and $\Sigma_{t'}$ are both asymmetric. Thus, we have to check that $\tilde{\Sigma} = \lambda^2 \Sigma_t$ and $\hat{\Sigma} = \Sigma_{t'}$ satisfy the two conditions stated in Lemma 13.

To do this, we let $\Sigma_t = \begin{pmatrix} p & r \\ r & q \end{pmatrix}$ with $p > q > 0$ and $r^2 < pq$. The posterior covariance matrix $\Sigma_{t'}$ at the later time t' can be calculated from the ‘‘prior’’ covariance matrix Σ_t , after focusing on θ_1 for $t' - t$ units of time. Thus

$$\Sigma_{t'} = \left(\Sigma_t^{-1} + \begin{pmatrix} t' - t & 0 \\ 0 & 0 \end{pmatrix} \right)^{-1} = \left(\frac{1}{pq - r^2} \begin{pmatrix} q + (t' - t)(pq - r^2) & -r \\ -r & p \end{pmatrix} \right)^{-1}.$$

Let $q' = q + (t' - t)(pq - r^2)$ with $q < q' < p$ (the latter inequality holds because $t' < t_1^*$). Then the above simplifies to

$$\Sigma_{t'} = \frac{pq - r^2}{pq' - r^2} \begin{pmatrix} p & r \\ r & q' \end{pmatrix}$$

The scaling factor λ is thus given by

$$\lambda^2(p + q + 2r) = \frac{pq - r^2}{pq' - r^2}(p + q' + 2r). \quad (17)$$

We now check the first condition in Lemma 13, which for $\tilde{\Sigma} = \lambda^2 \Sigma_t$ and $\hat{\Sigma} = \Sigma_{t'}$ becomes $\lambda^2(p - q) \geq \frac{pq - r^2}{pq' - r^2}(p - q')$. Using (17) to eliminate λ , it suffices to show that

$$\frac{p - q}{p + q + 2r} \geq \frac{p - q'}{p + q' + 2r}.$$

This inequality holds simply because $q' > q$.

We then turn to the second condition in Lemma 12, which in the current setting becomes $2\lambda^2 \left(\frac{pq - r^2}{p - r} \right) \leq 2 \left(\frac{pq - r^2}{pq' - r^2} \right) \cdot \left(\frac{pq' - r^2}{p - r} \right)$. This simplifies to $\lambda^2 \leq 1$, which clearly holds. Intuitively, since Σ_t and $\Sigma_{t'}$ are both posterior beliefs following the prior Σ , they ‘‘become symmetric’’ at the same posterior belief $\Sigma_{t_1^*}$. Thus the second condition in Lemma 13 holds with equality when we consider Σ_t versus $\Sigma_{t'}$. It follows that when comparing $\lambda^2 \Sigma_t$ and $\Sigma_{t'}$, the former prior belief leads to lower uncertainty when entering Stage 2.

Hence Lemma 13 applies, and we again have

$$\lambda k^*(c, \Sigma_t) \geq k^*(c, \lambda^2 \Sigma_t) \geq k^*(c, \Sigma_{t'}).$$

This proves $\frac{k^*(c, \Sigma_t)}{\sigma_t} \geq \frac{k^*(c, \Sigma_{t'})}{\sigma_{t'}}$ and the proposition.

O.4.5 Generalization to Unequal Learning Speeds

In this section we show that the conclusion of Proposition 1 further generalizes to situations where the information about the two unknown payoffs arrives with different levels of precision. Formally, fix $\alpha_1, \alpha_2 > 0$ and suppose a unit of time devoted to learning about the payoff v_i produces the signal $v_i + \mathcal{N}(0, \alpha_i^2)$. Any prior covariance matrix over (v_1, v_2) can be written as $\begin{pmatrix} \alpha_1^2 \Sigma_{11} & -\alpha_1 \alpha_2 \Sigma_{12} \\ -\alpha_1 \alpha_2 \Sigma_{21} & \alpha_2^2 \Sigma_{22} \end{pmatrix}$ where $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$ is also a positive definite matrix. This formulation of the problem will make the results below easier to state.

To transform this problem into our main model, let $\theta_1 = v_1/\alpha_1$ and $\theta_2 = -v_2/\alpha_2$, so that each unit of time devoted to θ_i produces a standard normal signal about it. Moreover, the prior covariance matrix over (θ_1, θ_2) is simply Σ , and the payoff-relevant state is $\omega = v_1 - v_2 = \alpha_1 \theta_1 + \alpha_2 \theta_2$. Throughout this section, we assume Σ and α satisfy Assumption 3, so that we can apply Theorem 1 to characterize optimal attention allocation.

Corollary 4. *Suppose Assumption 3 holds and $\alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{12} \geq \alpha_1 \Sigma_{12} + \alpha_2 \Sigma_{22}$. The agent's optimal information acquisition strategy $(\beta_1(t), \beta_2(t))$ in this generalized binary choice problem consists of two stages:*

- **Stage 1:** At all times

$$t < t_1^* = \frac{\alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{12} - \alpha_1 \Sigma_{21} - \alpha_2 \Sigma_{22}}{\alpha_2 \det(\Sigma)},$$

the agent optimally allocates all attention to θ_1 .

- **Stage 2:** At times $t \geq t_1^*$, the agent optimally uses the constant mixture $\left(\frac{\alpha_1}{\alpha_1 + \alpha_2}, \frac{\alpha_2}{\alpha_1 + \alpha_2}\right)$.

From this we can compute the agent's posterior variance of $\alpha_1 \theta_1 + \alpha_2 \theta_2$ at each time t . Generalizing Lemma 11, we have

$$\sigma_t^2 = \begin{cases} \frac{\alpha_1^2 \Sigma_{11} + \alpha_2^2 \Sigma_{22} + 2\alpha_1 \alpha_2 \Sigma_{12} + \alpha_2^2 \det(\Sigma)t}{1 + \Sigma_{11}t} & \text{if } t \leq t_1^*; \\ \frac{(\alpha_1 + \alpha_2)^2 \det(\Sigma)}{\Sigma_{11} + \Sigma_{22} - 2\Sigma_{12} + \det(\Sigma)t} & \text{if } t \geq t_1^*. \end{cases}$$

In particular, $\sigma_{t_1^*}^2 = \frac{(\alpha_1 + \alpha_2)\alpha_2 \det(\Sigma)}{\Sigma_{11} - \Sigma_{12}}$. We omit the detailed calculations.

