

Dynamically Aggregating Diverse Information ^{*}

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

March 1, 2020

Abstract

An agent has access to multiple information sources, each of which provides information about a different attribute of an unknown state. Information is acquired continuously—where the agent chooses both which sources to sample from, and also how to allocate resources across them—until an endogenously chosen time, at which point a decision is taken. We show that the optimal information acquisition strategy proceeds in stages, where resource allocation is constant over a fixed set of providers during each stage, and at each stage a new provider is added to the set. We additionally apply this characterization to derive results regarding: (1) endogenous information acquisition in a binary choice problem, and (2) equilibrium information provision by competing news sources.

^{*}We are grateful to Yash Deshpande, Mira Frick, Drew Fudenberg, Boyan Jovanovic, George Mailath, Konrad Mierendorff, Lan Min, Peter Norman Sørensen, Jakub Steiner, and Philipp Strack for helpful comments and suggestions. Xiaosheng Mu acknowledges the hospitality of the Cowles Foundation at Yale University, which hosted him during parts of this research.

[†]Department of Economics, University of Pennsylvania

[‡]Department of Economics, Columbia University

[§]Microsoft Research

1 Introduction

We study dynamic acquisition of information when a decision-maker has access to multiple kinds of information, and limited resources with which to acquire that information. Our decision-maker seeks to learn a Gaussian state, and we model each information source as a Brownian motion whose drift is an unknown *attribute* which is correlated with the state and with other attributes. This structure captures information acquisition in many economic settings, including for example:

- An investor wants to learn the value of a stock portfolio, and can acquire information about the value of each stock included in the portfolio.
- A voter wants to learn the quality of a candidate, and can acquire information about the views of various political pundits, including those biased in favor of and against the candidate.
- An analyst wants to forecast a macroeconomic variable such as GDP growth, and can acquire information about recent economic activity across industries and locations.

At every instant of time, the decision-maker allocates a fixed budget of attention across the information sources, where attention increases the amount/precision of information extracted from the source. This information is used for a future decision taken at an endogenously chosen stopping time.

Our model resembles, but does not fall under, the classic multi-armed bandit (MAB) framework (Gittins, 1979; Bergemann and Välimäki, 2008). To see this, recall that in MAB, the choice of which arm to pull plays the dual role of influencing the evolution of beliefs and also determining flow payoffs. In our setting, *information acquisition choices* influence the evolution of beliefs, whereas *actions*—taken separately—determine payoffs. Thus in our paper, information acquisition decisions are driven by learning concerns exclusively, and the exploration-exploitation trade-off central to bandit models does not appear.¹

The static version of our problem, in which the decision-maker acquires information at one instant only and takes an action immediately thereafter, is straightforward. Because normal signals can be completely Blackwell-ordered based on their precisions (Hansen and Torgersen, 1974), different attention allocations (i.e., different mixtures over the sources) can be compared based on how much they reduce the variance of the payoff-relevant state.

¹This 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).

Moreover, because these are Blackwell comparisons, the optimal attention allocation does not depend on the decision problem that the decision-maker faces. Our problem is also straightforward if information is acquired over a known interval of time, as the decision-maker should acquire information (in any order) to minimize uncertainty about the payoff-relevant state at the known end date.

But if the decision time is not known ex-ante, then the decision-maker may have to trade off between learning more about the state in a given period of time versus acquiring information that will result in better decisions later on. This trade-off arises because a given source not only provides information about the state; it also alters the information value of the other sources due to correlation. To solve the dynamic problem, the decision-maker has to take into account how acquisitions today change the value of information tomorrow; these dynamic externalities can be quite complex to describe.

Our contribution is to demonstrate that the optimal dynamic acquisition strategy can nevertheless be explicitly characterized, so long as the unknown attributes are not too strongly correlated. Under this strategy, the decision-maker initially exclusively observes 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 his 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. Similar to the solution for the static problem, the optimal information acquisition strategy holds for all decision problems.

The main idea in the proof is the following. Consider the class of information acquisition problems with a known (exogenously given) decision time t . Intertemporal trade-offs exist if the optimal acquisitions for some time t are “in conflict” with those for a later time $t + \Delta$, forcing the decision-maker to choose between what is best for the possible decision time t versus what is best for $t + \Delta$. Our key observation is that these trade-offs exist *only* if the best allocation of $t + \Delta$ units of attention involves lower attention to some source compared to the best allocation of t units. This would make it impossible for any sampling strategy to simultaneously optimize for decisions at times t and $t + \Delta$.

If however the optimal attention allocations across different times t are non-decreasing for each source, then between any t and $t + \Delta$, the decision-maker can simply choose the allocation that takes him from the optimal acquisitions for t to the optimal acquisitions for $t + \Delta$. Such a strategy would maximize learning about the payoff-relevant state at all times, a property that we call “uniformly optimal.” We show that so long as the different attributes are not too strongly correlated, a uniformly optimal strategy exists and has the

nested structure that we described above. See Section 5 for more detail.

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 comparative statics in 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 two applications:

The first setting that we consider is endogenous information acquisition for binary choice. A large literature in economics and neuroscience (originating with [Ratcliff and McKoon \(2008\)](#)) models a consumer’s decision process for choosing between two goods with unknown payoffs. Although this literature has primarily focused on optimal stopping times given exogenous information, a model in [Fudenberg et al. \(2018\)](#) endogenizes the information acquisition process. They show that if payoffs are Gaussian, independent and symmetric, then the decision-maker optimally mixes equally over the sources at every moment in time.

This model is nested in our framework as the case of two unknown attributes (the unknown payoffs), and the decision-maker wants to learn the difference of these attributes (as this is a sufficient statistic for which payoff is larger). A straightforward corollary of our main result generalizes the [Fudenberg et al. \(2018\)](#) result to correlated payoffs, asymmetric initial uncertainty, and asymmetric levels of source informativeness. In addition, we can use our characterization to derive new comparative static results with respect to these primitives. We find that an increase in initial uncertainty about either payoff results in uniformly more attention paid to learning about that payoff at every instant, while an increase in signal noise has an ambiguous effect (which we describe). We also consider a comparative static in prior correlation across the payoffs, and find that an increase in the size of correlation asymmetrically favors the source that the decision-maker attends to first. All of these are new and empirically testable predictions.

In our next application, we consider a game between strategic information sources. Media sources compete over readers’ attention by choosing the precision of the information they provide. For any given precision levels, our main results tell us the optimal time path of attention for the readers; this allows us to derive each source’s best-reply function. The key insight is that sources face a trade-off: providing informative articles increases the immediate competitive value of the source, but releasing information too quickly reduces the need for readers to continue to engage with the source. We find that equilibrium level of news informativeness is *higher* when the information providers are *less patient* and the information

they provide are *more positively correlated*.

This part of our analysis contributes to a growing literature about how competition across news sources affects the quality of news (Gentzkow and Shapiro, 2008). In particular, a large literature studies endogenous choice of media slant (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Chan and Suen, 2008; Perego and Yuksel, 2020), and recent models have additionally endogenized news informativeness (Galperti and Trevino, 2020; Chen and Suen, 2020). Our analysis focuses on this latter aspect of news informativeness. Different from the prior work, we consider the effect of information precision on the *time path* of people’s information demand, and how these dynamic considerations affect the informativeness of news. As far as we are aware, identification of how time preferences interact with the informational environment is new, and our tractable characterizations for the dynamic attention path are what allow us to study this.

1.1 Related Literature

We build on a large literature about optimal dynamic information acquisition. In contrast to an earlier focus in the literature on the choice of signal precisions (Moscarini and Smith, 2001), our framework characterizes the choice between many *kinds* of information, each providing information about a different unknown. Our model is closest in this respect to Fudenberg et al. (2018) and Gossner et al. (2019).² In Fudenberg et al. (2018), the agent can learn about the (independent) values of two goods by observing the evolution of diffusion processes, and in Gossner et al. (2019), the agent can learn about the values of each of K goods (again, independent) by observing Bernoulli signals.³ Compared to these papers, we study a setting where the agent dynamically learns about many *correlated* attributes.

There is not a large prior literature on dynamic learning in the presence of correlation. One interesting model is that of 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 in several productive ways: Garfagnini and Strulovici (2016) consider the optimal experimentation strategy for a forward-looking agent with acquisition costs, while Bardhi (2019) studies gen-

²Che and Mierendorff (2019) and Mayskaya (2019) also consider choice from a prescribed set of information sources, but they focus on Poisson signals that confirm either of two states.

³Gossner et al. (2019) study the consequences of attention manipulations, where the agent is forced to attend initially to one particular attribute. This interesting question bears certain high-level resemblances to our comparative statics in Section 6. However, we focus on consequences for the time path of attention, instead of consequences for the final decision (which good is chosen), as Gossner et al. (2019) do.

eral Gaussian sample paths and introduces potential conflict between an agent acquiring the information and a principal making the decision. These informational environments 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.⁴

In the context of learning about multiple attributes, [Klabjan et al. \(2014\)](#) and [Sanjurjo \(2017\)](#) study a search problem where each attribute value is perfectly learned upon a single inspection. Working with general distributions, these authors show that an attribute is “more attractive for discovery” than another attribute whenever its distribution is a mean-preserving spread of the latter. Besides having noisy Gaussian signals, the main distinction of our informational setting is that we allow for correlation across attributes—much of our analysis regards what this correlation implies for the optimal strategy.

Another strand of the literature considers agents who choose from completely flexible information structures at entropic (or more generally, “posterior-separable”) costs, such as in [Yang \(2015\)](#), [Steiner et al. \(2017\)](#), [Hébert and Woodford \(2018\)](#), [Morris and Strack \(2019\)](#), and [Zhong \(2019\)](#).⁵ Compared to these papers, our agent has access to a prescribed (physical) set of signals, and acquires information under an attention capacity constraint. Thus the different signals in our setting are equally costly to acquire regardless of the current belief, which is the key distinction from measuring information acquisition costs by the reduction of uncertainty.⁶

In previous work ([Liang et al., 2017](#)), we studied a related setting in discrete time, introduced the notion of “myopic information acquisition” and studied its (approximate) optimality properties.⁷ We did not obtain a characterization of the optimal strategy itself. Going beyond those results, the characterizations in the present paper precisely (and more generally) describe the optimal path of attention allocations, which are useful in applications

⁴We also note that our main result is unlikely to apply to their setting, since “close” attributes are very strongly correlated, violating the condition on the prior that we require.

⁵It is interesting that [Steiner et al. \(2017\)](#) also show 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.

⁶Formally, we consider a sequential sampling problem in which the flow cost of acquiring information only depends on the current time and not on the current belief.

⁷In the present paper as well as in [Liang et al. \(2017\)](#), we study the complete path of information acquisitions, but one corollary of the main results in these papers is that information acquisitions under a myopic procedure will be asymptotically efficient. In [Liang and Mu \(2020\)](#), we provide a more thorough analysis of the conditions on the informational environment under which myopic acquisitions lead to long-run (in)efficient learning.

as we illustrate. The technical methods in this paper also differ from the prior work—see Section 8 for further discussion.

Finally, this paper is related to recent work on data acquisition by firms. Azevedo et al. (2019) study allocation of resources (i.e., test users) to learn about the quality of multiple innovations. These authors show that the tail distribution of innovation quality crucially affects the (static) optimal experimentation strategy. Immorlica et al. (2018) consider dynamic allocations of a budget of data samples for learning about an evolving state, and demonstrate near-efficiency guarantees for certain classes of benchmark policies.

2 Model

An agent has access to K *information sources*, each of which is a diffusion process that provides information about an unknown *attribute* $\theta_k \in \mathbb{R}$. The random vector $(\theta_1, \dots, \theta_K)$ is jointly normal with a known prior mean vector 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 an affine function of the attributes:

Assumption 1. $\omega = \sum_{i=1}^K \alpha_i \theta_i + b$ for some weights $\alpha_1, \dots, \alpha_K \in \mathbb{R}$ and constant $b \in \mathbb{R}$.

It is equivalent to assume that ω is jointly normally distributed together with the θ_i , and that there is no residual uncertainty about ω given complete knowledge of the attribute values.⁸

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 each weight α_i is strictly positive. Intuitively, an attribute with zero payoff weight does not matter for learning about ω ; we verify this in Appendix D.6. The weights $\alpha_1, \dots, \alpha_K$ along with the prior covariance matrix Σ are the key primitives of our model.

Time is continuous, and the agent has a 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).⁹

⁸If $\omega, \theta_1, \dots, \theta_K$ are jointly normal, then the conditional distribution of $\omega \mid \theta_1, \dots, \theta_K$ is itself a normal distribution whose mean is a linear combination of $\theta_1, \dots, \theta_K$ and the prior mean of ω . The assumption of no residual uncertainty means that the conditional variance is zero, returning Assumption 1.

⁹Recent models that feature fixed budgets of attention include Fudenberg et al. (2018) and Che and Mierendorff (2019).

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

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

Above, each \mathcal{B}_i is an independent standard 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.¹⁰ In particular, 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 $\theta_i + \mathcal{N}(0, 1)$.¹¹

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 $(\beta_1(t), \dots, \beta_K(t))$ results in observation of $\theta_i + \mathcal{N}(0, 1/\beta_i(t))$ for each $i = 1, \dots, K$. Our results are somewhat easier to state for the continuous-time framework, so we take this to be our main model, but as we show in Section 8, each of our results directly implies a corresponding result for the discrete-time framework.

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 $\{X_i^{\leq t}\}_{i=1}^K$. An *information acquisition strategy* S is a map from observations $\{X_i^{\leq t}\}_{i=1}^K$ into $\Delta(\{1, \dots, K\})$, representing how the agent divides attention at each instant as a function of the observed diffusion processes. 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 sampling information at some flow cost, or to stop acquiring information 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 $u(a, \omega)$, where u is a known payoff function that depends on the action taken a and the payoff-relevant state ω . The agent's posterior belief about ω at this time determines the action that maximizes his expected flow payoff $\mathbb{E}[u(a, \omega)]$.

¹⁰Having constant informativeness *across* sources implies that it is *with loss* to further normalize the payoff weights α_i to be equal. Indeed, our subsequent results indicate that the case of equal weights is special. For example, with $K = 2$, the conclusions of Theorem 1 always hold when $\alpha_1 = \alpha_2$ but not in general.

¹¹Note that this definition also treats “attention” and “time” in the same way, in the 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.

To summarize, the agent chooses his information acquisition strategy and stopping time (S, τ) to maximize

$$\max_{S, \tau} \mathbb{E} \left[\max_a \mathbb{E}[u(a, \omega) | \mathcal{F}_\tau] - c(\tau) \right],$$

where $c(\tau)$ is a non-negative and weakly increasing function that measures the cost of waiting until time τ .¹² Our focus throughout this paper is on 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 Preliminaries

At every time t , the agent's past attention allocations integrate to a *cumulated attention vector*

$$q(t) = (q_1(t), \dots, q_K(t))' \in \mathbb{R}_+^K$$

describing how much attention has been paid to each source thus far. A useful property of Bayesian updating from Gaussian signals is that the agent's posterior covariance matrix about $(\theta_1, \dots, \theta_K)$ can be expressed as a function *solely* of $q(t)$, and in particular does not depend on the realizations of the diffusion processes. This posterior covariance matrix is

$$(\Sigma^{-1} + \text{diag}(q(t)))^{-1}, \tag{1}$$

where Σ is the prior covariance matrix over the attribute values, and $\text{diag}(q(t))$ is the diagonal matrix with entries $q_1(t), \dots, q_K(t)$. The above formula says that the posterior precision matrix (i.e., inverse of the posterior covariance matrix) is the sum of the prior precision matrix (Σ^{-1} in this case) and the signal precision matrix ($\text{diag}(q(t))$ in this case).

Due to the Gaussian structure of the problem, maximizing the informativeness of the learning process about ω up to time t is equivalent to minimizing the posterior variance of this payoff-relevant state. Using (1) and the decomposition of ω from Assumption 1, the agent's posterior variance about ω is

$$V(q) = \alpha'(\Sigma^{-1} + \text{diag}(q))^{-1}\alpha. \tag{2}$$

This function V is globally convex, differentiable, and decreasing in each q_i . See Appendix A for the proof of these and additional properties.

¹²Adding geometric or other forms of discounting to the model would not affect any of the results.

In the special case in which the agent stops at a fixed and known time T , every information acquisition strategy that minimizes posterior variance about ω at time T is optimal (see Section 5 for details), and the order of acquisitions does not matter. Note that the sequence of acquisitions *will* matter outside of this special setting.

4 Optimal Information Acquisition Strategy

We begin with the case of two information sources, as the simpler setting allows us to derive stronger results and explain certain key intuitions. Following this we present results for the case of any finite number of sources, as well as an extended outline of our proof strategy.

4.1 $K = 2$

Suppose there are two information sources and two attributes θ_1 and θ_2 . The agent seeks to learn $\omega = \alpha_1\theta_1 + \alpha_2\theta_2$, with each $\alpha_i > 0$. His prior over the unknown attributes is

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \right).$$

The covariances between the attributes and the payoff-relevant state are

$$cov_i := \text{Cov}(\omega, \theta_i) = \alpha_i \Sigma_{ii} + \alpha_j \Sigma_{ji}$$

and we assume that these covariances satisfy the following relationship:

Assumption 2. $cov_1 + cov_2 = \alpha_1(\Sigma_{11} + \Sigma_{12}) + \alpha_2(\Sigma_{21} + \Sigma_{22}) \geq 0$.

Since both variances Σ_{11}, Σ_{22} are positive, Assumption 2 can be understood as requiring that the covariance Σ_{12} is not too negative relative to the size of either variance. A sufficient condition is for the weights on the two attributes to be equal (i.e., $\alpha_1 = \alpha_2$), in which case Assumption 2 holds for all priors.¹³ Another sufficient condition is for the attributes to be positively correlated ($\Sigma_{12} = \Sigma_{21} \geq 0$), in which case Assumption 2 holds for all weights α_1 and α_2 .

Our first result establishes the optimal information acquisition strategy under this assumption.

¹³This follows from $2 \cdot |\Sigma_{12}| \leq 2 \cdot \sqrt{\Sigma_{11} \cdot \Sigma_{22}} \leq \Sigma_{11} + \Sigma_{22}$.

Theorem 1. Suppose Assumption 2 is satisfied.¹⁴ Define

$$t_i^* := \frac{\text{cov}_i - \text{cov}_j}{\alpha_j \det(\Sigma)}.$$

W.l.o.g. let $\text{cov}_i \geq \text{cov}_j$. Then an optimal information acquisition strategy 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 $t \leq t_i^*$, 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 $t > t_i^*$, the agent optimally allocates attention in the constant proportion $(\beta_1(t), \beta_2(t)) = \left(\frac{\alpha_1}{\alpha_1 + \alpha_2}, \frac{\alpha_2}{\alpha_1 + \alpha_2} \right)$.

Under mild assumptions on the primitives, this optimal strategy is in fact unique up to the stopping time τ (after which attention allocations obviously do not matter). We defer the technical discussion to Appendix A.3.

Thus there are two stages of information acquisition. In the first stage, which ends at some t^* , the agent allocates all of his attention to one of the attributes. After time t^* , he divides his attention across the attributes in a constant ratio across time. 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 realized Brownian motion paths. But whenever the agent continues acquiring information, his optimal attention allocations are as given above.

The characterization reveals that the optimal information acquisition strategy is completely determined from the prior covariance matrix Σ and the weight vector α . In particular, it does not depend on the agent's cost of waiting or the payoff function. Thus, when the prior belief satisfies Assumption 2, the optimal information acquisition strategy is constant across different objectives and also across different stopping rules. Relatedly, we can solve for the optimal stopping rule in this setting as if information acquisition were *exogenously* given by Theorem 1.

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

¹⁴The condition provided in Assumption 2 for $K = 2$ is not only sufficient but also necessary for our characterization to hold independently of the agent's payoff criterion—see Appendix B.1 and Proposition 4 for details.

Example 1 (Independent Attributes). First consider the case of independent attributes. For example, suppose the two unknown attribute values are distributed as

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} 6 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

under the agent's prior, and he wants to learn $\theta_1 + \theta_2$. Then, applying Theorem 1, the agent begins by putting 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). Now suppose the attributes are correlated; for example, the unknown attribute values are distributed as

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} 6 & 2 \\ 2 & 1 \end{pmatrix} \right)$$

under the agent's prior, and he wants to learn $\theta_1 + \theta_2$. Applying Theorem 1, the agent begins by putting all attention towards learning θ_1 . At time $t_1^* = \frac{5}{2}$, his posterior covariance matrix as given by (1) becomes $\begin{pmatrix} 3/8 & 1/8 \\ 1/8 & 3/8 \end{pmatrix}$, which makes the two attributes symmetric. After this time 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 agent wants to learn $\theta_1 + 2\theta_2$. As before, the agent begins by putting all attention towards learning θ_1 . Stage 1 ends at time $t_1^* = \frac{3}{2}$, when the posterior covariance matrix is $\begin{pmatrix} 3/5 & 1/5 \\ 1/5 & 2/5 \end{pmatrix}$. After this time, he optimally acquires information in the mixture $(1/3, 2/3)$.

To interpret the optimal strategy, first consider the case of equal payoff weights ($\alpha_1 = \alpha_2$), as in Examples 1 and 2. Then, the condition $cov_1 = \alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{21} \geq \alpha_1 \Sigma_{21} + \alpha_2 \Sigma_{12} = cov_2$ reduces to $\Sigma_{11} \geq \Sigma_{22}$. So Stage 1 involves a direct comparison of prior uncertainty about the two attributes, where the agent initially chooses to learn exclusively about the attribute over which he is more uncertain. More generally, we can measure value of information by how much it reduces the variance of the payoff-relevant state ω . Then the condition $cov_1 \geq cov_2$ equivalently says that the *marginal value* of learning about attribute θ_1 exceeds that of learning about θ_2 , according to the prior belief.

Suppose without loss of generality that $cov_1 \geq cov_2$, so that the agent initially learns exclusively about θ_1 , which has greater marginal value. As information about θ_1 accumulates, the marginal values of learning either attribute evolve, with the marginal value of θ_1 decreasing faster than θ_2 . Eventually, these *marginal values equalize*. From this point on, the agent optimally acquires information in a constant ratio that is proportional to the weight vector α . Dividing attention in this way achieves the most efficient aggregation of information about ω . Moreover, as we show in the proof, acquisition of information proportional to α maintains equal marginal values of the two information sources, so that acquiring information in this mixture remains optimal. We provide a more involved proof outline in Section 5, but the intuition is readily seen through the previous examples. In Examples 1 and 2, since the agent seeks to learn $\theta_1 + \theta_2$, the two attributes become symmetric once their posterior variances equalize. After that, equal attention allocation maintains symmetry and equal marginal values.

Although symmetry is lost in Example 3, the posterior covariance matrix $\begin{pmatrix} 3/5 & 1/5 \\ 1/5 & 2/5 \end{pmatrix}$ at time $t_1^* = \frac{3}{2}$ has the key property that the payoff-relevant state $\omega = \theta_1 + 2\theta_2$ is *independent* of $\theta_1 - \theta_2$, as they are jointly normal and have zero covariance.¹⁵ As we show in Lemma 5, this independence property implies equal marginal values.¹⁶ This explains why the agent is willing to mix at time t_1^* . The specific mixture (1/3, 2/3) ensures that every subsequent posterior covariance matrix continues to have the independence property. Hence equal marginal values are maintained, and the agent optimally follows this mixture at future times as well.

4.2 General K

We now consider the case of multiple attributes, where we will show that the results for the $K = 2$ case extend qualitatively.

A sufficient condition on the prior belief, parallel to the one stated in Assumption 2, is the following:

Assumption 3. *The prior covariance matrix satisfies $|\Sigma_{ij}| \leq \frac{1}{2K-3} \cdot \Sigma_{ii}, \forall i \neq j$.*

This condition requires that the size of the covariance between every pair of attribute values

¹⁵ $\text{Cov}(\theta_1 + 2\theta_2, \theta_1 - \theta_2) = \text{Var}(\theta_1) + \text{Cov}(\theta_1, \theta_2) - 2\text{Var}(\theta_2) = 3/5 + 1/5 - 2 \times 2/5 = 0$.

¹⁶Indeed, Lemma 5 shows that the marginal value of learning θ_i is given by γ_i^2 , where γ_i is the posterior covariance between ω and θ_i . Thus the marginal values are equal if and only if $\text{Cov}(\omega, \theta_1) = \pm \text{Cov}(\omega, \theta_2)$; that is, ω is independent of either $\theta_1 - \theta_2$ or $\theta_1 + \theta_2$.

is bounded by an expression depending on the variances.¹⁷ For the case of two attributes, we require only that the covariance Σ_{12} is smaller in magnitude than both variances Σ_{11} and Σ_{22} , which would imply our previous Assumption 2.¹⁸ In general, the condition in Assumption 3 is more restrictive for larger numbers of sources K .

To interpret the use of Assumption 3, note that prior covariances measure the *complementarity or substitution effects* across the information provided by different sources (i.e., whether information from one source increases or decreases the learning benefits from other sources). Assumption 3 limits the magnitude of such complementarity/substitution, so that the agent's short-run and long-run information acquisition incentives are aligned. In Appendix B.1, we provide a counterexample to illustrate that misalignment can occur when Assumption 3 is violated.

Under this assumption, the optimal information acquisition strategy is described as follows:

Theorem 2. *Suppose Assumption 3 is satisfied.¹⁹ Then, there exist times*

$$0 = t_0 \leq t_1 \leq \dots \leq t_{K-1} < t_K = +\infty$$

and nested sets

$$\emptyset = B_0 \subsetneq B_1 \subsetneq \dots \subsetneq B_{K-1} \subsetneq B_K = \{1, \dots, K\},$$

such that an optimal information acquisition strategy is described by a deterministic path of attention allocations $(\beta_1(t), \dots, \beta_K(t))_{t \geq 0}$. This path consists of K stages: For each $1 \leq k \leq K$, the instantaneous attention allocation 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 $t \geq t_{K-1}$ is proportional to α .

¹⁷Note that this condition requires the covariances to be not too negative, and also not too positive, which differs from the previous Assumption 2. Loosely, the difference between the $K = 2$ and $K > 2$ cases is that with $K > 2$, the relationship between any two sources (i.e., whether they are complements or substitutes) is affected by observation of other sources outside of this pair. In particular, two sources that were previously complementary can cease to be so when the agent (optimally) samples a third source, and their covariance can switch sign along the path of information acquisition. This does not happen with $K = 2$.

¹⁸However, when $K = 2$ our previous Assumption 2 is strictly weaker. So Theorem 1 does not follow as a corollary from Theorem 2 below.

¹⁹The condition we provide in Assumption 3 for general K is sufficient but not necessary for the characterization to hold. In fact, a weaker, but less interpretable, sufficient condition that we use in the proofs is the following: The inverse of the prior covariance matrix Σ^{-1} is *diagonally-dominant*; that is, $[\Sigma^{-1}]_{ii} \geq \sum_{j \neq i} |[\Sigma^{-1}]_{ij}|$ for all $1 \leq i \leq K$. Nonetheless, the constant $\frac{1}{2K-3}$ in Assumption 3 is tight, in a sense that we formalize in Appendix D.7.

The times t_k as well as the attention allocations (and their support B_k) at each stage can be determined directly from the primitives α and Σ , and are history-independent. In Appendix C, we explain how to compute these times and sets. Theorem 2 thus tells us that the agent can reduce the dynamic information acquisition problem to a sequence of K static problems, each of which involves finding the optimal constant ratio 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 or waiting cost.

We again demonstrate this result in an example:

Example 4. Suppose there are three unknown attributes, and the agent wants to learn $\omega = \theta_1 + \theta_2 + \theta_3$. The agent's prior over these attribute values is

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \end{pmatrix}, \begin{pmatrix} 4 & 0 & 0 \\ 0 & 4 & -1 \\ 0 & -1 & 3 \end{pmatrix} \right)$$

Note that this prior satisfies Assumption 2.

The optimal information acquisition strategy consists of three stages:

Stage 1. The agent initially puts all attention towards learning θ_1 . To interpret, notice that negative correlation between attributes θ_2 and θ_3 reduces the overall uncertainty about the sum $\theta_2 + \theta_3$; thus, the marginal value of learning θ_1 is initially higher than learning either θ_2 or θ_3 . The agent attends only to θ_1 until time $t_1 = \frac{1}{12}$, at which point his posterior covariance matrix becomes

$$\begin{pmatrix} 3 & 0 & 0 \\ 0 & 4 & -1 \\ 0 & -1 & 3 \end{pmatrix},$$

as given by (1). This posterior belief has the property that $\omega = \theta_1 + \theta_2 + \theta_3$ is independent of $\theta_1 - \theta_2$, so as discussed the marginal values of learning θ_1 and learning θ_2 have equalized. Since the posterior variance of θ_3 is smaller than θ_2 , the marginal value of learning θ_3 is strictly lower.

Stage 2. The agent next splits his attention between learning θ_1 and learning θ_2 in the constant proportion $(4/7, 3/7)$. These acquisitions reduce the marginal value of learning θ_1 and the marginal value of learning θ_2 *at the same rate*, thus maintaining the equality between

these marginal values. At time $t_2 = \frac{13}{44}$, the agent's posterior covariance matrix is

$$\begin{pmatrix} 11/5 & 0 & 0 \\ 0 & 44/15 & -11/15 \\ 0 & -11/15 & 44/15 \end{pmatrix}.$$

The marginal values of learning all three attributes have become the same, since at this time $\omega = \theta_1 + \theta_2 + \theta_3$ is independent of both $\theta_1 - \theta_2$ and $\theta_1 - \theta_3$.

Stage 3. From this time on, the agent acquires information evenly from each source via the constant attention allocation $(1/3, 1/3, 1/3)$.

4.3 Arbitrary Priors

Suppose the prior belief does not satisfy the assumptions given above; can we still say anything about optimal information acquisition? It turns out that under optimal sampling from *any* prior belief, the agent's posterior beliefs will eventually satisfy Assumption 3. In fact, optimal sampling is not required: Along any path in which each source receives infinite attention (which is necessary for complete learning of ω and thus satisfied under optimal sampling), the agent's beliefs will enter the set of beliefs defined by Assumption 3.

Formally, consider the cumulated attention vector $q(t)$ introduced earlier. We then have:

Lemma 1. *Starting from any prior belief, the optimal information acquisition strategy has the property that the induced cumulated attentions $q_i(t) \rightarrow \infty$ for each $1 \leq i \leq K$ as $t \rightarrow \infty$.*²⁰

Lemma 2. *Suppose $q_i(t) \rightarrow \infty$ for each $1 \leq i \leq K$. Then, the agent's posterior beliefs satisfy Assumption 3 at all sufficiently late times.*

Once Assumption 3 is met, the characterization given in Theorem 2 holds (taking the “prior” to be the posterior belief at that time). In particular, we can conclude from Lemma 1, Lemma 2 and Theorem 2 that:

Proposition 1. *Starting from any prior belief, the optimal information acquisition strategy is eventually a constant attention allocation (across all sources) proportional to the weight vector α .*

²⁰We note that starting from a general prior belief, $q_i(t)$ can be a random variable depending on past signal realizations. Thus the lemma asserts that each source receives infinite attention along every history.