Next, note that Lemma 12 holds without change, so we just need to study how the stopping boundary $k^*(c, \Sigma)$ varies with Σ in this more general setting. Naturally, we say that two prior covariance matrices $\tilde{\Sigma}$ and $\hat{\Sigma}$ induce the same prior uncertainty if the prior variances $\alpha_1^2 \tilde{\Sigma}_{11} + \alpha_2^2 \tilde{\Sigma}_{22} + 2\alpha_1 \alpha_2 \tilde{\Sigma}_{12}$ and $\alpha_1^2 \hat{\Sigma}_{11} + \alpha_2^2 \hat{\Sigma}_{22} + 2\alpha_1 \alpha_2 \hat{\Sigma}_{12}$ are equal. The key Lemma 13 above is then generalized as follows:

Lemma 17. *Let $\tilde{\Sigma}$ and $\hat{\Sigma}$ be two covariance matrices that induce the same prior uncertainty. Suppose further that the following two conditions holds:*

- (a) $\alpha_1 \tilde{\Sigma}_{11} + \alpha_2 \tilde{\Sigma}_{12} - \alpha_1 \tilde{\Sigma}_{21} - \alpha_2 \tilde{\Sigma}_{22} \geq \alpha_1 \hat{\Sigma}_{11} + \alpha_2 \hat{\Sigma}_{12} - \alpha_1 \hat{\Sigma}_{21} - \alpha_2 \hat{\Sigma}_{22} \geq 0;$
- (b) $(\tilde{\sigma}_{t_1^*}^2)^2 = \frac{(\alpha_1 + \alpha_2) \alpha_2 \det(\tilde{\Sigma})}{\tilde{\Sigma}_{11} - \tilde{\Sigma}_{12}} \leq \frac{(\alpha_1 + \alpha_2) \alpha_2 \det(\hat{\Sigma})}{\hat{\Sigma}_{11} - \hat{\Sigma}_{12}} = (\hat{\sigma}_{t_1^*}^2)^2.$

Then the posterior variances satisfy $\tilde{\sigma}_t^2 \leq \hat{\sigma}_t^2$ for all $t \geq 0$. Consequently $k^(c, \tilde{\Sigma}) \geq k^*(c, \hat{\Sigma})$.*

That is, with potentially unequal payoff weights α_1 and α_2 , prior ‘‘asymmetry’’ is not simply measured by the difference between Σ_{11} and Σ_{22} . Rather, it is measured by the difference in initial marginal values $cov_1 = \alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{12}$ and $cov_2 = \alpha_1 \Sigma_{21} + \alpha_2 \Sigma_{22}$. Condition (a) thus requires this asymmetry to be larger under $\tilde{\Sigma}$ than under $\hat{\Sigma}$. This turns out to be equivalent to $\tilde{T}'(0) \leq \hat{T}'(0)$, where the hitting time T is the same as defined in the proof of Lemma 13.²⁶ Together with condition (b), this implies $\tilde{T}'(v) \leq \hat{T}'(v)$ for all v and thus $\tilde{T}(v) \leq \hat{T}(v)$, just as in that proof.

Using the same notion of asymmetry, we obtain direct generalizations of Lemma 14 and Lemma 15 as well. Finally, Lemma 16 continues to hold since its proof does not depend on payoff weights. These lemmata allow us to replicate the proof of Proposition 1 with only minor modifications. Thus, to summarize, we have the following result:

Proposition 9. *Consider the binary choice problem with general signal variances α_1^2 and α_2^2 . Denote the agent’s prior covariance matrix over the payoffs (v_1, v_2) as $\begin{pmatrix} \alpha_1^2 \Sigma_{11} & -\alpha_1 \alpha_2 \Sigma_{12} \\ -\alpha_1 \alpha_2 \Sigma_{21} & \alpha_2^2 \Sigma_{22} \end{pmatrix}$. Then whenever Σ and α satisfy Assumption 3, the agent’s choice accuracy in this problem is (weakly) higher at earlier stopping times.*

We reiterate that Assumption 3 is guaranteed if $\alpha_1 = \alpha_2$ (as we assumed previously), or if $\Sigma_{12} \geq 0$ (i.e., the unknown payoffs v_1 and v_2 are negatively correlated).

O.5 Proof of Proposition 3

For $t \leq T$, $\hat{n}_i(t) = 0$ for every $i > 1$, so the result $\hat{n}_i(t) \leq n_i(t)$ trivially holds. Below we consider a fixed time $t > T$. We can further assume $n_1(t) < T$, because otherwise the proof of Proposition 2

²⁶From the formulae for σ_t^2 we can compute $T(v)$ and $T'(v)$. Writing $v^* = \sigma_0^2 - \sigma_{t_1^*}^2 = \frac{(cov_1 - cov_2)cov_1}{\Sigma_{11} - \Sigma_{12}}$, we have

$$T'(v) = \begin{cases} \frac{1}{(cov_1 - \frac{\Sigma_{11}}{cov_1} v)^2} & \text{if } v \in [0, v^*]; \\ \frac{(\alpha_1 + \alpha_2)^2}{(\sigma_0^2 - v)^2} & \text{if } v \in [v^*, \sigma_0^2]. \end{cases}$$

(in the main text) shows that $\hat{n}(t)$ coincides with $n(t)$. Given this assumption, we claim that $\hat{n}_1(t)$ is exactly equal to T . Indeed, let $\bar{t} \geq t$ be the first time at which $n_1(\bar{t}) = T$; such a time exists by monotonicity and continuity of $n_1(\cdot)$. Then $\hat{n}(\bar{t})$ coincides with $n(\bar{t})$, which in particular implies $\hat{n}_1(\bar{t}) = n_1(\bar{t}) = T$. Monotonicity of $\hat{n}_1(\cdot)$ then implies $\hat{n}_1(t) \leq T$. But by assumption $\hat{n}_1(t) \geq T$, so equality must hold.

Next, we connect the two vectors $n(t)$ and $\hat{n}(t)$ by a continuous path. For each $x \in [n_1(t), T]$, we define q^x as the cumulative attention vector (at time t) resulting from a hypothetical attention manipulation that forces the agent to observe source 1 for x units of time. That is,

$$q^x = (q_1^x, \dots, q_K^x) = \underset{q_1, \dots, q_K \geq 0: \sum_i q_i = t \text{ and } q_1 \geq x}{\operatorname{argmin}} V(q).$$

Clearly, $n(t) = q^{n_1(t)}$ and $\hat{n}(t) = q^T$. By the same argument as in the previous paragraph, $q_1^x = x$ holds for x in this range. So in defining q^x we can replace the constraint $q_1^x \geq x$ with equality.