Thus, in general, the optimal information acquisition strategy will eventually have the properties described earlier: independence of signal realizations, of the payoff function and of the waiting cost. The notion of “eventual” is uniform on these dimensions. Specifically, we show in the proof that there exists \bar{t} depending only on α and Σ , such that the optimal attention allocation at any time $t \geq \bar{t}$ is proportional to α .

5 Proof Outline for Theorems 1 and 2

The plan of the proof is to first define a *uniformly optimal* strategy, which minimizes the agent’s posterior variance about ω at every possible stopping time. When uniformly optimal strategies exist, they are the optimal information acquisition strategy. We then show that under the assumption on the prior belief that we provide, uniformly optimal strategies do exist, and have the structure that we characterize.

Definition of a uniformly optimal strategy. For every time t , define the *t-optimal* attention vector to be the allocation of t units of attention that minimizes posterior variance about ω .²¹

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

We will say that an attention allocation strategy is uniformly optimal if it integrates to the *t-optimal* vector at *every* t .

Definition 1. Say that an information acquisition strategy S is uniformly optimal if the induced cumulated attention vector at each time t is $n(t)$, independently of signal realizations.

This is a strong property, and existence of such a strategy is in general not guaranteed.

When a uniformly optimal strategy exists, it is optimal. By definition, if a cumulated attention vector is *t-optimal*, it implies that the agent has learned as much about ω as possible in the interval $[0, t)$. Thus, if the agent stops acquiring information at time t (and takes the optimal action), his expected flow payoff is maximized among all strategies that deterministically stop at t . The form of the payoff function u does not matter because, due to normal beliefs, achieving minimum posterior variance means that the agent’s information up to time t is Blackwell more informative than under any other strategy (Blackwell, 1951; Hansen and Torgersen, 1974).

²¹We show in Lemma 6 that this minimizer is unique.

Requiring that $q(t)$ is t -optimal at *every* time t then implies that the information acquisition strategy is most informative about ω at every history and maximizes expected payoffs given any *exogenous* stopping time. In our Gaussian environment, such a strategy also maximizes expected payoffs even when the stopping time can be endogenously chosen; this follows from a result of Greenshtein (1996) (see Lemma 7 in the appendix). It follows that whenever a uniformly optimal strategy exists, it must be the optimal strategy in our problem.²² It remains to show that under Assumption 3, a uniformly optimal strategy does exist, and has the structure described in Theorem 2.

Existence of a uniformly optimal strategy. To show that a uniformly optimal strategy exists, we make use of the following simple lemma:

Lemma 3. *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, we require that for every $t' > t$, the optimal allocation of t' units of attention devotes a higher amount of attention to *all* sources compared to the optimal allocation of t units. This is necessary and sufficient for a single information acquisition strategy to achieve the optimal cumulated attention vectors at both times.

Whether or not this condition is satisfied turns out to depend on the cross-partials of the posterior variance function V , as given by (2). When information from the different sources are complements—meaning that additional information about one attribute improves the value to additional information about another—the agent optimally chooses a *positive* mixture to take advantage of the complementarity. In contrast, if more information about attribute i *decreases* the marginal value of information about attribute j , then the agent may prefer to re-allocate attention away from attribute i towards attribute j . This can lead the optimal allocation of $t + \Delta$ units to involve less attention towards attribute i than the optimal allocation of t units. The consequence is a failure of monotonicity in the t -optimal vectors $n(t)$, precluding existence of a uniformly optimal strategy. See Appendix B.1 for such an example.

Assumptions 2 and 3 control the sizes of the cross-partials of V at the prior belief, and have the implication that at all *subsequent* beliefs along the optimal sampling path, different

²²While it is possible to write down the Bellman equation for this control problem, the value function (as a function of the current belief) is high-dimensional and difficult to solve for explicitly, especially if we do not have any structure on $u(\cdot)$ and $c(\cdot)$. Our argument based on Blackwell comparisons gets to the optimal policy (i.e., attention allocation) without going through the value function. See also Appendix A.2.

sources are complements whenever their marginal values are highest. This ensures that the agent will always acquire signals in positive mixtures,²³ so that a uniformly optimal strategy exists.

Structure of the uniformly optimal strategy. When a uniformly optimal strategy exists, the instantaneous attention allocations $\beta(t)$ are simply the time-derivatives of the t -optimal vectors $n(t)$. Under this strategy, the agent divides attention at every moment across learning those attributes that maximize the *instantaneous* marginal reduction of posterior variance V .

When there is a single attribute that maximizes reduction in V , as is the case for a generic prior belief, the agent optimally allocates all attention towards learning the corresponding attribute. As beliefs about that attribute, say attribute i , become more precise, the marginal value of learning about i decreases continuously relative to the marginal value of learning other attributes. Eventually the marginal value of learning about i will equal the marginal value of learning about some other attribute j .

At this point, there are multiple attributes that yield the same marginal value of reduction in variance. Since V is differentiable, directional derivatives can be written as a convex combination of the partial derivatives in each of the coordinate directions. Hence, all mixtures over i and j lead to the same, maximal, *instantaneous* reduction in uncertainty about the state. However, these mixtures have different implications for the marginal values of the different sources at *future* instants. For the dynamic problem, the agent thus optimally turns from the “first-order” comparison of marginal values to a “second-order” comparison of mixtures.²⁴ We demonstrate that there is a unique mixture over i and j that maintains equivalence of their marginal values, and this mixture is selected in the optimal dynamic rule. Technically, we derive the (second-order) optimal mixture by working with the Hessian matrix of V ; see Lemmata 5, 11 and 12 in the appendix.

²³The formal version of this claim is Lemma 10 in the appendix. Note that complementarity or substitution of two sources is captured by the relevant cross-partial derivative of the posterior variance function V , given in Lemma 5.

²⁴We mention that the idea of trying to maximize the marginal value of learning is known in the operations research literature as *knowledge-gradient*; see for example Frazier et al. (2008, 2009). These papers establish the asymptotic optimality of knowledge-gradient strategies when the agent seeks to select the best one out of K unknown payoffs. Although we also study a (correlated) Gaussian environment, we have a different decision problem based on a weighted sum of the unknowns, and the two settings overlap only when $K = 2$ as we discuss in Section 6. Moreover, our Theorems 1 and 2 show that knowledge gradient is *exactly* optimal in many situations. In this sense our results complement those of Frazier et al. (2008, 2009), which give general bounds on the potential loss of adopting knowledge-gradient.

The rest of the proof follows similarly: as uncertainty about attributes i and j decrease, eventually their marginal values equal those of a third attribute. At this point the agent expands his observation set to include the new source(s), and we can repeat the same reasoning. This yields the “nested-set” property in Theorems 1 and 2.

6 Application 1: Binary Choice

The framework we study relates to a large body of work regarding “binary choice tasks,” in which an agent has a choice between two goods with payoffs v_1 and v_2 , and can devote effort towards learning about these payoffs before making his decision. The leading model in this domain, the drift-diffusion model (Ratcliff and McKoon, 2008), supposes that the agent observes a Brownian motion whose drift depends on which good yields the higher payoff. In our framework, this model corresponds to a case in which the agent’s prior belief is supported on two points—either $(v_1, v_2) = (v_L, v_H)$ or $(v_1, v_2) = (v_H, v_L)$ where $v_H > v_L$ are known quantities. Thus the agent has uncertainty over which good is better, but not over how much better it is.²⁵ Fudenberg et al. (2018) recently proposed a variation on this model to allow for the latter kind of uncertainty. In their *uncertain drift-diffusion* model, the agent has a jointly normal prior over (v_1, v_2) , and has access to two Brownian motions with drifts corresponding to these unknown payoffs.

Both the classic drift-diffusion model and also Fudenberg et al. (2018) focus primarily on deriving the optimal stopping rule given exogenous information, which we do not pursue here. Fudenberg et al. (2018) additionally consider a version of their model in which the agent endogenously acquires information by choosing attention allocations (subject to an budget constraint) that scale the drifts of the two Brownian motions.²⁶ Since the payoff difference $v_1 - v_2$ is a sufficient statistic for the agent’s decision, this corresponds exactly to our framework with $K = 2$, $\theta_1 = v_1$, $\theta_2 = -v_2$, and equal payoff weights $\alpha_1 = \alpha_2 = 1$.

These authors show that if the agent’s prior is both *independent and symmetric*—that is, $\Sigma = I$ —then the agent optimally devotes equal attention to both payoffs at all times. We now show how our Theorem 1 generalizes this result in two directions: arbitrary priors (Section 6.1) and asymmetric information precision about the two payoffs (Section 6.2).

²⁵That is, the classic DDM assumes that $|v_1 - v_2|$ is known to the agent.

²⁶See Section E of Fudenberg et al. (2018).

6.1 General Prior Covariance Matrix

Suppose the agent's prior is

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \right).$$

Here, $\rho \in (-1, 1)$ captures the prior correlation between the two unknown payoffs. Recall that for the case of equal payoff weights $\alpha_1 = \alpha_2 = 1$, Theorem 1 characterizes the optimal information acquisition strategy starting from any prior. We thus obtain the following corollary:

Corollary 1. *Suppose $\sigma_1 \geq \sigma_2$. The agent's optimal information acquisition strategy $(\beta_1(t), \beta_2(t))$ in this binary choice problem consists of two stages:*

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

$$t \leq t_1^* = \frac{1/\sigma_2^2 - 1/\sigma_1^2}{1 - \rho^2},$$

the agent optimally allocates all attention to the first information source (about v_1).

- **Stage 2:** *At times $t > t_1^*$, the agent optimally allocates half of his attention to each information source.*

When $\Sigma = I$, the threshold is $t_1^* = 0$, so that the agent splits his attention evenly from the beginning. This returns Theorem 5 in Fudenberg et al. (2018). Corollary 1 demonstrates that two aspects of their characterization generalize: Starting from an arbitrary prior covariance matrix Σ , the agent will *eventually* acquire information according to the constant proportion $(\frac{1}{2}, \frac{1}{2})$. Moreover, this proportion is optimal from the beginning whenever the two unknown payoffs have the same initial uncertainty. But if the prior belief is ex-ante “asymmetric,” the agent initially devotes all attention to learning about the payoff he deems more uncertain.²⁷

Corollary 1 additionally allows us to derive new comparative statics in the prior belief.

Corollary 2. *Suppose $\sigma_1 \geq \sigma_2$. Then, holding all else equal:*

- *an increase in σ_1 results in uniformly higher attention towards source 1 (i.e., $\beta_1(t)$ is weakly smaller at every t);*

²⁷We note additionally that the Fudenberg et al. (2018) result does not characterize “off-equilibrium” attention allocations (where the agent has paid unequal attention to the two sources in the past). In contrast, our corollary above applies to all prior beliefs and thus allows for characterization of optimal information acquisition following any history, including those in which the agent has previously behaved sub-optimally.

- an increase in σ_2 results in uniformly lower attention towards source 1;
- an increase in $|\rho|$ results in uniformly higher attention towards source 1.

Since the long-run frequencies with which the sources are viewed is independent of Σ (Stage 2 in Corollary 1), changes in the prior only affect the attention strategy by changing t_1^* , the time at which the agent switches from observing source 1 to source 2. The first two comparative statics are intuitive: since $\sigma_1 > \sigma_2$, the agent initially has greater uncertainty about the first payoff. As we increase this difference in prior uncertainty—either by increasing σ_1 or decreasing σ_2 —Stage 1 increases in length and the threshold t_1^* moves later. To understand the third comparative static, note that as the degree of correlation $|\rho|$ increases in magnitude, information about the first payoff becomes more revealing about the second payoff. Thus, everything else equal, it takes longer for the agent’s uncertainty about the first payoff to “catch up” with his uncertainty about the second payoff; so t_1^* increases. All of these are new predictions enabled by our previous results, and are testable empirically.

6.2 Asymmetric Levels of Informativeness

We can alternatively enrich the Fudenberg et al. (2018) setting by allowing the informativeness of the two sources to be different, which would be the case if for example it was easier to obtain information about one of the payoffs than the other. Formally, suppose that $(\theta_1, \theta_2)' \sim \mathcal{N}((\mu_1, \mu_2)', I)$ as in Fudenberg et al. (2018), but each diffusion process X_i evolves as

$$dX_i^t = \beta_i(t)\theta_i dt + \zeta_i \sqrt{\beta_i(t)} \mathcal{B}_i^t.$$

$\zeta_i > 0$ moderates the informativeness of the process, and larger ζ_i corresponds to a more noisy source. Under this setup, a unit of attention paid to source i delivers a normal signal of the form

$$\theta_i + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \zeta_i^2).$$

The Fudenberg et al. (2018) setting is nested as $\zeta_1 = \zeta_2 = 1$.

To map this setting into our main model, we normalize the noise terms to have *unit variances* as follows: Define $\tilde{\theta}_i = \theta_i/\zeta_i$, so that each unit of attention spent on source i equivalently generates a *standard normal* signal about $\tilde{\theta}_i$. Under this transformation, the payoff-relevant state is $\zeta_1\tilde{\theta}_1 + \zeta_2\tilde{\theta}_2$, and the agent’s prior covariance matrix over $(\tilde{\theta}_1, \tilde{\theta}_2)$ is

$$\tilde{\Sigma} = \begin{pmatrix} 1/\zeta_1^2 & 0 \\ 0 & 1/\zeta_2^2 \end{pmatrix}.$$

Assumption 2 is satisfied in this transformed problem; thus the optimal attention choices $(\beta_1(t), \beta_2(t))$ are again characterized by Theorem 1.

Corollary 3. *Suppose $\zeta_1 \leq \zeta_2$. The agent's optimal information acquisition strategy $(\beta_1(t), \beta_2(t))$ in this binary choice problem consists of two stages:*

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

$$t \leq t_1^* = \zeta_1(\zeta_2 - \zeta_1),$$

the agent optimally allocates all attention to source 1.

- **Stage 2:** *At times $t > t_1^*$, the agent optimally allocates his attention in the constant fraction $\left(\frac{\zeta_1}{\zeta_1 + \zeta_2}, \frac{\zeta_2}{\zeta_1 + \zeta_2}\right)$.*

When $\zeta_1 = \zeta_2$ so that the sources are equally informative, the threshold is $t_1^* = 0$ and the mixture at Stage 2 is $(1/2, 1/2)$, again returning Theorem 5 in Fudenberg et al. (2018). But when the sources have different levels of informativeness, then the agent initially devotes all attention to learning from the more informative source, which however receives lower attention in the long run.

This corollary permits study of how changes in ζ_i , the noisiness of a source, affect the time path of attention. Recall that in the previous section we considered a similar comparative static regarding initial uncertainty (Corollary 2). Comparison of the two corollaries reveals that prior noise and signal noise affect attention allocation in different ways. In contrast to the straightforward comparative static in σ_1 reported in Corollary 2, the effect of a local increase in ζ_1 has two, potentially competing, effects: (1) it changes the length of Stage 1 (i.e., t_1^*), and (2) it also affects the long-run frequencies with which the two sources are viewed in Stage 2.

The direction of the second effect is clear: increasing the noise level ζ_1 always results in a higher long-run share of viewership for source 1. But the first effect on the length of Stage 1 can be ambiguous: making its information more noisy simultaneously reduces the marginal value of source 1, but also reduces the speed at which its marginal value shrinks to that of source 2's. From the closed-form expression for t_1^* in Corollary 3, we see that

$$\frac{\partial t_1^*}{\partial \zeta_1} \geq 0 \iff \zeta_1 \leq \zeta_2/2.$$

Thus, when ζ_1 is quite small relative to ζ_2 , it holds that $\frac{\partial t_1^*}{\partial \zeta_1} > 0$. In this regime, increasing the noisiness of source 1 makes it receive higher attention both in Stage 1 and in Stage 2. As ζ_1 increases beyond the threshold $\zeta_2/2$, further increasing this noise level leads to lower

attention paid to source 1 in Stage 1, and higher attention in Stage 2. We are not aware of prior literature that studies this effect of information precision on the *time path* of people’s information demand.

Finally, it is straightforward to consider a model incorporating both the generalizations of Section 6.1 and 6.2; we defer this analysis to Appendix F.

7 Application 2: Competing News Sources

Next, we apply our results to study information provision in a setting with strategic information providers. Specifically, we are interested in how competition affects the quality of information, when sources strategically determine the precision of the information that they provide.

To fix ideas, suppose a politician has been associated with two potential cases of misconduct in office: negligence in handling sensitive military information and use of public office to advance personal goals. The severity of each of these acts is unknown, and the public expects them to be correlated: e.g., politicians who are careless with sensitive materials are more likely to abuse power, and vice versa. Two online news sources respectively have connections with military personnel and with staff in the White House, and report on the corresponding misconduct case. These sources primarily earn revenue by running ads, so they aim to maximize time spent on their site. The choice variable is the informativeness of articles on their site, i.e., how quickly to reveal what they know.

Formally, a representative news reader seeks to learn the sum of attributes θ_1 and θ_2 , and his prior over these parameters is

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \right),$$

where $\rho \in (-1, 1)$ measures prior correlation between θ_1 and θ_2 . All of our results extend to a mass of readers sharing this common prior. We assume that the prior covariance is not too negative compared with the prior variances; specifically, we require:

Assumption 4. $\sigma_1 + \rho\sigma_2 \geq 0$ and $\sigma_2 + \rho\sigma_1 \geq 0$.

This is guaranteed if the prior is symmetric ($\sigma_1 = \sigma_2$) or positively correlated ($\rho \geq 0$).

Each of two news sources $i = 1, 2$ (freely) chooses a standard deviation ζ_i , where a unit of time spent on its site generates the signal

$$\theta_i + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \zeta_i^2).$$

Note that there is no cost for the news sources to provide more informative articles. Nonetheless, as we demonstrate below, in equilibrium the sources will choose strictly positive noise levels ζ_i .

The news reader has some underlying decision to make at a future date (e.g., whether or not to support the politician), and optimally allocates attention given ζ_1 and ζ_2 , which are fixed across time. Denote his optimal allocation at time t by $(\beta_1(t), \beta_2(t))$. Each news source i 's payoff is the discounted average attention paid to that source $\int_0^\infty r e^{-rt} \beta_i(t) dt$, where r is a (common) discount rate. We can interpret this as reduced form for advertising revenue, where each news source receives profit proportional to the amount of viewership.²⁸

For any fixed (ζ_1, ζ_2) , we can transform the reader's information acquisition problem to our main model by normalizing the signals to have unit noise variances and scaling the states θ_1, θ_2 accordingly (as we did in Section 6.2). For this transformed problem, Theorem 1 characterizes the full time path of attention. In Stage 1, the higher marginal value source receives all viewership; whereas in Stage 2, the reader mixes over both sources. If source 1 is selected in Stage 1, then its payoff is

$$U_1(\zeta_1, \zeta_2) = \int_0^{t_1^*} r e^{-rt} dt + \int_{t_1^*}^\infty r e^{-rt} \frac{\zeta_1}{\zeta_1 + \zeta_2} dt,$$

while source 2's payoffs is

$$U_2(\zeta_1, \zeta_2) = \int_{t_1^*}^\infty r e^{-rt} \frac{\zeta_2}{\zeta_1 + \zeta_2} dt,$$

where t_1^* is the switch-point as described in Theorem 1.

The key tension is between optimizing for greater long-run viewership—where larger noise ζ_i increases the long-run frequency $\frac{\zeta_i}{\zeta_i + \zeta_j}$ —versus competing to be chosen in the short-run—which encourages smaller ζ_i . Intuitively, more precise information improves the competitive value of the source at the beginning of time, but reduces the value of continual engagement with the source. This trade-off is not straightforward, as the importance of being chosen first depends on t_1^* (the length of Stage 1), which is itself endogenous to the chosen noise levels ζ_1 and ζ_2 .

The following proposition characterizes the equilibrium:

Proposition 2. *Under Assumption 4, the unique equilibrium between two competing news sources is a pure strategy equilibrium (ζ_1^*, ζ_2^*) with*

$$\zeta_1^* = \sigma_1(\sigma_1 + \rho\sigma_2)z \quad \text{and} \quad \zeta_2^* = \sigma_2(\sigma_2 + \rho\sigma_1)z,$$

²⁸Here, for the sake of illustrating the equilibrium, we are considering the case where the reader samples forever. In the politician example, this would be reasonable if the election is far away.

where

$$z = \sqrt{\frac{\sigma_1 \sigma_2 (1 - \rho^2)}{r(\sigma_1 + \rho \sigma_2)(\sigma_2 + \rho \sigma_1)(\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2)}}.$$

Given these equilibrium choices of noise levels, the reader optimally mixes over the sources in the constant fraction $\left(\frac{\zeta_1^*}{\zeta_1^* + \zeta_2^*}, \frac{\zeta_2^*}{\zeta_1^* + \zeta_2^*}\right)$ at every moment.

These expressions simplify substantially if we suppose that the prior covariance matrix is symmetric (i.e., the reader is initially equally uncertain about θ_1 and θ_2):

Corollary 4. *If $\sigma_1 = \sigma_2 = \sigma$, then the unique equilibrium is (ζ^*, ζ^*) where*

$$\zeta^* = \sigma \cdot \sqrt{\frac{1 - \rho}{2r}}.$$

Our first observation is that in equilibrium, there is no “Stage 1” of information gathering: the reader immediately begins mixing in a constant fraction over the sources. This is *despite* the possibility of initial asymmetry in how well each attribute is understood. Thus, Proposition 2 reveals that in equilibrium, sources choose noise levels that exactly offset this prior asymmetry, equalizing their marginal values from the beginning.

Asymmetry in σ_i does, however, impact how the reader mixes over the sources and the profits that the sources receive, as we discuss in the subsequent corollary.²⁹

Corollary 5 (Division of Attention). *Equilibrium attention paid to source 1, $\frac{\zeta_1^*}{\zeta_1^* + \zeta_2^*}$*

- (a) *exceeds equilibrium attention paid to source 2 if and only if $\sigma_1 \geq \sigma_2$;*
- (b) *is increasing in σ_1 and decreasing in σ_2 ;*
- (c) *is decreasing in ρ if $\sigma_1 \geq \sigma_2$ and increasing in ρ if $\sigma_1 \leq \sigma_2$;*
- (d) *is independent of r .*

Part (a) says that the source providing information about the less-understood attribute receives more attention at every moment in time (and thus also receives higher profit). Intuitively, greater initial uncertainty increases the marginal value of learning from the corresponding source, giving this source a competitive advantage. But the result is more subtle than it seems, since this asymmetry is in the prior belief only. From Corollary 1, we know that if the sources were to provide *equally* informative signals ($\zeta_1 = \zeta_2$), then the reader

²⁹For this set of comparative statics, we assume that changes in σ_1, σ_2 and ρ maintain Assumption 4, so that Proposition 2 continues to hold.

would eventually mix equally between these sources regardless of the prior. What Part (a) says, then, is that the initial advantage conferred to a source is turned into a persistent advantage in this strategic setting: this source can afford to provide noisier information, and can thus capture more attention at every moment. The greater this initial asymmetry, the larger the persistent advantage, as described in Part (b) of the corollary.

Part (c) says that attention paid to the more frequented source is decreasing in the correlation ρ . Thus, when attributes are positively correlated—as in our example, where negligence increases the probability of corruption and vice versa—attention is more equal across the sources, and the strategic advantage conferred to the source with the more uncertain attribute is lower. In contrast, when attributes are negatively correlated—so that a higher level on one attribute implies a lower value on the other—then initial asymmetries are exaggerated in equilibrium, and the amount of attention paid to the sources becomes more unequal. To the best of our knowledge, this relationship between the direction of correlation, and how competitive the setting is, has not been noted.³⁰ Finally, Part (d) says that the share of attention is independent of the sources’ common discount rate.³¹

Our final result in this section is about the overall quality of information, and how this depends on the primitives. In equilibrium, the sampling procedure described in Proposition 2 leads to posterior variances about ω approximately given by $\frac{(\zeta_1^* + \zeta_2^*)^2}{t}$ at large times t .³² Thus, the sum of standard deviations $\zeta_1^* + \zeta_2^*$ is an appropriate measure of aggregate noise in equilibrium.

Corollary 6 (Informativeness of News). *Equilibrium aggregate noise level, $\zeta_1^* + \zeta_2^*$*

(a) *is decreasing in the discount rate r ;*

(b) *is increasing in the prior variance σ_1 if $\sigma_1 \geq \sigma_2$ (and otherwise increasing in σ_2);*

(c) *is decreasing in the prior correlation ρ .*

³⁰More formally, the noise terms ζ_1 and ζ_2 must make equal $Cov(\frac{\theta_1}{\zeta_1}, \omega) = \sigma_1^2/\zeta_1 + \rho\sigma_1\sigma_2/\zeta_1$ and $Cov(\frac{\theta_2}{\zeta_2}, \omega) = \sigma_2^2/\zeta_2 + \rho\sigma_1\sigma_2/\zeta_2$. Large positive ρ implies a smaller asymmetry in ζ_1 and ζ_2 , while large negative ρ implies a greater asymmetry.

³¹If the discount rate differed across sources, then this asymmetry would matter for the equilibrium share of attention.

³²To see this, recall the transformation $\tilde{\theta}_i = \theta_i/\zeta_i$ (described in Section 6.2), which maps this game with endogenous noise variances to our main model with unit variances. Under this transformation, the sources provide standard normal signals about $\tilde{\theta}_1$ and $\tilde{\theta}_2$, and the payoff-relevant state ω can be rewritten as $\zeta_1\tilde{\theta}_1 + \zeta_2\tilde{\theta}_2$. The asymptotic approximation of posterior variances then follows from Claim 1 in Liang and Mu (2020).

Part (a) says that the more patient the information providers are (larger r), the less precise their signals will be in equilibrium. This is because less patient information providers compete over short-run profits (i.e., being chosen in Stage 1), and thus prefer precise signals, while patient providers compete for long-run profits (i.e., long-run frequency), and thus prefer imprecise signals. Part (b) and Part (c) can be similarly explained by this trade-off between short-run and long-run profits. Recall that the length of Stage 1 is determined by how long it takes for the marginal value of information about the less-understood attribute to “catch up” with the other. When prior uncertainty about the less-understood attribute is larger, so is the initial difference in the two sources’ marginal values. Thus Stage 1 becomes shorter, and sources prioritize Stage 2 by providing noisier information, as described in Part (b) of the corollary.

To understand the effect of correlation on equilibrium informativeness, note that when the reader initially learns about the less-understood attribute 1, he also learns about the other attribute 2 due to correlation. Thus the marginal values of *both* information sources decrease. As ρ increases, the two sources become closer to substitutes, and the marginal value of source 2 decreases *faster*. Hence it takes longer for the marginal value of source 1 to equalize the marginal value of source 2, implying a longer Stage 1. The sources thus have stronger incentives to be chosen first, and they provide more precise information in equilibrium.

These results contribute to a large literature regarding how competition affects information provision ([Gentzkow and Shapiro, 2006](#)). Here we are specifically interested in how providers with *imperfectly correlated* information interact strategically. Imperfect correlation means that some information is common across the sources, but that each source also has a monopoly on a residual component of the unknown that readers would like to know. This places our setting intermediate between monopolists—who can fully extract rents by noising up information—and perfect competition—where firms compete away rents by providing precise information. What our analysis in this section reveals is that information providers compete for readers in the short-run and exploit readers in the long-run. Thus, a crucial force determining the equilibrium quality of information is how information providers trade off between these two time periods. We find that news is of higher quality when information providers are *less forward-looking*, when the information they provide is *more positively correlated*, and when *prior uncertainty is lower*, as each of these increases the relative importance of the short-run competition.

Finally, in Appendix I we generalize these insights to a game where $K > 2$ news sources compete, and where readers seek to learn $\theta_1 + \dots + \theta_K$. Assuming a symmetric prior over

these attributes, we can generalize Corollary 4 to show that $\zeta^* = \sigma \cdot \sqrt{\frac{1-\rho}{Kr}}$ is the equilibrium noise level in the unique symmetric pure strategy equilibrium. Note that for this problem, the transformed prior covariance matrix $\tilde{\Sigma}$ does *not* in general satisfy Assumption 3.³³ Nonetheless, we are able to directly compute the uniformly optimal strategy (defined in Section 5) and show that our K -stage characterization of optimal information acquisition extends to this setting.

8 Discrete Time

Although our main model is in continuous time, our results have immediate implications for a related discrete-time model (previously described in Remark 1): Suppose time is discrete with period length Δ , and at each period $t = 0, \Delta, 2\Delta, \dots$, the agent allocates a budget of Δ precision across the K information sources. Choice of precision levels $\pi_1(t), \dots, \pi_K(t)$ (subject to $\sum \pi_i(t) \leq \Delta$) produces independent observation of the signals

$$Y_i = \theta_i + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}\left(0, \frac{1}{\pi_i(t)}\right).$$

At each period t , the agent also decides whether or not to stop and take an action. Our main results can be directly mapped into statements for this setting.

Corollary 7. *Suppose Assumption 3 holds. Then at each period $t = 0, \Delta, 2\Delta, \dots$, the optimal mixture over signals is $(\pi_1(t), \dots, \pi_K(t))$ where*

$$\pi_i(t) = \int_t^{t+\Delta} \beta_i(t) dt,$$

with $\beta_i(t)$ being the optimal attention allocation for the continuous-time model that is described in Theorem 2.

In a companion piece, Liang et al. (2017), we discretize not only time but also information acquisitions: at each period t , the agent has to choose one of the K signals with precision Δ , without the ability to mix. The necessity of an integer approximation complicates characterization of the full sequence of signal choices. Nevertheless, we provide conditions under which myopic acquisition is optimal or eventually optimal.

³³This can happen if the endogenously chosen noise levels ζ_1, \dots, ζ_K are very different.

9 Conclusion

Information acquisition is a classic problem within economics, but there are relatively few *dynamic* models that are tractable and admit explicit characterizations. In this paper we present a class of dynamic information acquisition problems whose solution can be explicitly characterized in closed-form. It turns out 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 is not too strong. In return, we can allow for a great deal of generality in other aspects of the problem, such as arbitrary payoff functions and/or many patterns of correlation and asymmetry across the unknown variables. In the present paper, we show how our characterization of optimal information acquisition can be used to derive new results in two example economic settings. We believe that the tractability of the solution opens the door to additional interesting applications.

Appendix

A Preliminaries

A.1 Posterior Variance Function

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

Lemma 4. *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

The next lemma calculates the first and second derivatives of the posterior variance function V :

Lemma 5. *Given a cumulated attention vector $q \geq \mathbf{0}$, define*

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

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

Proof. By Lemma 5, 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 $q \geq \mathbf{0}$. So V is convex. \square

These technical properties are used to show that for each t , the t -optimal vector $n(t)$ is unique:

Lemma 6. *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 4), $\lim_{\lambda \rightarrow \infty} V(|q^\lambda|) = \lim_{\lambda \rightarrow \infty} V(q^\lambda) = V(r)$. 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 5, $\partial_i V(r) = -\gamma_i^2$ for $1 \leq i \leq K$. Thus $\gamma_1 = \dots = \gamma_l = 0$. Since γ is not the zero vector,³⁴ 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

³⁴This follows because α is not the zero vector, by assumption.