To prove the proposition, it suffices to show that as x decreases from T to $n_1(t)$, q_i^x weakly increases for each $i > 1$. Similar to our proof of Theorem 2, the proof strategy here will be to use the Hessian matrix of V to compute the derivative of the vector q^x with respect to x . For this we fix $x > n_1(t)$, and assume for now that q_i^x is strictly positive for each $i > 1$. Then the first-order condition for the constrained optimality of q^x yields $\partial_2 V(q^x) = \dots = \partial_K V(q^x)$. If the vector q^x is left-differentiable at x , then for any y slightly smaller than x , q^y must also satisfy the above equal marginal value condition (for every source $i > 1$).

Thus the left derivative of q^x is a vector $u \in \mathbb{R}^K$ that satisfies $u_1 = 1$, $u_1 + \dots + u_K = 0$ and

$$\operatorname{Hess}_V(q^x) \cdot u = \lambda(c, 1, \dots, 1)', \quad (18)$$

for some $\lambda, c \in \mathbb{R}$ that will be determined later.

Under the differentiability assumption, we can solve for u as follows. Note from Lemma 3 that $\partial_{ij} V = 2\gamma_i \gamma_j [(\Sigma^{-1} + Q)^{-1}]_{ij}$, where we save notation by writing $Q = \operatorname{diag}(q^x)$ from now on. Then, we have the matrix identity

$$\operatorname{Hess}_V(q^x) = 2 \operatorname{diag}(\gamma) \cdot (\Sigma^{-1} + Q)^{-1} \cdot \operatorname{diag}(\gamma), \quad (19)$$

where γ is as usual the vector $(\Sigma^{-1} + Q)^{-1} \cdot \alpha$. As shown in the proof of Proposition 7, γ has strictly positive coordinates. Recalling $\partial_i V = -\gamma_i^2$, we thus have $\gamma_2 = \dots = \gamma_K > 0$. γ_1 cannot be larger, since then q_1^x should be larger than x to minimize V . γ_1 cannot be equal to the other sources either, since in that case the vector q^x would satisfy the first-order condition for the unconstrained variance minimization problem, leading to $q^x = n(t)$ and $x = n_1(t)$. So $\gamma_2 = \dots = \gamma_K > \gamma_1 > 0$, which will be useful below.

Now, using (18) and (19), we have

$$\begin{aligned}
u &= \lambda \cdot Hess^{-1} \cdot (c, 1, \dots, 1)' = 2\lambda \cdot \text{diag}(1/\gamma) \cdot (\Sigma^{-1} + Q) \cdot \text{diag}(1/\gamma) \cdot (c, 1, \dots, 1)' \\
&= 2\lambda \cdot \text{diag}(1/\gamma) \cdot (\Sigma^{-1} + Q) \cdot \left(\frac{c}{\gamma_1}, \frac{1}{\gamma_2}, \dots, \frac{1}{\gamma_2} \right)' \\
&= \frac{2\lambda}{\gamma_2^2} \cdot \text{diag}(1/\gamma) \cdot (\Sigma^{-1} + Q) \cdot \left(\frac{c\gamma_2^2}{\gamma_1}, \gamma_2, \dots, \gamma_2 \right)'.
\end{aligned}$$

If we rewrite $\frac{c\gamma_2^2}{\gamma_1}$ as $\gamma_1 - b$ for some $b \in \mathbb{R}$, then the vector $\left(\frac{c\gamma_2^2}{\gamma_1}, \gamma_2, \dots, \gamma_2 \right)'$ differs from γ only in that the first coordinate is smaller by b . Using $(\Sigma^{-1} + Q)\gamma = \alpha$, we thus obtain

$$(\Sigma^{-1} + Q) \cdot \left(\frac{c\gamma_2^2}{\gamma_1}, \gamma_2, \dots, \gamma_2 \right)' = (\alpha_1 - b[\Sigma^{-1} + Q]_{11}, \alpha_2 - b[\Sigma^{-1} + Q]_{21}, \dots, \alpha_K - b[\Sigma^{-1} + Q]_{K1})'$$

By the symmetry of $\Sigma^{-1} + Q$, it follows that

$$u_i = \frac{2\lambda}{\gamma_2^2} \cdot \frac{\alpha_i - b[\Sigma^{-1} + Q]_{1i}}{\gamma_i}. \tag{20}$$

Recall $u_1 = 1$, which reflects $q_1^y = y$ for every y . Thus $\lambda \neq 0$. Moreover, recall $\sum_{i=1}^K u_i = 0$, a consequence of the fact that $\sum_{i=1}^K q_i^y = t$ for every y . Thus, we obtain

$$b \sum_{i=1}^K \frac{[\Sigma^{-1} + Q]_{1i}}{\gamma_i} = \sum_{i=1}^K \frac{\alpha_i}{\gamma_i}. \tag{21}$$

The RHS is clearly positive. On the LHS, using $\gamma_2 = \dots = \gamma_K > \gamma_1 > 0$ and $[\Sigma^{-1} + Q]_{11} > 0$,

$$\begin{aligned}
\sum_{i=1}^K \frac{[\Sigma^{-1} + Q]_{1i}}{\gamma_i} &= \frac{[\Sigma^{-1} + Q]_{11}}{\gamma_1} + \sum_{i>1} \frac{[\Sigma^{-1} + Q]_{1i}}{\gamma_i} \\
&> \frac{[\Sigma^{-1} + Q]_{11} \cdot \gamma_1}{\gamma_2^2} + \sum_{i>1} \frac{[\Sigma^{-1} + Q]_{1i} \cdot \gamma_i}{\gamma_2^2} \\
&= \frac{1}{\gamma_2^2} \sum_{i=1}^K [\Sigma^{-1} + Q]_{1i} \cdot \gamma_i \\
&= \frac{\alpha_1}{\gamma_2^2} > 0,
\end{aligned}$$

where the last equality uses $\sum_{i=1}^K (\Sigma^{-1} + Q)_{1i} \cdot \gamma_i = \alpha_1 > 0$ as a consequence of $(\Sigma^{-1} + Q)\gamma = \alpha$.

Thus, (21) implies the crucial inequality $b > 0$. It follows that for this b and any $i > 1$, $\frac{\alpha_i - b[\Sigma^{-1} + Q]_{1i}}{\gamma_i}$ is strictly positive (by assumption $[\Sigma^{-1} + Q]_{1i} \leq 0$). From $\sum_{i=1}^K u_i = 0$, we further know that $\frac{\alpha_1 - b[\Sigma^{-1} + Q]_{11}}{\gamma_1}$ is strictly negative. We can then use (20) to determine the unique value of $\lambda < 0$ that makes $u_1 = 1$. Since $\lambda < 0$ and $b > 0$, (20) yields $u_i < 0$ for every $i > 1$. Hence, starting from q^x , any local decrease in the amount of attention towards source 1 increases the optimal amount of attention towards every other source. This is what we desire to show.