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 7. *A uniformly optimal strategy is dynamically optimal regardless of the payoff function $u(\cdot)$ or the waiting cost function $c(\cdot)$.*

Proof. This is essentially a continuous-time version of Theorem 3.1 in Greenshtein (1996), which establishes a Blackwell ordering over *sequential* experiments for dynamic decision problems. In our environment with normal signals about an one-dimensional unknown (our payoff-relevant state ω), this theorem implies that a sequence of signals Blackwell-dominates another if and only if the former sequence leads to uniformly lower posterior variances. While the general result of Greenshtein (1996) covers decision problems in which the agent takes multiple actions, a simpler proof suffices for the class of stopping problems considered in this paper. The argument follows the proof of Theorem 5 in Fudenberg et al. (2018), with some modifications. For completeness we reproduce this proof below, using our notation.

Fix any attention strategy S and denote by $\mathbb{E}^S[\cdot]$ the associated expectation operator, and by $\mathbb{E}^{S^*}[\cdot]$ the expectation operator associated with the uniformly optimal strategy S^* . The optimal stopping rule τ (under S) is a solution to

$$\sup_{\tau} \mathbb{E}^S[\max_a \mathbb{E}[u(a, \omega) \mid \mathcal{F}_{\tau}] - c(\tau)]. \quad (3)$$

By the Dambis–Dubins–Schwartz Theorem (see for example Theorem 1.6 in Chapter V of Revuz and Yor (1999)), there exists a Brownian motion $(B_{\nu})_{\nu \in [0, v_0]}$ such that

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

where v_0 denotes the prior variance of ω , and the random variable v_t is the posterior variance at time t under strategy S . This change of variables is a time change where the new scale is the posterior variance.

For each $v \in (0, v_0]$, define the stochastic process $\phi_v := \inf\{t : v_t \leq v\}$. If the agent stops with posterior variance v , his posterior expectation of ω is the value of $B_{v_0 - v}$. Denote by $U(\cdot, \cdot)$ his maximum expected payoff when taking the optimal action given this belief, where the arguments are the expected value and variance of ω . Then by (3), the value of the agent can be rewritten as³⁵

$$\sup_v \mathbb{E}[U(B_{v_0 - v}, v) - c(\phi_v)].$$

³⁵This generalizes Fudenberg et al. (2018), where the U function is simply its first argument (in the special case of binary choice).

As the posterior variance v_t is greater than the minimum posterior variance v_t^* under S^* at all times t , we have that

$$\phi_v \geq \phi_v^* := \inf\{t : v_t^* \leq v\} \quad \forall v.$$

Consequently, the value under strategy S is smaller than the value under S^* :

$$\begin{aligned} \sup_{\tau} \mathbb{E}^S[\max_a \mathbb{E}[u(a, \omega) \mid \mathcal{F}_{\tau}]] &= \sup_v \mathbb{E}[U(B_{v_0-v}, v) - c(\phi_v)] \\ &\leq \sup_v \mathbb{E}[U(B_{v_0-v}, v) - c(\phi_v^*)] \\ &= \sup_{\tau} \mathbb{E}^{S^*}[\max_a \mathbb{E}[u(a, \omega) \mid \mathcal{F}_{\tau}]]. \end{aligned} \tag{4}$$

□

We also have a simple converse result:

Lemma 8. *Fixing Σ , α and the payoff function $u(\cdot)$. Suppose an information acquisition strategy is optimal for all cost functions $c(\cdot)$, then it is uniformly optimal.*

Proof. Take an arbitrary time t and consider the cost function with $c(\tau) = 0$ 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 cost function, the induced cumulated attention vector must achieve t -optimality. Varying t yields the result. □

A.3 Uniqueness of Optimal Information Acquisition

By Lemma 7, whenever a uniformly optimal strategy exists, it is the optimal information strategy regardless of the form of $u(\cdot)$ and $c(\cdot)$. As we show in later appendices, Assumptions 2 and 3 guarantee existence. The results in Theorems 1 and 2 thus characterize the uniformly optimal strategy.

Without further assumptions on u and c , there could exist other optimal information acquisition strategies. For example, consider the cost function $c(\cdot)$ used in the proof of Lemma 8. Under this cost function, the agent always stops at some fixed time τ . Hence any strategy that achieves the τ -optimal vector $n(\tau)$ gives the same, maximal amount of information about ω at the stopping time. All such strategies are optimal for this problem, and we cannot identify the attention allocation at any particular instant before τ . Uniform optimality, in particular t -optimality for $t < \tau$, is not necessary for optimal information acquisition here.

Nonetheless, such counterexamples are non-generic. A careful inspection of the proof of Lemma 7 suggests that whenever $c(\tau)$ is *strictly* increasing in τ , an attention allocation

strategy S does as well as the uniformly optimal strategy S^* if and only if the following holds:

For every $v > 0$ such that the agent stops with positive density at posterior variance v under S , the posterior variances under S decrease to v at the same time as under S^ .*

That is, we require $\phi_v = \phi_v^*$ whenever the posterior variance v is realized under the stopping rule.

We now introduce an assumption on the agent's stopping rule:

Assumption 5. *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 measure of continuation histories such that the agent optimally stops in the interval $(t, t']$.*

To see how this condition implies $S = S^*$ up to the stopping time, let us suppose for contradiction that after some history, the strategy S deviates from uniform optimality. Then, along this history, the posterior variances under S in the interval $(t, t']$ are strictly *larger* than under S^* (for some t' slightly bigger than t). By assumption, the agent stops in this interval with positive probability. Thus we can take any posterior variance v achieved in this interval, and deduce that v is reached *slower* under S than under S^* . As discussed above, this is sufficient to show that S performs strictly worse than S^* .

In summary, we have the following result:

Proposition 3. *Suppose the waiting cost $c(\cdot)$ is strictly increasing, and Assumption 5 is satisfied. Then, any optimal information acquisition strategy coincides with the uniformly optimal strategy at every history where the agent has not stopped.*

We note that although Assumption 5 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 in turn guaranteed if extreme values of ω agree on the optimal action, and if the marginal cost of waiting is bounded away from zero. These conditions on the primitives are rather weak, and are satisfied in most natural applications of the model (e.g., binary choice with constant marginal waiting cost).

B Proof of Theorem 1

Define cov_1, cov_2 as in the statement of Theorem 1:

$$cov_1 = \alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{21}; \quad cov_2 = \alpha_1 \Sigma_{12} + \alpha_2 \Sigma_{22}.$$

Further define $x_i = \alpha_i \det(\Sigma)$ to ease notation.

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

$$\begin{aligned}
\gamma &:= (\Sigma^{-1} + Q)^{-1} \cdot \alpha \\
&= (\Sigma^{-1} \cdot (I + \Sigma Q))^{-1} \cdot \alpha \\
&= (I + \Sigma Q)^{-1} \cdot \Sigma \cdot \alpha \\
&= (I + \Sigma Q)^{-1} \cdot \begin{pmatrix} \text{cov}_1 \\ \text{cov}_2 \end{pmatrix} \\
&= \frac{1}{\det(I + \Sigma Q)} \begin{pmatrix} 1 + q_2 \Sigma_{22} & -q_2 \Sigma_{12} \\ -q_1 \Sigma_{21} & 1 + q_1 \Sigma_{11} \end{pmatrix} \cdot \begin{pmatrix} \text{cov}_1 \\ \text{cov}_2 \end{pmatrix} \\
&= \frac{1}{\det(I + \Sigma Q)} \begin{pmatrix} x_1 q_2 + \text{cov}_1 \\ x_2 q_1 + \text{cov}_2 \end{pmatrix}.
\end{aligned}$$

By Lemma 5, this implies the marginal values of the two sources are given by:

$$\begin{aligned}
\partial_1 V(q_1, q_2) &= \frac{-(x_1 q_2 + \text{cov}_1)^2}{\det^2(I + \Sigma Q)}, \\
\partial_2 V(q_1, q_2) &= \frac{-(x_2 q_1 + \text{cov}_2)^2}{\det^2(I + \Sigma Q)}.
\end{aligned} \tag{5}$$

Note that Assumption 2 translates into $\text{cov}_1 + \text{cov}_2 \geq 0$. 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 $\text{cov}_1 \geq \text{cov}_2$, then cov_1 is non-negative. Let $t_1^* = \frac{\text{cov}_1 - \text{cov}_2}{x_2}$. Then when $q_1 + q_2 \leq t_1^*$ we always have

$$x_1 q_2 + \text{cov}_1 \geq \text{cov}_1 \geq x_2 q_1 + \text{cov}_2,$$

since $x_1 q_2 \geq 0$ and $x_2 q_1 \leq x_2(q_1 + q_2) \leq x_2 t_1^* = \text{cov}_1 - \text{cov}_2$. We also have

$$x_1 q_2 + \text{cov}_1 \geq -(x_2 q_1 + \text{cov}_2),$$

since $x_1 q_2, x_2 q_1 \geq 0$ and by assumption $\text{cov}_1 + \text{cov}_2 \geq 0$. Thus, (5) 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 function V . That is, $n(t) = (t, 0)$ for $t \leq t_1^*$.

For $t > t_1^*$, observe that (5) 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 (5) and the budget constraint $n_1(t) + n_2(t) = t$, yields the solution

$$n(t) = \left(\frac{x_1 t + \text{cov}_1 - \text{cov}_2}{x_1 + x_2}, \frac{x_2 t - \text{cov}_1 + \text{cov}_2}{x_1 + x_2} \right).$$

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.

B.1 Counterexample

The following example illustrates how and why Theorem 1 might fail:

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 cumulated attention vectors $n(t)$ (see Section 5) 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 $t \leq 1/4$, 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 $\omega = \theta_1 + 4\theta_2$ is independent of $\theta_1 + \theta_2$ (see Footnote 16).³⁶ However, to

³⁶The key difference between this counterexample and Example 3 is that here ω is independent of the sum $\theta_1 + \theta_2$, rather than the difference $\theta_1 - \theta_2$. Although both cases imply equal marginal values, it turns out that independence between ω and $\theta_1 - \theta_2$ is necessary for $n(t)$ to be monotonic. To this end, Assumptions 2 and 3 essentially rule out the other possibility that ω is independent of $\theta_1 + \theta_2$.

maintain equal marginal values at future instants, it is actually optimal to take attention away from attribute 1 and re-distribute 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 cumulated attention toward attribute 1 is decreasing in this interval.³⁷

This failure of monotonicity occurs because at $t = 1/4$, the two sources of information strongly substitute one another—by Lemma 5 in the appendix, the cross-partial $\partial_{12}V = 96/343 > 0$, suggesting that the marginal value of either source (as measured by *reduction* in the posterior variance V) is *lower* after having learned from the other source. Consequently, there does not exist a uniformly optimal strategy in this example (Lemma 3). 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 (Lemma 8).

B.2 Necessity of Assumption 2

We show here that 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 4. *Suppose Assumption 2 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 characterize the t -optimal attention vector $n(t)$:

1. If $t \leq \frac{-(cov_1+cov_2)}{x_2}$, then $x_1 q_2 + cov_1$ is negative and has larger absolute value than $x_2 q_1 + cov_2$ (which is positive) whenever $q_1 + q_2 = t$. By (5), this means $\partial_1 V(q_1, q_2) \leq \partial_2 V(q_1, q_2)$, and so $n(t) = (t, 0)$. In words, with a very small budget, it is optimal to devote all attention to source 1.
2. If $\frac{-(cov_1+cov_2)}{x_2} < t < \frac{-(cov_1+cov_2)}{x_1}$, then $\partial_1 V(0, t) < \partial_2 V(0, t)$ 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 $n(t) = (\frac{-x_1 t - cov_1 - cov_2}{x_2 - x_1}, \frac{x_2 t + cov_1 + cov_2}{x_2 - x_1})$.

³⁷Subsequently, at times $t \in (1, 3]$, the t -optimal vector is $n(t) = (0, t)$, allocating all attention to attribute 2. Finally, at times $t \geq 3$, $n(t) = (\frac{t-3}{5}, \frac{4t+3}{5})$, allocating attention proportional to α .

3. If $\frac{-(cov_1+cov_2)}{x_1} \leq t \leq \frac{cov_2-cov_1}{x_1}$, then $(x_2q_1+cov_2)^2 - (x_1q_2+cov_1)^2 = (cov_2-cov_1-x_1q_2+x_2q_1) \cdot (cov_1+cov_2+x_1q_2+x_2q_1) \geq 0$ whenever $q_1+q_2 = t$. Thus $\partial_1 V(q_1, q_2) \geq \partial_2 V(q_1, q_2)$, implying that the t -optimal attention vector should be $n(t) = (0, t)$.
4. Finally, if $t > \frac{cov_2-cov_1}{x_1}$, then it holds that $\partial_1 V(0, t) < \partial_2(0, t)$ and $\partial_1 V(t, 0) > \partial_2(t, 0)$. So $n(t)$ is interior and satisfies the first-order condition

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

since both terms are now positive. This together with $n_1(t) + n_2(t) = t$ yields the solution $n(t) = (\frac{x_1 t + cov_1 - cov_2}{x_1 + x_2}, \frac{x_2 t - cov_1 + cov_2}{x_1 + x_2})$ and completes the analysis.

Note that in Case 2 above, as t increases in the range, $n_1(t)$ *actually decreases*. This proves that a uniformly optimal strategy does not exist. \square

C An Algorithm for Finding the Optimal Information Acquisition Strategy when $K > 2$

The next appendix provides a detailed proof of Theorem 2. Here we give an outline and show how the times t_k and sets B_k defined in Theorem 2 can be found recursively. 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 are the sources whose marginal reduction of posterior variance is highest (see Lemma 5).

2. **(Computation of the constant attention allocation in stage k .)** If $|B_k| > k$ then stage k is degenerate, and we proceed to stage $k+1$ with $Q_k = Q_{k-1}$. Otherwise we can re-order the attributes so that the k attributes in B_k are the first k 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 $k \times k$ top-left submatrix of Σ and Σ_{TR} to be the $k \times (K-k)$ top-right block. Finally let

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

be a $k \times 1$ vector. The agent's optimal attention allocation in stage k is proportional to α^k ; that is,

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

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 .³⁸ 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 $Q_k = Q^k(t_k)$.

D Proof of Theorem 2

D.1 Weaker Assumption

Given Lemma 7, it is sufficient to show that the t -optimal vector $n(t)$ is weakly increasing in t , and that its time-derivative is locally constant as described in the theorem. We will in fact prove the same result under the following weaker assumption:

Assumption 6. *The inverse of the prior covariance matrix Σ^{-1} is diagonally-dominant. That is,*

$$[\Sigma^{-1}]_{ii} \geq \sum_{j \neq i} |[\Sigma^{-1}]_{ij}| \quad \forall 1 \leq i \leq K.$$

³⁸This smallest time can be computed as follows. For each $j > k$, consider the following (polynomial) equation in t :

$$(e'_j \cdot (\Sigma^{-1} + Q^k(t))^{-1} \cdot \alpha)^2 = (e'_1 \cdot (\Sigma^{-1} + Q^k(t))^{-1} \cdot \alpha)^2.$$

Any solution $t > t_{k-1}$ is a time at which source j would have the same marginal value as sources $1, \dots, k$. Such a solution t necessarily exists, since at $t = t_{k-1}$ the LHS is smaller by assumption, while at $t = \infty$ the LHS is bigger as the RHS is 0.

Let $s(j)$ be the smallest solution to the above equation, for each fixed $j > k$. Then $t_k := \min_{j > k} s(j)$ is the earliest time after t_{k-1} such that the sources having the greatest marginal value are a strict superset of the first k sources.

This is implied by Assumption 3 via the following lemma.

Lemma 9. *Suppose the prior covariance matrix Σ satisfies Assumption 3, 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

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

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

D.2 Technical Property of γ

The following technical lemma will be repeatedly used.

Lemma 10. *Suppose Σ^{-1} is diagonally-dominant. Given an arbitrary attention vector q , define γ as in Lemma 5 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 $-\gamma_i = |\gamma_j|$ 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 $i \in B$. The result that these γ_i are the same follows from the definition that their absolute values are maximal. \square

D.3 The Last Stage

To prove Theorem 2 under Assumption 6, 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 11. *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*

$$n(t) = n(\underline{t}) + \frac{t}{\alpha_1 + \dots + \alpha_K} \cdot \alpha. \quad ^{39}$$

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 cumulated 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 5. Applying Lemma 10, 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

D.4 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

³⁹That is, $n_i(t) = n_i(\underline{t}) + \frac{t}{\alpha_1 + \dots + \alpha_K} \cdot \alpha_i$ for each i .

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 state the lemma under a slightly stronger assumption that Σ^{-1} is *strictly* diagonally-dominant. Later we will discuss how the lemma should be modified without this strictness.

Lemma 12. *Suppose Σ^{-1} is strictly diagonally-dominant. Choose any time \underline{t} and denote*

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

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 $\operatorname{argmin}_i \partial_i V(n(\bar{t}))$ is a strict superset of B . When $|B| = K$, it holds that $\bar{t} = \infty$ and β is proportional to α , as given by Lemma 11.

Proof. 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 10, 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 11. Specifically, given the prior covariance matrix Σ , we can choose another basis of the attributes $\theta_1, \dots, \theta_k, \hat{\theta}_{k+1}, \dots, \hat{\theta}_K$ with two properties:

1. each $\hat{\theta}_j$ ($j > k$) is a linear combination of the original attributes $\theta_1, \theta_2, \dots, \theta_K$;
2. $\operatorname{Cov}[\theta_i, \hat{\theta}_j] = 0$ 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 $\hat{\theta}$ the vector $(\hat{\theta}_{k+1}, \dots, \hat{\theta}_K)'$. The payoff-relevant state $\omega = \alpha' \cdot \theta$ can thus be rewritten as $\tilde{\alpha}' \cdot \tilde{\theta} + \hat{\alpha}' \cdot \hat{\theta}$ for some constant coefficient vectors $\tilde{\alpha} \in \mathbb{R}^k$ and $\hat{\alpha} \in \mathbb{R}^{K-k}$. Using property 2 above, we can solve for $\tilde{\alpha}$ from Σ, α and B :

$$\tilde{\alpha} = (\Sigma_{TL})^{-1} \cdot (\Sigma_{TL}, \Sigma_{TR}) \cdot \alpha \tag{6}$$

where Σ_{TL} represents the $k \times k$ top-left submatrix 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 $\hat{\theta}$ 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 about ω is simply the variance about $\tilde{\alpha}' \cdot \tilde{\theta}$ plus the variance about $\hat{\alpha}' \cdot \hat{\theta}$. 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 about $\tilde{\alpha}' \cdot \tilde{\theta}$.

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 11: 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.⁴⁰ Moreover, at $q = n(\underline{t})$, the definition of the set B implies that these k partial derivatives are smaller (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 cumulated 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 desired.

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 submatrix of Q , and

$$\tilde{\gamma} = ((\Sigma_{TL})^{-1} + \tilde{Q})^{-1} \cdot \tilde{\alpha}. \quad (7)$$

We will show that $\tilde{\gamma}$ is just the first k coordinates of γ . Indeed, observe that $((\Sigma_{TL})^{-1} + \tilde{Q})^{-1}$ is also the $k \times k$ top-left submatrix of $(\Sigma^{-1} + Q)^{-1}$.⁴¹ Using (6) and (7), we have

$$\begin{aligned} \tilde{\gamma} &= [(\Sigma^{-1} + Q)^{-1}]_{TL} \cdot (\Sigma_{TL})^{-1} \cdot (\Sigma_{TL}, \Sigma_{TR}) \cdot \alpha \\ &= [(\Sigma^{-1} + Q)^{-1}]_{TL} \cdot (\alpha_1, \dots, \alpha_k)' + [(\Sigma^{-1} + Q)^{-1}]_{TL} \cdot (\Sigma_{TL})^{-1} \cdot \Sigma_{TR} \cdot (\alpha_{k+1}, \dots, \alpha_K)'. \end{aligned}$$

On the other hand, from $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$ we have

$$\begin{aligned} (\gamma_1, \dots, \gamma_k)' &= ([(\Sigma^{-1} + Q)^{-1}]_{TL}, [(\Sigma^{-1} + Q)^{-1}]_{TR}) \cdot \alpha \\ &= [(\Sigma^{-1} + Q)^{-1}]_{TL} \cdot (\alpha_1, \dots, \alpha_k)' + [(\Sigma^{-1} + Q)^{-1}]_{TR} \cdot (\alpha_{k+1}, \dots, \alpha_K)'. \end{aligned}$$

Comparing the above two formulas, $\tilde{\gamma}$ is the first k coordinates of γ so long as

$$[(\Sigma^{-1} + Q)^{-1}]_{TL} \cdot (\Sigma_{TL})^{-1} \cdot \Sigma_{TR} = [(\Sigma^{-1} + Q)^{-1}]_{TR},$$

⁴⁰To be rigorous, the conclusion should be about the function $\tilde{V}(q_1, \dots, q_k)$, which is the posterior variance about $\tilde{\alpha}'\tilde{\theta}$ in the smaller problem. But as discussed, this differs from $V(q_1, \dots, q_k, 0, \dots, 0)$ by a constant.

⁴¹This holds because $(\Sigma^{-1} + Q)^{-1} = Q^{-1} - Q^{-1}(Q^{-1} + \Sigma)^{-1}Q^{-1}$. Note that Q^{-1} is a block matrix: its $k \times k$ top-left block is \tilde{Q}^{-1} , and its $k \times (K - k)$ top-right block is zeros (its bottom-right block can be seen as the diagonal matrix with infinities). So the top-left block of $Q^{-1} - Q^{-1}(Q^{-1} + \Sigma)^{-1}Q^{-1}$ is simply $\tilde{Q}^{-1} - \tilde{Q}^{-1}[(Q^{-1} + \Sigma)^{-1}]_{TL}\tilde{Q}^{-1}$, which in turn is equal to $\tilde{Q}^{-1} - \tilde{Q}^{-1}(\tilde{Q}^{-1} + \Sigma_{TL})^{-1}\tilde{Q}^{-1} = ((\Sigma_{TL})^{-1} + \tilde{Q})^{-1}$.

which indeed holds.⁴²

Hence $\tilde{\gamma}_i = \gamma_i$ for $1 \leq i \leq k$, and it is the same positive number by Lemma 10. Rewriting (7) 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 . A theorem of Carlson and Markham (1979) says that if Σ^{-1} is (strictly) diagonally-dominant, then so is $(\Sigma_{TL})^{-1}$ for any principal submatrix Σ_{TL} . Consequently the row sums of $(\Sigma_{TL})^{-1}$ are all strictly positive, implying that $\tilde{\alpha}_i > 0$. \square

D.5 Completing the Proof

We now apply Lemma 12 repeatedly to prove Theorem 2. Continuing to assume strict diagonal dominance, we can apply Lemma 12 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 $B^1 = \operatorname{argmin}_{1 \leq i \leq K} \partial_i V(\mathbf{0})$. Applying Lemma 12 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. Since the sets $\emptyset = B^0, B^1, B^2, \dots$ are nested by construction, we eventually have $B^m = \{1, \dots, K\}$ for some m , and consequently $t^m = \infty$.

Note that $B^{l+1} - B^l$ need not be a singleton for each l (i.e., two sources can simultaneously become new minimizers of $\partial_i V$). Thus m can be smaller than K , and the nested sets B^1, \dots, B^m and increasing times t^1, \dots, t^m do not necessarily satisfy the conclusion of Theorem 2. However, this is easy to resolve by including “redundant” times. Formally, we set $t_k = t^l$ for any k satisfying $|B^l| \leq k < |B^{l+1}|$. We also choose B_1, \dots, B_K such that $B_{k+1} - B_k$ is a singleton for each k , and $B_k = B^l$ whenever $k = |B^l|$. The nested sets B_1, \dots, B_K and weakly increasing times t_1, \dots, t_K then satisfy the conclusions of Theorem 2. This completes the characterization under the assumption that Σ^{-1} is *strictly* diagonally-dominant.

⁴²Consider the identity $(\Sigma^{-1} + Q)^{-1} \cdot (\Sigma^{-1} + Q) = I_K$. The top-right block of the product is zeros, so by block matrix multiplication we have

$$[(\Sigma^{-1} + Q)^{-1}]_{TL} \cdot (\Sigma^{-1} + Q)_{TR} = -[(\Sigma^{-1} + Q)^{-1}]_{TR} \cdot (\Sigma^{-1} + Q)_{BR}.$$

Next consider the identity $\Sigma \cdot (\Sigma^{-1} + Q) = I_K + \Sigma(Q)$. The top-right block is again zeros, and we similarly deduce

$$\Sigma_{TL} \cdot (\Sigma^{-1} + Q)_{TR} = -\Sigma_{TR} \cdot (\Sigma^{-1} + Q)_{BR}.$$

These two equalities together yield the desired result.

D.6 Weak Diagonal Dominance and Zero Weights

Here we demonstrate how to prove Theorem 1 assuming only that Σ^{-1} is weakly diagonally-dominant. The difficulty that arises with this change is that in the proof of Lemma 12, we cannot conclude that the optimal attention allocation 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 12 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 assume without loss 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 (6).

We want to show that if $\hat{\alpha}$ is supported on the same set of sources as $\tilde{\alpha}$ —i.e., more sources maximize the marginal value, but the observation set is unchanged—then $\hat{\alpha}$ in fact coincides with $\tilde{\alpha}$ on their support. Indeed, by definition of $\hat{\alpha}$ (going back to the proof of Lemma 12) we can write

$$\omega = \sum_{i=1}^{\hat{k}} \hat{\alpha}_i \theta_i + \text{residual term orthogonal to } \theta_1, \dots, \theta_{\hat{k}}.$$

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

$$\omega = \sum_{i=1}^{\tilde{k}} \hat{\alpha}_i \theta_i + \text{residual term orthogonal to } \theta_1, \dots, \theta_{\tilde{k}},$$

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.

We mention that our proof of Theorem 2 (and Theorem 1) extends without change to cases where some payoff weights are zero, rather than strictly positive. In fact, because any source with zero weight receives no attention in the long run, it *never* receives any attention under the optimal strategy in environments where our characterization applies.⁴³ Thus these sources can be simply dropped from the model without affecting our results.

D.7 Tightness of $\frac{1}{2K-3}$

Finally, we provide an example to show that the constant $\frac{1}{2K-3}$ in Assumption 3 is *tight for the existence of a uniformly optimal strategy*.

Proposition 5. *For any $\rho > \frac{1}{2K-3}$, there exists a prior covariance matrix Σ satisfying $|\Sigma_{ij}| \leq \rho \cdot \Sigma_{ii}$ for all $i \neq j$, as well as some weight vector α , such that uniform optimality is unachievable.*

Proof. Let Σ have diagonal entries 1 and off-diagonal entries $-\rho$, with $\rho > \frac{1}{2K-3}$. We also choose $\alpha_2 = \dots = \alpha_K = 1$, and α_1 equal to a small positive number.

For this problem, we will show that the t -optimal vector $n(t)$ is *not* monotonic over time, which implies the result via Lemma 3. Note that the last $K - 1$ sources have symmetric prior and symmetric payoff weights. Thus, the posterior variance function $V(q_1, q_2, \dots, q_K)$ is symmetric in its last $K - 1$ arguments. This implies that the t -optimal vector $n(t)$ must satisfy $n_2(t) = \dots = n_K(t)$; otherwise its permutations would have the same posterior variance, contradicting uniqueness of $n(t)$.

Minimizing the posterior variance at time t thus simplifies to the following problem:

$$(n_1, n_2) \in \underset{q_1, q_2 \geq 0, q_1 + (K-1)q_2 = t}{\operatorname{argmin}} V(q_1, q_2, \dots, q_2).$$

That is, the agent optimally divides attention between source 1 and the remaining sources, which always receive equal attention.

The posterior belief of such an agent can be derived by Bayesian updating on the following K normal signals: $\theta_1 + \mathcal{N}\left(0, \frac{1}{q_1}\right)$ and $\theta_i + \mathcal{N}\left(0, \frac{1}{q_2}\right)$ for $2 \leq i \leq K$. We can show that in

⁴³It may receive finite attention when our assumptions are violated.

terms of predicting the payoff-relevant state $\alpha_1\theta_1 + \sum_{i>1}\theta_i$, the agent's belief is the same as if he had observed just *two* signals: $\theta_1 + \mathcal{N}\left(0, \frac{1}{q_1}\right)$, and $\frac{1}{K-1} \sum_{i>1}\theta_i + \mathcal{N}\left(0, \frac{1}{(K-1)q_2}\right)$.⁴⁴

Given this equivalence, we can relate t -optimal vectors in this problem with K sources to t -optimal vectors in a smaller problem with just two sources. Specifically, define $\theta_1^* = \theta_1$, $\theta_2^* = \frac{1}{K-1} \sum_{i>1}\theta_i$, $\alpha_1^* = \alpha_1$, $\alpha_2^* = K-1$. Then the payoff-relevant state can be rewritten as

$$\omega = \alpha_1^* \cdot \theta_1^* + \alpha_2^* \cdot \theta_2^*.$$

The discussion in the preceding paragraph shows that the posterior variance function V^* in this $K=2$ problem satisfies

$$V^*(q_1, q_2) = V\left(q_1, \frac{q_2}{K-1}, \dots, \frac{q_2}{K-1}\right),$$

because on both sides the posterior variance is derived assuming that the agent had observed the two signals $\theta_1 + \mathcal{N}\left(0, \frac{1}{q_1}\right)$ and $\frac{1}{K-1} \sum_{i>1}\theta_i + \mathcal{N}\left(0, \frac{1}{(K-1)q_2}\right)$. Hence, t -optimality in this smaller problem is equivalent to t -optimality in the original problem.