To complete the proof, we now dispense with two previous assumptions in the analysis. First, we have assumed $q_i^x > 0$ for each $i > 1$. In general, relabeling the attributes if necessary, we can assume that among the sources $2, \dots, K$, the ones with maximal marginal value (i.e., those with maximal γ_i) are sources $2 \sim k$ for some k . Then the first-order condition for the constrained optimality of q_i^x requires $q_i^x = 0$ for $i > k$. Moreover, by the same argument as above, γ_1 is strictly smaller. By continuity, sources $2 \sim k$ also maximize the marginal value at any q_i^y where y is slightly smaller than x . So $q_i^y = 0$ for $i > k$ also holds. Thus, locally we can reduce the problem with K sources to a smaller problem with only the first k sources, similar to the proof of Theorem 2.

In this smaller problem, the payoff weight vector $\tilde{\alpha}$ is given by (11), and it is strictly positive as shown in the proof of Theorem 2. The prior covariance matrix becomes Σ_{TL} , which is the $k \times k$ top-left principal sub-matrix of Σ . Since Σ^{-1} is an M -matrix, so is $(\Sigma_{TL})^{-1}$.²⁷ The vector $\tilde{\gamma} = [(\Sigma_{TL})^{-1} + \tilde{Q}]^{-1} \cdot \tilde{\alpha} \in \mathbb{R}^k$ is the first k coordinates of γ , so we have $\tilde{\gamma}_2 = \dots = \tilde{\gamma}_k > \tilde{\gamma}_1 > 0$. Hence, the above procedure for finding u directly applies to this smaller problem, and yields a vector $\tilde{u} \in \mathbb{R}^k$ with $\tilde{u}_1 = 1$, $\tilde{u}_i < 0$ for $1 < i \leq k$ and $\tilde{u}_1 + \dots + \tilde{u}_k = 0$. In the original problem, the left-derivative of q^x is thus the vector $(\tilde{u}_1, \dots, \tilde{u}_k, 0, \dots, 0)'$. Once again, locally decreasing attention towards source 1 weakly increases attention towards every other source. To be fully rigorous, we also note that by the Maximum Theorem, q^x varies continuously with x . So the preceding form of local monotonicity implies global monotonicity (without continuity, q_i^x may only be piece-wise monotone).

Our second simplifying assumption is left-differentiability, which we now argue is without loss. Let x be the infimum of those numbers $\hat{x} \in [n_1(t), T]$ such that q^y is left-differentiable in y for $y > \hat{x}$. By properties of the infimum, x is in fact the *smallest* number $\geq n_1(t)$ such that q^y is left-differentiable for $y > x$. Now suppose $x > n_1(t)$, and we will deduce a contradiction. Specifically, let sources $2 \sim k$ have the maximal marginal value at q^x . We can use the above procedure to find the vectors $\tilde{u} \in \mathbb{R}^k$ and $(\tilde{u}_1, \dots, \tilde{u}_k, 0, \dots, 0)' \in \mathbb{R}^K$, such that if q^x is perturbed slightly in the direction $-u$, the resulting attention vector maintains the equal marginal value property across sources $2 \sim k$. Formally, for any attention vector q with total attention t such that sources $2 \sim k$ have the maximal marginal value, we can solve for $b, \lambda \in \mathbb{R}$ and $\tilde{u} \in \mathbb{R}^k$ from (20) and (21) (with $\tilde{\alpha}$, Σ_{TL} and $\text{diag}(q)$ replacing α , Σ and Q in those equations). It is easy to see that these solutions vary continuously with q , so we can write $u = f(q)$ for a continuous function f . This allows us to define the following system of ordinary differential equations with a right boundary condition

²⁷Rewriting $(\Sigma_{TL})^{-1}$ as the Schur complement of Σ^{-1} with respect to its bottom-right block, the result follows from the fact that M -matrices are closed under Schur complements. This fact can be proved by the same induction argument as in Carlson and Markham (1979).

(where the derivative at x is interpreted as the left-derivative).

$$q'(y) = f(q(y)) \text{ for every } y \in (x - \epsilon, x], \text{ and } q(x) = q^x.$$

By Peano's Existence Theorem, this system of ODE admits a solution when ϵ is sufficiently small. Note that by construction, for any y in the interval $(x - \epsilon, x]$, we have $q_1(y) = y$, $\sum_{i=1}^K q_i(y) = t$ and $q_i(y) = 0$ for $i > k$. Moreover, at the vector $q(y)$ sources $2 \sim k$ have equal marginal values, which are maximal if ϵ is sufficiently small. Hence $q(y)$ satisfies the Kuhn-Tucker conditions for minimizing $V(q)$ subject to $q_1 \geq y$, $q_i \geq 0$ and $\sum_{i=1}^K q_i = t$. Since V is convex, these conditions are sufficient, so that $q(y)$ coincides with q^y . But then we see that q^y is left-differentiable for any $y > x - \epsilon$, contradicting the definition of x .

This completes the entire proof of Proposition 3.

O.6 Proof of Proposition 4

O.6.1 Reduction to the Main Model

We first show it is without loss to consider $\sigma_b = 1$. Suppose the result holds for $\sigma_b = 1$, then for a general σ_b , we can write $b = \sigma_b \cdot \hat{b}$ and $\omega = \sigma_b \cdot \hat{\omega}$ for some random variables \hat{b} and $\hat{\omega}$, where \hat{b} has unit variance. Note that the two signals $\omega + \phi_1 b + \mathcal{N}(0, \zeta_1^2)$ and $\omega - \phi_2 b + \mathcal{N}(0, \zeta_2^2)$ are informationally equivalent to $\hat{\omega} + \phi_1 \hat{b} + \mathcal{N}\left(0, \left(\frac{\zeta_1}{\sigma_b}\right)^2\right)$ and $\hat{\omega} - \phi_2 \hat{b} + \mathcal{N}\left(0, \left(\frac{\zeta_2}{\sigma_b}\right)^2\right)$. We have thus transformed the general problem into one with payoff-relevant state $\hat{\omega}$ and unknown benefit \hat{b} , where \hat{b} has unit variance. The result for this case then pins down equilibrium choices of ϕ_i and $\frac{\zeta_i}{\sigma_b}$, which then yield the equilibrium in the general case.

Hence, for the rest of the proof, we assume $\sigma_b = 1$. Define $\theta_1 = \frac{1}{\zeta_1}(\omega + \phi_1 b)$ and $\theta_2 = \frac{1}{\zeta_2}(\omega - \phi_2 b)$. Observe that $\omega + \phi_1 b + \mathcal{N}(0, \zeta_1^2)$ is informationally equivalent to $\frac{1}{\zeta_1}(\omega + \phi_1 b) + \mathcal{N}(0, 1)$. Thus a unit of time spent on source i produces a standard normal signal about the corresponding θ_i , which returns our main model. The prior covariance matrix for (θ_1, θ_2) is

$$\Sigma = \begin{pmatrix} \frac{\sigma_\omega^2 + \phi_1^2}{\zeta_1^2} & \frac{\sigma_\omega^2 - \phi_1 \phi_2}{\zeta_1 \zeta_2} \\ \frac{\sigma_\omega^2 - \phi_1 \phi_2}{\zeta_1 \zeta_2} & \frac{\sigma_\omega^2 + \phi_2^2}{\zeta_2^2} \end{pmatrix}$$

and the payoff-relevant state can be written as $\omega = \alpha_1 \theta_1 + \alpha_2 \theta_2$ with payoff weights $\alpha_1 = \zeta_1 \cdot \frac{\phi_2}{\phi_1 + \phi_2}$ and $\alpha_2 = \zeta_2 \cdot \frac{\phi_1}{\phi_1 + \phi_2}$.

Below we derive the reader's optimal attention allocation using our main results. Throughout we assume $\zeta_1 \leq \zeta_2$, so that source 1 is more precise. Under this assumption, we have $cov_1 = \sigma_\omega^2 / \zeta_1 \geq \sigma_\omega^2 / \zeta_2 = cov_2$, which implies that source 1 is attended to in the first stage. Moreover, by

Theorem 1 (which applies since cov_1 and cov_2 are both positive), the length of the first stage is

$$t_1^* = \frac{cov_1 - cov_2}{\alpha_2 \det(\Sigma)} = \frac{\zeta_1(\zeta_2 - \zeta_1)}{\phi_1(\phi_1 + \phi_2)}, \quad (22)$$

where we used $\det(\Sigma) = \sigma_\omega^2 \left(\frac{\phi_1 + \phi_2}{\zeta_1 \zeta_2} \right)^2$.

Thus, the reader optimally attends only to source 1 until time t_1^* , and afterwards gives $\frac{\alpha_1}{\alpha_1 + \alpha_2}$ fraction of his attention to source 1. We can then write the two sources' payoffs as follows:

$$\begin{aligned} U_1 &= \int_0^{t_1^*} r e^{-rt} dt + \int_{t_1^*}^{\infty} r e^{-rt} \left(\frac{\zeta_1 \phi_2}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right) dt - \lambda(1 - \phi_1)^2 \\ &= 1 - e^{-rt_1^*} \left(\frac{\zeta_2 \phi_1}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right) - \lambda(1 - \phi_1)^2; \\ U_2 &= e^{-rt_1^*} \left(\frac{\zeta_2 \phi_1}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right) - \lambda(1 - \phi_2)^2. \end{aligned} \quad (23)$$

These payoffs define a stage game between the two sources, and our goal is to characterize its equilibrium. Our strategy below is to use first-order conditions to pin down what an equilibrium must be. We will then verify the equilibrium by checking all possible deviations.

O.6.2 Solving for Equilibrium Precisions ζ_1^* , ζ_2^*

We show here that the precision choices ζ_1 , ζ_2 must be equal in any equilibrium. Suppose not, then small changes in ζ_1 and ζ_2 do not affect our standing assumption that $\zeta_1 \leq \zeta_2$. Thus we can take the first-order conditions for ζ_1 and ζ_2 . Observe that $\frac{\partial t_1^*}{\partial \zeta_1} = \frac{\zeta_2 - 2\zeta_1}{\phi_1(\phi_1 + \phi_2)}$. So we can compute that

$$\frac{\partial U_1}{\partial \zeta_1} = e^{-rt_1^*} \left(\frac{\zeta_2 \phi_1}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right) \left(r \frac{\zeta_2 - 2\zeta_1}{\phi_1(\phi_1 + \phi_2)} + \frac{\phi_2}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right).$$

The FOC then requires that

$$r \frac{2\zeta_1 - \zeta_2}{\phi_1(\phi_1 + \phi_2)} = \frac{\phi_2}{\zeta_1 \phi_2 + \zeta_2 \phi_1}. \quad (24)$$

Now consider the FOC for source 2. Observe that $\frac{\partial t_1^*}{\partial \zeta_2} = \frac{\zeta_1}{\phi_1(\phi_1 + \phi_2)}$. So

$$\frac{\partial U_2}{\partial \zeta_2} = e^{-rt_1^*} \left(\frac{\zeta_1 \phi_1}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right) \left(-r \frac{\zeta_2}{\phi_1(\phi_1 + \phi_2)} + \frac{\phi_2}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right).$$

The FOC then requires that

$$r \frac{\zeta_2}{\phi_1(\phi_1 + \phi_2)} = \frac{\phi_2}{\zeta_1 \phi_2 + \zeta_2 \phi_1}. \quad (25)$$

(24) and (25) together imply

$$r \frac{2\zeta_1 - \zeta_2}{\phi_1(\phi_1 + \phi_2)} = r \frac{\zeta_2}{\phi_1(\phi_1 + \phi_2)},$$

which simplifies to $\zeta_1 = \zeta_2$ and leads to a contradiction.

Hence $\zeta_1 = \zeta_2$ must hold in equilibrium. In this case the first-order conditions derived above need not hold with equality, because the payoffs in (23) are derived under the assumption that $\zeta_1 \leq \zeta_2$, so that the same payoff expressions apply only to downward deviations of ζ_1 and upward deviations of ζ_2 . Given this, the first-order conditions become inequalities $\frac{\partial U_1}{\partial \zeta_1} \geq 0$ and $\frac{\partial U_2}{\partial \zeta_2} \leq 0$ (evaluated at the equilibrium choices). These translate into the following inequality versions of (24) and (25):

$$\begin{aligned} r \frac{2\zeta_1 - \zeta_2}{\phi_1(\phi_1 + \phi_2)} &\leq \frac{\phi_2}{\zeta_1\phi_2 + \zeta_2\phi_1}; \\ r \frac{\zeta_2}{\phi_1(\phi_1 + \phi_2)} &\geq \frac{\phi_2}{\zeta_1\phi_2 + \zeta_2\phi_1}. \end{aligned}$$

Since we already know $\zeta_1 = \zeta_2$, the two inequalities above must both hold equal, and we further deduce that

$$\zeta_1 = \zeta_2 = \sqrt{\frac{\phi_1\phi_2}{r}}. \quad (26)$$

Note also that given $\zeta_2 = \sqrt{\frac{\phi_1\phi_2}{r}}$, choosing ζ_1 to be any smaller number cannot be profitable for source 1. This is because as ζ_1 decreases, the term $r \frac{\zeta_2 - 2\zeta_1}{\phi_1(\phi_1 + \phi_2)} + \frac{\phi_2}{\zeta_1\phi_2 + \zeta_2\phi_1}$ appearing in $\frac{\partial U_1}{\partial \zeta_1}$ increases and remains positive. So the choice $\zeta_1 = \sqrt{\frac{\phi_1\phi_2}{r}}$ is robust to any downward deviation (in this variable). Similarly, given this ζ_1 , choosing any larger ζ_2 is not profitable for source 2. By symmetry, source 1 (respectively source 2) also cannot profit from upward (respectively downward) deviations in precision.