In this smaller problem, the prior covariance matrix Σ^* about (θ_1^*, θ_2^*) is

$$\Sigma^* = \begin{pmatrix} 1 & -\rho \\ -\rho & \frac{1-(K-2)\rho}{K-1} \end{pmatrix}.$$

In particular, since $\rho > \frac{1}{2K-3}$, $\Sigma_{21}^* + \Sigma_{22}^*$ is negative. Thus if $\alpha_1^* = \alpha_1$ is sufficiently small, this $K=2$ problem violates Assumption 2. By Proposition 4, we conclude that the t -optimal cumulated attention vectors are not monotonic over time. The same holds for the original problem, completing the proof. \square

E Proof of Proposition 1

E.1 Proof Outline

As discussed in the main text, we only need to prove that each source receives infinite attention (Lemma 1) and that Theorem 2 applies at any posterior belief after each source is

⁴⁴This can be proved by directly computing the posterior covariance matrix. Alternatively, note that the signal $\frac{1}{K-1} \sum_{i>1}\theta_i + \mathcal{N}\left(0, \frac{1}{(K-1)q_2}\right)$ is the *average* of the $K-1$ signals $\theta_i + \mathcal{N}\left(0, \frac{1}{q_2}\right)$ for $2 \leq i \leq K$ considered initially, so it contains weakly less information. However, it is sufficient for learning $\omega = \alpha_1\theta_1 + \sum_{i>1}\theta_i$ for the following reason: 1) it is sufficient for learning $\sum_{i>1}\theta_i$, and 2) conditional on this sum, the original $K-1$ signals $\theta_i + \mathcal{N}\left(0, \frac{1}{q_2}\right)$ only provide information about the differences $\theta_i - \theta_j$ (with $i, j > 1$). These differences are independent from θ_1 conditional on $\sum_{i>1}\theta_i$ (they are in fact independent from both), so the extra information does not affect the belief about θ_1 conditional on $\sum_{i>1}\theta_i$.

sufficiently sampled. The latter is easy: Observe that the agent’s posterior *precision* matrix is given by $\Sigma^{-1} + Q$, where Q is the diagonal matrix with entries q_1, \dots, q_K . As $q_i \rightarrow \infty$ to each i , clearly the matrix $\Sigma^{-1} + Q$ is diagonally-dominant. So the conclusion of Theorem 2 holds.⁴⁵

It remains to prove Lemma 1. This is in turn implied by the following lemma:

Lemma 13. *Fix Σ and α . Given any $\underline{q} \in \mathbb{R}_+$, there exists $\bar{q} \in \mathbb{R}_+$ such that the cumulated attention vectors $q(t)$ under the optimal strategy have the following property: Whenever $q_i(t) < \underline{q}$ for some source i , it holds that $q_j(t) \leq \bar{q}$ for every source j .*

Taking the contrapositive, this result says that whenever a source j has received attention more than \bar{q} , then *each* source i has received attention at least \underline{q} . Since there necessarily exists such a source j as $t \rightarrow \infty$, the consequence is that all sources must eventually receive cumulated attention $\geq \underline{q}$. This lemma thus implies Lemma 1.

We now sketch how we prove the above lemma. First it is clear that the result for any \underline{q} follows from the result for any larger \underline{q} . So we will assume \underline{q} is large (to be formalized later). We will then prove the result by choosing \bar{q} even larger (also determined later). Suppose for contradiction that after some history, the cumulated attention vector satisfies $q_i(t_0) < \underline{q}$ and $q_j(t_0) > \bar{q}$. By relabeling the signals, we can assume that

$$q_1(t_0), \dots, q_k(t_0) < \underline{q} \leq q_{k+1}(t_0), \dots, q_{K-1}(t_0); \quad q_K(t_0) > \bar{q}.$$

That is, the cumulated attention devoted to each of the first k sources is “deficient,” whereas source K has received “excessive” attention. We can further assume that source K continues to receive positive attention in some interval $(t_0, t_0 + \epsilon]$; otherwise we can replace t_0 by an earlier time without changing these conditions.

Our proof method will be to construct a profitable deviation strategy (of how to allocation attention) following this history, so that optimality is violated. Thanks to the main theorem of Greenshtein (1996), any deviation strategy is profitable so long as it *decreases the posterior variance about ω at all future times*. Given a deviation strategy, let $\tilde{q}(t)$ denote the induced cumulated attention vector, which is distinguished from $q(t)$. Then the deviation is profitable whenever the following inequality holds:⁴⁶

$$V(\tilde{q}(t)) \leq V(q(t)), \quad \forall t \geq t_0.$$

⁴⁵This argument shows that Assumption 6 is satisfied when each q_i is large. It can be shown that in fact, the stronger Assumption 3 is also satisfied if we take q_i even larger (i.e. Lemma 2 holds).

⁴⁶Such a deviation is strictly profitable if in addition $V(\tilde{q}(t)) < V(q(t))$ holds strictly for $t \in (t_0, t_0 + \epsilon]$, which is verified below.

E.2 The Deviation

We now construct such a deviation. Take any time $T \geq t_0$, there are three cases:

- (a) Suppose that the original strategy S devotes positive attention to source K at time T . Then under the deviation strategy, the agent *diverts this attention (evenly) toward those sources i with $\tilde{q}_i(T) < \underline{q}$* .⁴⁷ If no such source exists, the deviation strategy devotes the same amount of attention to source K .
- (b) Suppose that the original strategy devotes attention to some source in $k+1, \dots, K-1$. Then the deviation strategy devotes the same attention to this source.
- (c) Suppose that the original strategy devotes attention to source $i \leq k$. If $\tilde{q}_i(T) < \underline{q}$ or $\tilde{q}_i(T) = q_i(t)$, then the deviation strategy also observes source i . Otherwise we have $\tilde{q}_i(T) = \underline{q} > q_i(T)$, and in this case the deviation strategy *diverts this amount of attention to source K instead*.

To interpret, the deviation strategy starts to deviate at time t_0 , when some source K has been observed too often compared to some other sources $1, \dots, k$. Following that history, the deviation refrains from observing source K and instead devotes attention to sources $1, \dots, k$, until all of these “deficient” sources are no longer deficient, after which the deviation strategy agrees with the original strategy in the amount of attention allocated to source i .

E.3 Four Kinds of Sources

Our end goal is to show that at any time $T \geq t_0$, either $\tilde{q}(T) = q(T)$, or $V(\tilde{q}(T)) < V(q(T))$. This will show that the deviation is profitable. But to do that, we first provide a categorization of the different sources and their cumulated attention vectors (under the deviation strategy versus the original strategy).

1. For sources $i \in I_1 \subset \{1, \dots, k\}$, we have $q_i < \tilde{q}_i < \underline{q}$ (henceforth we fix T and use q_i to denote $q_i(T)$). By construction, these sources have received equal attention diverted from source K , under the deviation strategy. So for some $x > 0$ it holds that

$$\tilde{q}_i = q_i + x, \quad \forall i \in I_1.$$

⁴⁷Formally, when the time derivative of $q_K(T)$ is positive, we set the time derivative of $\tilde{q}_K(T)$ to be zero, and compensate it by increasing the time derivatives of $\tilde{q}_i(T)$ for those signals i insufficiently observed.

2. For sources $i \in I_2 \subset \{1, \dots, k\}$, we have $q_i < \tilde{q}_i = \underline{q}$. These are the sources that have reached the target level \underline{q} under the deviation strategy, but not under the original strategy. Let x_i denote the difference $\tilde{q}_i - q_i$, then by construction we have $x_i \leq x$, which is defined above.
3. For sources $i \in I_3$, we have $q_i = \tilde{q}_i \geq \underline{q}$. These include the sources $k+1, \dots, K-1$, which the deviation strategy does not affect. Also included are those sources in $1, \dots, k$ that have reached cumulated attention \underline{q} under both the original and deviation strategies.
4. Finally source K is the only source with $q_i > \tilde{q}_i$. In fact we have

$$q_K - \tilde{q}_K = \sum_{i < K} (\tilde{q}_i - q_i) = |I_1| \cdot x + \sum_{i \in I_2} x_i.$$

Suppose $\tilde{q} \neq q$, then either I_1 or I_2 is non-empty. We will use this characterization to show $V(\tilde{q}) < V(q)$.

E.4 Comparison of Posterior Variances

The following technical lemma is needed, and we prove it at the end:

Lemma 14. *There exists a positive constant C_H depending only on Σ and α , such that for all $q_1, \dots, q_K \geq 0$,*

$$\partial_i V(q) \geq \frac{-C_H}{q_i^2}, \quad \forall 1 \leq i \leq K.$$

Moreover, there exists another positive constant C_L such that the following holds when \underline{q} is large:

If $q_1, \dots, q_K \geq \underline{q}$, then

$$\partial_i V(q) \leq \frac{-C_L}{q_i^2}, \quad \forall 1 \leq i \leq K.$$

And if some $q_i < \underline{q}$, then there exists j such that

$$q_j < \underline{q} \quad \text{and} \quad \partial_j V(q) \leq \frac{-C_L}{\underline{q}^2}.$$

To prove $V(\tilde{q}) < V(q)$, first consider the case that I_1 (defined in the previous subsection) is the empty set. Let $j \in I_2$ be the source that maximizes $x_j = \tilde{q}_j - q_j$. We then have

$$V(\tilde{q}) = V(\tilde{q}_j, \tilde{q}_{-j}) \leq V(q_j, \tilde{q}_{-j}) + (\tilde{q}_j - q_j) \cdot \partial_j V(\tilde{q}) \leq V(q_j, \tilde{q}_{-j}) - \frac{x_j \cdot C_L}{\underline{q}^2} \leq V(q_1, \dots, q_{K-1}, \tilde{q}_K) - \frac{x_j \cdot C_L}{\underline{q}^2}. \quad (8)$$

The first inequality uses the convexity of V . The second inequality uses the second part of Lemma 14 (which applies because $\tilde{q}_i \geq \underline{q}$ for all i when I_1 is empty), as well as $\tilde{q}_j = \underline{q}$ (since $j \in I_2$). The last inequality uses the monotonicity of V and $\tilde{q}_i \geq q_i$ for all but the last source.

On the other hand, we also have

$$V(q) \geq V(q_1, \dots, q_{K-1}, \tilde{q}_K) + (q_K - \tilde{q}_K) \cdot \partial_K V(q_1, \dots, q_{K-1}, \tilde{q}_K) \geq V(q_1, \dots, q_{K-1}, \tilde{q}_K) - \frac{(K-1)x_j \cdot C_H}{(\tilde{q}_K)^2}, \quad (9)$$

where the first inequality is by convexity, and the second uses the first part of Lemma 14 and $q_K - \tilde{q}_K = \sum_{i \in I_2} x_i \leq (K-1)x_j$ by our choice of j .

Recall that $\tilde{q}_K \geq \bar{q}$. Thus whenever \bar{q} is much larger compared to \underline{q} , the above inequalities (8) and (9) imply that $V(\tilde{q}) < V(q)$, as we desire to show.

Next we consider the case where I_1 is non-empty. By the third part of Lemma 14, we can choose $j \in I_1$ such that $\partial_j V(\tilde{q}) \leq \frac{-C_L}{\underline{q}^2}$. Then, similar to (8) we have

$$V(\tilde{q}) \leq V(q_1, \dots, q_{K-1}, \tilde{q}_K) - \frac{x \cdot C_L}{\underline{q}^2},$$

with x replacing the role of x_j . Likewise, we have the following analogue of (9):

$$V(q) \geq V(q_1, \dots, q_{K-1}, \tilde{q}_K) - \frac{(K-1)x \cdot C_H}{(\tilde{q}_K)^2},$$

where we used $q_K - \tilde{q}_K = |I_1| \cdot x + \sum_{i \in I_2} x_i \leq (K-1)x$.

Hence we are once again able to deduce $V(\tilde{q}) < V(q)$ so long as $\tilde{q}_K \geq \bar{q}$ is much larger than \underline{q} . This completes the proof of Proposition 1 modulo Lemma 14.

E.5 Proof of Lemma 14

In light of Lemma 5, the key will be to estimate the size of the different coordinates of $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$.

For the first part, note that the matrix norm of the posterior covariance matrix $(\Sigma^{-1} + Q)^{-1}$ is bounded above (by the norm of the prior covariance matrix Σ). Thus for any possible q , the vector γ is bounded. We now write

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

Comparing the i -th coordinate on both sides, we have $\alpha_i = e'_i \cdot \Sigma^{-1} \cdot \gamma + q_i \gamma_i$. This then implies that the product $q_i \gamma_i$ is bounded across different possible q . Since $\partial_i V(q) = -\gamma_i^2$, the first part of Lemma 14 is proved.

For the second part, we use the matrix identity

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

So $\gamma_i = e'_i \cdot (\Sigma^{-1} + Q)^{-1} \cdot \alpha = \frac{\alpha_i}{q_i} - \frac{1}{q_i} \cdot e'_i \cdot (\Sigma + Q^{-1})^{-1} \cdot Q^{-1} \cdot \alpha$. If q_1, \dots, q_K are all large, then the term being subtracted is at most $\frac{\alpha_i}{2q_i}$, because the matrix norm of $(\Sigma + Q^{-1})^{-1}$ is bounded above and the norm of Q^{-1} is small. Thus $\gamma_i \geq \frac{\alpha_i}{2q_i}$, implying that $\partial_i V \leq \frac{-\alpha_i^2}{4q_i^2}$. The second part of the lemma holds for $C_L = \min_i \frac{\alpha_i^2}{4}$.

For the third part, let $q_1, \dots, q_m < \underline{q} \leq q_{m+1}, \dots, q_K$. Suppose for the sake of contradiction that $\partial_i V(q) > \frac{-C_L}{q^2}$ for each $1 \leq i \leq m$, with C_L defined above. Then $|\gamma_i| < \frac{\alpha_i}{2\underline{q}} < \frac{\alpha_i}{2q_i}$ for $1 \leq i \leq m$. Thus, $\alpha_i - q_i \gamma_i > \frac{\alpha_i}{2}$. We now rewrite $\alpha = (\Sigma^{-1} + Q) \cdot \gamma$ as

$$\Sigma \cdot (\alpha - Q\gamma) = \gamma.$$

Since the i -th coordinate of $\alpha - Q\gamma$ is simply $\alpha_i - q_i \gamma_i$, we deduce that the vector norm of $\alpha - Q\gamma$ is bounded away from zero. So the above identity suggests that the norm of γ is also bounded away from zero. However, for $1 \leq i \leq m$ we have $|\gamma_i| < \frac{\alpha_i}{2\underline{q}}$ by hypothesis, and for $i > m$ we know from the first part that $|\gamma_i| \leq \frac{\sqrt{C_H}}{q_i} \leq \frac{\sqrt{C_H}}{\underline{q}}$. Hence the norm of γ is in fact close to zero when \underline{q} is large. This leads to a contradiction and completes the proof.

F Supplementary Material to Section 6

We consider here a binary choice problem that generalizes both Section 6.1 and 6.2. Suppose the agent's prior belief is

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \right),$$

and the observed diffusion processes X_i evolve as

$$dX_i^t = \beta_i(t)\theta_i dt + \zeta_i \sqrt{\beta_i(t)} \mathcal{B}_i^t.$$

As in Section 6.2, $\zeta_i > 0$ represents the noise level of source i .

Using the same transformation as in Section 6.2, we have that the payoff-relevant state is $\zeta_1 \tilde{\theta}_1 + \zeta_2 \tilde{\theta}_2$, and the agent's prior covariance matrix over $(\tilde{\theta}_1, \tilde{\theta}_2)$ is

$$\tilde{\Sigma} = \begin{pmatrix} \frac{\sigma_1^2}{\zeta_1^2} & \frac{\rho\sigma_1\sigma_2}{\zeta_1\zeta_2} \\ \frac{\rho\sigma_1\sigma_2}{\zeta_1\zeta_2} & \frac{\sigma_2^2}{\zeta_2^2} \end{pmatrix}.$$

Assumption 2 for this transformed problem requires

$$\frac{\sigma_1(\sigma_1 + \rho\sigma_2)}{\zeta_1} + \frac{\sigma_2(\sigma_2 + \rho\sigma_1)}{\zeta_2} \geq 0,$$

which is guaranteed if $\sigma_1 = \sigma_2$ or $\rho \geq 0$ or $\zeta_1 = \zeta_2$.

If the above inequality holds, then the characterization in Theorem 1 applies to this problem, and we obtain:

Corollary 9. *Suppose*

$$\frac{\sigma_1(\sigma_1 + \rho\sigma_2)}{\zeta_1} \geq \left| \frac{\sigma_2(\sigma_2 + \rho\sigma_1)}{\zeta_2} \right|.$$

The agent's optimal information acquisition strategy $(\beta_1(t), \beta_2(t))$ in the binary choice problem consists of two stages:

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

$$t \leq t_1^* = \frac{\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_1\zeta_2 - \sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1^2}{\sigma_1^2\sigma_2^2(1 - \rho^2)},$$

the agent optimally allocates all attention to source 1.