O.6.3 Solving for Equilibrium Biases ϕ_1^* , ϕ_2^*

We now fix precision choices and characterize equilibrium levels of bias. Since $\zeta_1 = \zeta_2$ in equilibrium, we have $t_1^* = 0$, meaning that there is no stage 1. Hence the two sources' payoffs simplify to

$$\begin{aligned} U_1 &= \frac{\phi_2}{\phi_1 + \phi_2} - \lambda(1 - \phi_1)^2; \\ U_2 &= \frac{\phi_1}{\phi_1 + \phi_2} - \lambda(1 - \phi_2)^2. \end{aligned} \quad (27)$$

In this smaller game, we will show that there is a (unique) pure strategy equilibrium if and only if $\lambda \geq \frac{9}{16}$, in which case the equilibrium involves $\phi_1 = \phi_2 = \frac{1}{2} \left(1 + \sqrt{1 - \frac{1}{2\lambda}}\right)$.

The first-order conditions $\frac{\partial U_i}{\partial \phi_i} = 0$ give

$$2\lambda(1 - \phi_1) = \frac{\phi_2}{(\phi_1 + \phi_2)^2}; \quad 2\lambda(1 - \phi_2) = \frac{\phi_1}{(\phi_1 + \phi_2)^2}. \quad (28)$$

In addition, the second-order conditions $\frac{\partial^2 U_i}{\partial \phi_i^2} \leq 0$ give

$$2\lambda \geq \frac{2\phi_2}{(\phi_1 + \phi_2)^3}; \quad 2\lambda \geq \frac{2\phi_1}{(\phi_1 + \phi_2)^3}. \quad (29)$$

Comparing each equality in (28) with the corresponding one in (29) yields

$$2(1 - \phi_1) \leq \phi_1 + \phi_2; \quad 2(1 - \phi_2) \leq \phi_1 + \phi_2. \quad (30)$$

Moreover, multiplying the two equalities in (28) yields $(1 - \phi_1)\phi_1 = (1 - \phi_2)\phi_2$. So either $\phi_1 = \phi_2$ or $\phi_1 + \phi_2 = 1$. In the former case (30) implies $\phi_1 = \phi_2 \geq \frac{1}{2}$. In the latter case (30) implies $\phi_1 = \phi_2 = \frac{1}{2}$. Thus we always have $\phi_1 = \phi_2 \geq \frac{1}{2}$.

Now since $\phi_1 = \phi_2$, we can solve from (28) that ϕ_1 satisfies $4\phi_1(1 - \phi_1) = \frac{1}{2\lambda}$, which is equivalent to $(2\phi_1 - 1)^2 = 1 - \frac{1}{2\lambda}$. Using $\phi_1 \geq \frac{1}{2}$, we deduce that the only possible equilibrium is $\phi_1 = \phi_2 = \frac{1}{2} \left(1 + \sqrt{1 - \frac{1}{2\lambda}}\right)$, which we denote by ϕ^* . Clearly, for this solution to be a real number, a necessary condition is $\lambda \geq \frac{1}{2}$.

Below we show these choices are an equilibrium in this smaller game (involving only ϕ_1, ϕ_2) if and only if $\lambda \geq \frac{9}{16}$. Indeed, with these choices source 1's payoff is $\frac{1}{2} - \lambda(1 - \phi^*)^2$. This must be higher than choosing ϕ_1 close to 0, which would yield a payoff close to $1 - \lambda$. Thus we can deduce the inequality $\phi^*(2 - \phi^*) \geq \frac{1}{2\lambda}$. But recall that ϕ^* satisfies the equation $4\phi^*(1 - \phi^*) = \frac{1}{2\lambda}$. So we deduce that $2 - \phi^* \geq 4(1 - \phi^*)$, or $\phi^* \geq \frac{2}{3}$. It follows that $\frac{1}{2\lambda} = 4\phi^*(1 - \phi^*) \leq \frac{8}{9}$, which implies $\lambda \geq \frac{9}{16}$ as a necessary condition.

Conversely, suppose $\lambda \geq \frac{9}{16}$ holds. We need to show that

$$f(\phi_1) = \frac{\phi^*}{\phi_1 + \phi^*} - \lambda(1 - \phi_1)^2$$

is maximized at $\phi_1 = \phi^*$. Since we have derived the solution via the first- and second-order conditions, it holds that $f'(\phi^*) = 0$ and $f''(\phi^*) \leq 0$. Moreover, note that $f'''(\phi_1) = \frac{-6\phi^*}{(\phi_1 + \phi^*)^4} < 0$, so $f'(\phi_1)$ is a strictly concave function. The fact that $f'(\phi^*) = 0$ and $f''(\phi^*) \leq 0$ thus imply that $f'(\phi_1) \leq 0$ for all $\phi_1 \geq \phi^*$. Hence $f(\phi_1)$ is decreasing for $\phi_1 \geq \phi^*$.

On the other hand, since $f'(\phi_1)$ is concave and $f'(\phi^*) = 0$, there are two possibilities for the behavior of f' on the interval $[0, \phi^*]$: Either $f'(0) \geq 0$ and thus f' is non-negative on this whole interval, or $f'(0) < 0$ and f' crosses zero exactly once from below. This means f is either increasing on $[0, \phi^*]$, or first decreasing and then increasing. Hence the maximum of f on this interval must occur at the extreme points. When $\lambda \geq \frac{9}{16}$, we have $\phi^* = \frac{1}{2} \left(1 + \sqrt{1 - \frac{1}{2\lambda}}\right) \geq \frac{2}{3}$. Thus

$$f(\phi^*) = \frac{1}{2} - \lambda(1 - \phi^*)^2 \geq \frac{1}{2} - \frac{1}{9}\lambda \geq 1 - \lambda = f(0).$$

This shows that the function f is maximized at ϕ^* whenever $\lambda \geq \frac{9}{16}$.

O.6.4 Verifying the Equilibrium

Summarizing the above analysis, we have shown that the only possible pure strategy equilibrium is $\phi_1 = \phi_2 = \phi^*$ and $\zeta_1 = \zeta_2 = \frac{\phi^*}{\sqrt{r}}$, where the latter follows from (26). We have also shown that if

$\lambda \geq \frac{9}{16}$, then given source 2's equilibrium choices, source 1 does not have a profitable deviation in ζ_1 or in ϕ_1 alone. However, without further assumptions, it is possible for source 1 to profit from choosing ζ_1 and ϕ_1 both away from the target equilibrium. To illustrate, let $\lambda = \frac{9}{16}$ and $r = 1$. Then source 2's equilibrium choices are $\phi_2 = \zeta_2 = \phi^* = \frac{2}{3}$. By also choosing $\phi_1 = \zeta_1 = \frac{2}{3}$, source 1 obtains payoff $\frac{1}{2} - \lambda(1 - \phi^*)^2 = \frac{7}{16} = 0.4375$. Suppose instead source 1 chooses $\phi_1 = \frac{1}{6}$ and $\zeta_1 = \frac{1}{3}$. Then $t_1^* = \frac{\zeta_1(\zeta_2 - \zeta_1)}{\phi_1(\phi_1 + \phi_2)} = \frac{4}{5}$, and source 1 receives long-run attention $\frac{\zeta_1\phi_2}{\zeta_1\phi_2 + \zeta_2\phi_1} = \frac{2}{3}$. In this case source 1's payoff is higher:

$$1 - \frac{1}{3}e^{-\frac{4}{5}} - \frac{9}{16}\left(1 - \frac{1}{6}\right)^2 \approx 0.4596.$$

So the target equilibrium would not be an equilibrium when double deviations are considered.

However, we will show that when $\lambda \geq 1.6$, then given $\phi_2 = \phi^*$ and $\zeta_2 = \frac{\phi^*}{\sqrt{r}}$, source 1 cannot profitably deviate to any ϕ_1 and $\zeta_1 \leq \frac{\phi^*}{\sqrt{r}}$. And given $\phi_1 = \phi^*$ and $\zeta_1 = \frac{\phi^*}{\sqrt{r}}$, source 2 cannot profitably deviate to any ϕ_2 and $\zeta_2 \geq \frac{\phi^*}{\sqrt{r}}$. Thanks to the symmetry of the target equilibrium, verifying these will be sufficient to show it is indeed an equilibrium.

First consider source 1. We have the elementary inequality

$$e^{-rt_1^*} \geq 1 - rt_1^* = \frac{\phi_1(\phi_1 + \phi_2) - r\zeta_1(\zeta_2 - \zeta_1)}{\phi_1(\phi_1 + \phi_2)},$$

which is tight if and only if $\zeta_1 = \zeta_2$ and $t_1^* = 0$. Thus from (23) we have

$$U_1 \leq 1 - \frac{\phi_1(\phi_1 + \phi_2) - r\zeta_1(\zeta_2 - \zeta_1)}{\phi_1(\phi_1 + \phi_2)} \cdot \frac{\zeta_2\phi_1}{\zeta_1\phi_2 + \zeta_2\phi_1} - \lambda(1 - \phi_1)^2.$$

Plugging in $\phi_2 = \phi^*$ and $\zeta_2 = \frac{\phi^*}{\sqrt{r}}$, the above simplifies to

$$\begin{aligned} U_1 &\leq 1 - \frac{\phi_1(\phi_1 + \phi^*) - \sqrt{r}\zeta_1\phi^* + r\zeta_1^2}{(\phi_1 + \phi^*)(\phi_1 + \sqrt{r}\zeta_1)} - \lambda(1 - \phi_1)^2 \\ &= \frac{(\phi_1 + 2\phi^* - \sqrt{r}\zeta_1)\sqrt{r}\zeta_1}{(\phi_1 + \phi^*)(\phi_1 + \sqrt{r}\zeta_1)} - \lambda(1 - \phi_1)^2. \end{aligned}$$

We now optimize this upper-bound over ζ_1 , and then over ϕ_1 . Let $y = \phi_1 + \sqrt{r}\zeta_1$. Note that the range of y is $y \in [\phi_1, \phi_1 + \phi^*]$, since $\sqrt{r}\zeta_1 \leq \sqrt{r}\zeta_2 = \phi^*$. Using y , we can rewrite the above inequality as

$$U_1 \leq \frac{(2\phi_1 + 2\phi^* - y)(y - \phi_1)}{(\phi_1 + \phi^*)y} - \lambda(1 - \phi_1)^2. \quad (31)$$

Thus, in terms of y , we would like to maximize

$$\frac{(2\phi_1 + 2\phi^* - y)(y - \phi_1)}{y} = -y - \frac{2(\phi_1 + \phi^*)\phi_1}{y} + 2\phi^* + 3\phi_1$$

This is a single-peaked function in y , with global maximum occurring at $y = \sqrt{2(\phi_1 + \phi^*)\phi_1} \geq \phi_1$. We distinguish two cases. If $\sqrt{2(\phi_1 + \phi^*)\phi_1} \geq \phi_1 + \phi^*$, then the maximum of this function of y over the interval $[\phi_1, \phi_1 + \phi^*]$ occurs at $y = \phi_1 + \phi^*$. Plugging back into (31), we have

$$U_1 \leq \frac{\phi^*}{\phi_1 + \phi^*} - \lambda(1 - \phi_1)^2.$$

But in previous analysis we have already shown that the above function of ϕ_1 is maximized at ϕ^* whenever $\lambda \geq \frac{9}{16}$, so we are done in this case.

The more difficult case is $y = \sqrt{2(\phi_1 + \phi^*)\phi_1} \leq \phi_1 + \phi^*$, which corresponds to $\phi_1 \leq \phi^*$. Here the global maximum is achievable on the interval $y \in [\phi_1, \phi_1 + \phi^*]$, so we instead deduce from (31) the following:

$$U_1 \leq \frac{3\phi_1 + 2\phi^* - 2\sqrt{2(\phi_1 + \phi^*)\phi_1}}{\phi_1 + \phi^*} - \lambda(1 - \phi_1)^2.$$

We denote the function on the RHS as $g(\phi_1)$, and aim to show $g(\phi_1) \leq g(\phi^*)$ for all $\phi_1 \leq \phi^*$. Note that

$$g'(\phi_1) = \frac{\phi^*}{(\phi_1 + \phi^*)^2} \cdot \left(1 - \sqrt{\frac{2(\phi_1 + \phi^*)}{\phi_1}}\right) + 2\lambda(1 - \phi_1).$$

Thus $g'(\phi^*) = \frac{-1}{4\phi^*} + 2\lambda(1 - \phi^*) = 0$. Moreover, we claim that g' is concave on the interval $[0, \phi^*]$.

To see this, let us write

$$\phi^* g'(\phi_1) = \frac{1}{\left(\frac{\phi_1}{\phi^*} + 1\right)^2} \cdot \left(1 - \sqrt{2 + \frac{2\phi^*}{\phi_1}}\right) + 2\phi^* \lambda(1 - \phi_1).$$

The second term $2\phi^* \lambda(1 - \phi_1)$ is linear in ϕ_1 and thus does not affect convexity/concavity. As for the first term, we can rewrite it as

$$h(z) = \frac{1}{(1+z)^2} \cdot \left(1 - \sqrt{2 + \frac{2}{z}}\right)$$

with $z = \frac{\phi_1}{\phi^*}$. It thus suffices to show h is concave on the interval $[0, 1]$. For this we compute that

$$h''(z) = -3 \frac{-8\sqrt{z^5(z+1)} + 8\sqrt{2}z^3 + 12\sqrt{2}z^2 + 5\sqrt{2}z + \sqrt{2}}{4z^{2.5}(z+1)^{4.5}},$$

which is negative for $z \leq 1$ because $8\sqrt{z^5(z+1)} \leq 8\sqrt{2z^5} \leq 8\sqrt{2}z^2$.

Hence $g'(\phi^*) = 0$ and g' is concave on $[0, \phi^*]$. It follows that either g' is non-negative on this whole interval, or it is first negative then positive. So g is either increasing, or first decreasing and then increasing. It thus remains to show $g(0) \leq g(\phi^*)$, which reduces to $2 - \lambda \leq \frac{1}{2} - \lambda(1 - \phi^*)^2$. Since $\lambda \geq 1.6$, we have $\phi^* = \frac{1}{2} \left(1 + \sqrt{1 - \frac{1}{2\lambda}}\right) > \frac{3}{4}$. Thus we indeed have

$$2 - \lambda \leq \frac{1}{2} - \frac{1}{16}\lambda < \frac{1}{2} - \lambda(1 - \phi^*)^2,$$

where the first inequality again uses the assumption $\lambda \geq 1.6$. This completes the proof that source 1 does not have any profitable deviation.

We now turn to source 2's incentives. We use another elementary inequality

$$e^{rt_1^*} \geq 1 + rt_1^* = 1 + r \frac{\zeta_1(\zeta_2 - \zeta_1)}{\phi_1(\phi_1 + \phi_2)}.$$

Given $\phi_1 = \phi^*$ and $\zeta_1 = \frac{\phi^*}{\sqrt{r}}$, we can simplify the above as

$$e^{rt_1^*} \geq 1 + \frac{\sqrt{r}(\zeta_2 - \zeta_1)}{\phi_1 + \phi_2} = \frac{\phi_2 + \sqrt{r}\zeta_2}{\phi_2 + \phi^*}.$$

So from (23) we obtain an upper bound for U_2 :

$$\begin{aligned} U_2 &= e^{-rt_1^*} \left(\frac{\zeta_2 \phi_1}{\zeta_1 \phi_2 + \zeta_2 \phi_1} \right) - \lambda(1 - \phi_2)^2 \\ &\leq \frac{\phi_2 + \phi^*}{\phi_2 + \sqrt{r}\zeta_2} \cdot \frac{\sqrt{r}\zeta_2}{\phi_2 + \sqrt{r}\zeta_2} - \lambda(1 - \phi_2)^2 \\ &= (\phi_2 + \phi^*) \frac{\sqrt{r}\zeta_2}{(\phi_2 + \sqrt{r}\zeta_2)^2} - \lambda(1 - \phi_2)^2 \end{aligned} \quad (32)$$

Again we will optimize this upper bound first over ζ_2 , then over ϕ_2 . In terms of ζ_2 , we want to minimize

$$\frac{(\phi_2 + \sqrt{r}\zeta_2)^2}{\sqrt{r}\zeta_2} = \sqrt{r}\zeta_2 + \frac{\phi_2^2}{\sqrt{r}\zeta_2} + 2\phi_2.$$

The global minimum of this single-dipped function occurs at $\sqrt{r}\zeta_2 = \phi_2$, but since $\sqrt{r}\zeta_2 \geq \sqrt{r}\zeta_1 = \phi^*$, there are two cases to consider. In the first case, $\phi_2 \leq \phi^*$. Then the minimum subject to $\sqrt{r}\zeta_2 \geq \phi^*$ occurs precisely at $\sqrt{r}\zeta_2 = \phi^*$. In this case the upper bound (32) becomes $U_2 \leq \frac{\phi^*}{\phi_2 + \phi^*} - \lambda(1 - \phi_2)^2$, and as we have shown previously this is maximized at $\phi_2 = \phi^*$.

In the other case, we have $\phi_2 \geq \phi^*$. Then the optimal ζ_2 that maximizes the upper bound (32) is $\sqrt{r}\zeta_2 = \phi_2$. (32) then becomes

$$U_2 \leq \frac{\phi_2 + \phi^*}{4\phi_2} - \lambda(1 - \phi_2)^2.$$

We will show that the derivative of this function of ϕ_2 is negative for any $\phi_2 \geq \phi^*$, so that $\phi_2 = \phi^*$ is again the optimal choice. This comparison reduces to $2\lambda(1 - \phi_2) \leq \frac{\phi^*}{4\phi_2^2}$, which is equivalent to $8\lambda(1 - \phi_2)\phi_2^2 \leq \phi^*$. Since equality obtains when $\phi_2 = \phi^*$, we just need to show $8\lambda(1 - \phi_2)\phi_2^2 \leq 8\lambda(1 - \phi^*)(\phi^*)^2$. After factoring out $\phi_2 - \phi^*$, this inequality reduces to

$$\phi_2 + \phi^* \leq \phi_2^2 + \phi_2\phi^* + (\phi^*)^2.$$

This holds whenever $\phi^* \geq \frac{2}{3}$, since $\phi_2 \geq \phi^*$ implies $\phi_2^2 + \phi_2\phi^* + (\phi^*)^2 \geq \frac{3}{2}\phi^*(\phi_2 + \phi^*)$. Hence whenever $\lambda \geq \frac{9}{16}$, source 2 does not have profitable deviations either.