- **Stage 2:** *At times $t > t_1^*$, the agent optimally allocates his attention in the constant fraction $\left(\frac{\zeta_1}{\zeta_1 + \zeta_2}, \frac{\zeta_2}{\zeta_1 + \zeta_2}\right)$.*

We calculate t_1^* according to the definition in Theorem 1, as follows:

$$\begin{aligned} t_1^* &= \frac{\zeta_1\tilde{\Sigma}_{11} + \zeta_2\tilde{\Sigma}_{12} - \zeta_1\tilde{\Sigma}_{21} - \zeta_2\tilde{\Sigma}_{22}}{\zeta_2 \cdot \det(\tilde{\Sigma})} \\ &= \left(\frac{\sigma_1(\sigma_1 + \rho\sigma_2)}{\zeta_1} - \frac{\sigma_2(\sigma_2 + \rho\sigma_1)}{\zeta_2} \right) \bigg/ \left(\zeta_2 \cdot \frac{\sigma_1^2\sigma_2^2(1 - \rho^2)}{\zeta_1^2\zeta_2^2} \right) \\ &= \frac{\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_1\zeta_2 - \sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1^2}{\sigma_1^2\sigma_2^2(1 - \rho^2)}. \end{aligned}$$

G Proof of Proposition 2

Once ζ_1, ζ_2 are chosen, we can follow the analysis in Appendix F to transform the problem into our main model. Given the assumption that $\sigma_1 + \rho\sigma_2$ and $\sigma_2 + \rho\sigma_1$ are both positive, Assumption 2 is satisfied. Thus the reader's optimal attention allocation is characterized by

Corollary 9. In particular, if $t_1^* \geq 0 \geq t_2^*$, then in equilibrium source 1 is chosen exclusively until time t_1^* , after which the reader mixes in the fraction $(\frac{\zeta_1}{\zeta_1 + \zeta_2}, \frac{\zeta_2}{\zeta_1 + \zeta_2})$. Source 1's payoff is

$$\begin{aligned} U_1(\zeta_1, \zeta_2) &= \int_0^{t_1^*} r e^{-rt} dt + \int_{t_1^*}^{\infty} r e^{-rt} \frac{\zeta_1}{\zeta_1 + \zeta_2} dt \\ &= 1 - e^{-rt_1^*} \cdot \frac{\zeta_2}{\zeta_1 + \zeta_2}, \end{aligned}$$

while source 2's payoff is

$$U_2(\zeta_1, \zeta_2) = \int_{t_1^*}^{\infty} r e^{-rt} \frac{\zeta_2}{\zeta_1 + \zeta_2} dt = e^{-rt_1^*} \cdot \frac{\zeta_2}{\zeta_1 + \zeta_2}.$$

To derive a candidate equilibrium, we set $\frac{\partial U_1(\zeta_1, \zeta_2)}{\partial \zeta_1}$ and $\frac{\partial U_2(\zeta_1, \zeta_2)}{\partial \zeta_2}$ to zero and solve for ζ_1, ζ_2 . Specifically, from Corollary 9 we have

$$t_1^* = \frac{\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_1\zeta_2 - \sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1^2}{\sigma_1^2\sigma_2^2(1 - \rho^2)}.$$

It follows that

$$\frac{\partial t_1^*}{\partial \zeta_1} = \frac{\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_2 - 2\sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1}{\sigma_1^2\sigma_2^2(1 - \rho^2)}; \quad (10)$$

$$\frac{\partial t_1^*}{\partial \zeta_2} = \frac{\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_1}{\sigma_1^2\sigma_2^2(1 - \rho^2)}. \quad (11)$$

We then have

$$\begin{aligned} \frac{\partial U_1(\zeta_1, \zeta_2)}{\partial \zeta_1} &= r e^{-rt_1^*} \cdot \frac{\partial t_1^*}{\partial \zeta_1} \cdot \frac{\zeta_2}{\zeta_1 + \zeta_2} - e^{-rt_1^*} \cdot \frac{-\zeta_2}{(\zeta_1 + \zeta_2)^2} \\ &= e^{-rt_1^*} \cdot \frac{\zeta_2}{(\zeta_1 + \zeta_2)^2} \cdot \left(r \cdot \frac{\partial t_1^*}{\partial \zeta_1} \cdot (\zeta_1 + \zeta_2) + 1 \right). \end{aligned} \quad (12)$$

So $\frac{\partial U_1(\zeta_1, \zeta_2)}{\partial \zeta_1} = 0$ if and only if

$$r \cdot \frac{\partial t_1^*}{\partial \zeta_1} \cdot (\zeta_1 + \zeta_2) = -1.$$

Substituting in the expression for $\frac{\partial t_1^*}{\partial \zeta_1}$ from (10), this implies

$$-r \cdot [\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_2 - 2\sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1] \cdot (\zeta_1 + \zeta_2) = \sigma_1^2\sigma_2^2(1 - \rho^2). \quad (13)$$

Similarly we have

$$\begin{aligned} \frac{\partial U_2(\zeta_1, \zeta_2)}{\partial \zeta_2} &= -r e^{-rt_1^*} \cdot \frac{\partial t_1^*}{\partial \zeta_2} \cdot \frac{\zeta_2}{\zeta_1 + \zeta_2} + e^{-rt_1^*} \cdot \frac{\zeta_1}{(\zeta_1 + \zeta_2)^2} \\ &= e^{-rt_1^*} \cdot \frac{1}{(\zeta_1 + \zeta_2)^2} \left(-r \cdot \frac{\partial t_1^*}{\partial \zeta_2} \cdot \zeta_2(\zeta_1 + \zeta_2) + \zeta_1 \right). \end{aligned} \quad (14)$$

So $\frac{\partial U_2(\zeta_1, \zeta_2)}{\partial \zeta_2} = 0$ if and only if

$$r \cdot \frac{\partial t_1^*}{\partial \zeta_2} \cdot \zeta_2(\zeta_1 + \zeta_2) = \zeta_1.$$

Substituting in (11), this implies

$$r \cdot \sigma_1(\sigma_1 + \rho\sigma_2) \cdot \zeta_2(\zeta_1 + \zeta_2) = \sigma_1^2\sigma_2^2(1 - \rho^2). \quad (15)$$

Comparing this with (13), we see that the RHS are equal, so the LHS should also be equal. Simplifying, we obtain

$$\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_2 = \sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1.$$

Hence $t_1^* = 0$ in the candidate equilibrium. Additionally, the equilibrium noise levels ζ_1^* and ζ_2^* are related via

$$\begin{aligned} \zeta_1^* &= \sigma_1(\sigma_1 + \rho\sigma_2)z \\ \zeta_2^* &= \sigma_2(\sigma_2 + \rho\sigma_1)z \end{aligned}$$

for some $z \geq 0$. Plugging these expressions into (15) we have that

$$z = \sqrt{\frac{\sigma_1\sigma_2(1 - \rho^2)}{r(\sigma_1 + \rho\sigma_2)(\sigma_2 + \rho\sigma_1)(\sigma_1^2 + 2\rho\sigma_1\sigma_2 + \sigma_2^2)}}.$$

Next, we will show this pair of ζ_1^*, ζ_2^* is an equilibrium. Since the formulae are symmetric, we only check source 1 does not have an incentive to deviate to some $\zeta_1 \neq \zeta_1^*$. First consider $\zeta_1 < \zeta_1^*$. In this case source 1 is deviating to more precise information, which makes it the source listened to in Stage 1. The change in profit $\frac{\partial U_1(\zeta_1, \zeta_2^*)}{\partial \zeta_1}$ is still given by (12). In particular, this derivative has the same sign as $r \cdot \frac{\partial t_1^*}{\partial \zeta_1} \cdot (\zeta_1 + \zeta_2^*) + 1$, which in turn has the same sign as

$$r \cdot [\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_2^* - 2\sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1] \cdot (\zeta_1 + \zeta_2^*) + \sigma_1^2\sigma_2^2(1 - \rho^2)$$

The way we solved for ζ_1^*, ζ_2^* ensures that the above expression equals 0 when $\zeta_1 = \zeta_1^*$. As ζ_1 decreases, the term $\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_2^* - 2\sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1$ in the brackets becomes less negative (or positive). Since the other factor $(\zeta_1 + \zeta_2^*)$ becomes smaller, the overall product $[\sigma_1(\sigma_1 + \rho\sigma_2)\zeta_2^* - 2\sigma_2(\sigma_2 + \rho\sigma_1)\zeta_1] \cdot (\zeta_1 + \zeta_2^*)$ also becomes less negative (or positive). We thus know that for all $\zeta_1 < \zeta_1^*$, $\frac{\partial U_1}{\partial \zeta_1} > 0$. Hence source 1 strictly prefers ζ_1^* to any smaller ζ_1 .

If instead the deviation is to some bigger ζ_1 , then the consequence is that source 2 is now listened to in Stage 1. In this case source 1's payoff is *not* given by the above calculations. Rather, it is

$$\tilde{U}_1(\zeta_1, \zeta_2) = e^{-rt_2^*} \cdot \frac{\zeta_1}{\zeta_1 + \zeta_2}.$$

We can show that $\frac{\partial \tilde{U}_1(\zeta_1, \zeta_2^*)}{\partial \zeta_1}$ has the same sign as $-r \cdot \frac{\partial t_2^*}{\partial \zeta_1} \cdot \zeta_1(\zeta_1 + \zeta_2^*) + \zeta_2^*$; this is similar to (14), except with subscripts flipped. Further plugging in the expression for t_2^* , we obtain that $\frac{\partial \tilde{U}_1(\zeta_1, \zeta_2^*)}{\partial \zeta_1}$ has the same sign as

$$-r \cdot \sigma_2(\sigma_2 + \rho\sigma_1) \cdot \zeta_1(\zeta_1 + \zeta_2^*) + \sigma_1^2\sigma_2^2(1 - \rho^2).$$

Again, by construction the above expression equals 0 when $\zeta_1 = \zeta_1^*$. As ζ_1 increases, this expression becomes negative and so $\frac{\partial \tilde{U}_1(\zeta_1, \zeta_2^*)}{\partial \zeta_1} < 0$. Therefore source 1 also has no incentive to deviate to higher noise.

This argument shows that (ζ_1^*, ζ_2^*) is in fact a strict equilibrium, in the sense that ζ_i^* is the unique best response of source i to ζ_j^* . Since the game has *constant sum* of 1 (as total attention is 1 at every moment), we conclude that (ζ_1^*, ζ_2^*) is the unique equilibrium, pure or mixed.

H Other Proofs for Section 7

H.1 Proof of Corollary 5

From Proposition 2, we have

$$\frac{\zeta_1^*}{\zeta_2^*} = \frac{\sigma_1(\sigma_1 + \rho\sigma_2)}{\sigma_2(\sigma_2 + \rho\sigma_1)},$$

which is independent of r . This proves Part (d). Subtracting 1 from both sides, we have

$$\frac{\zeta_1^*}{\zeta_2^*} - 1 = \frac{\sigma_1^2 - \sigma_2^2}{\sigma_2(\sigma_2 + \rho\sigma_1)}.$$

The RHS is positive precisely when $\sigma_1 \geq \sigma_2$. Thus we deduce that $\zeta_1^* \geq \zeta_2^*$ if and only if $\sigma_1 \geq \sigma_2$, as claimed in Part (a).

Moreover, if $\sigma_1 \geq \sigma_2$, then as ρ increases the denominator $\sigma_2(\sigma_2 + \rho\sigma_1)$ increases, which implies that the fraction $\frac{\sigma_1^2 - \sigma_2^2}{\sigma_2(\sigma_2 + \rho\sigma_1)}$ (on the RHS of the above display) decreases. Thus $\frac{\zeta_1^*}{\zeta_2^*}$ decreases, and so does $\frac{\zeta_1^*}{\zeta_1^* + \zeta_2^*}$. Conversely, if $\sigma_1 \leq \sigma_2$, then an increase in ρ leads to an increase in source 1's equilibrium attention $\frac{\zeta_1^*}{\zeta_1^* + \zeta_2^*}$. This proves Part (c).

Lastly we prove Part (b). It suffices to show that $\frac{\zeta_1^*}{\zeta_2^*} = \frac{\sigma_1(\sigma_1 + \rho\sigma_2)}{\sigma_2(\sigma_2 + \rho\sigma_1)}$ is increasing in σ_1 . Once we do this, then by symmetry $\frac{\zeta_2^*}{\zeta_1^*}$ is increasing in σ_2 , so that $\frac{\zeta_1^*}{\zeta_2^*}$ is decreasing in σ_2 . We have

$$\frac{\partial \left(\frac{\sigma_1(\sigma_1 + \rho\sigma_2)}{\sigma_2(\sigma_2 + \rho\sigma_1)} \right)}{\partial \sigma_1} = \frac{\rho\sigma_1^2 + 2\sigma_1\sigma_2 + \rho\sigma_2^2}{\sigma_2(\sigma_2 + \rho\sigma_1)^2}$$

The numerator is positive because it can be written as the sum of $\sigma_1(\sigma_2 + \rho\sigma_1)$ and $\sigma_2(\sigma_1 + \rho\sigma_2)$, both of which are positive. This proves Part (b).

H.2 Proof of Corollary 6

From Proposition 2, we compute that

$$\zeta_1^* + \zeta_2^* = \sqrt{\frac{\sigma_1 \sigma_2 (\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2)(1 - \rho^2)}{r(\sigma_1 + \rho \sigma_2)(\sigma_2 + \rho \sigma_1)}},$$

which immediately implies Part (a).

To prove Part (b), suppose $\sigma_1 \geq \sigma_2$ and we will show that the RHS above is increasing in σ_1 . Ignoring constants, we simply need to show the following expression is increasing in σ_1 :

$$\frac{\sigma_1(\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2)}{(\sigma_1 + \rho \sigma_2)(\sigma_2 + \rho \sigma_1)}.$$

Using $\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2 = \sigma_1(\sigma_1 + \rho \sigma_2) + \sigma_2(\sigma_2 + \rho \sigma_1)$, this expression simplifies to

$$\frac{\sigma_1^2}{\sigma_2 + \rho \sigma_1} + \frac{\sigma_1 \sigma_2}{\sigma_1 + \rho \sigma_2}.$$

Its derivative with respect to σ_1 can be computed as

$$\frac{\rho \sigma_1^2 + 2\sigma_1 \sigma_2}{(\sigma_2 + \rho \sigma_1)^2} + \frac{\rho \sigma_2^2}{(\sigma_1 + \rho \sigma_2)^2}.$$

If $\rho \geq 0$, then both terms are positive and we are done. So suppose $\rho < 0$, in which case the second term has a *negative numerator*. Note also that the second denominator $(\sigma_1 + \rho \sigma_2)^2$ is bigger than the first denominator $(\sigma_2 + \rho \sigma_1)^2$ since $\sigma_1 \geq \sigma_2$. We can thus replace the second denominator with the first, and obtain the following inequality:

$$\frac{\rho \sigma_1^2 + 2\sigma_1 \sigma_2}{(\sigma_2 + \rho \sigma_1)^2} + \frac{\rho \sigma_2^2}{(\sigma_1 + \rho \sigma_2)^2} \geq \frac{\rho \sigma_1^2 + 2\sigma_1 \sigma_2}{(\sigma_2 + \rho \sigma_1)^2} + \frac{\rho \sigma_2^2}{(\sigma_2 + \rho \sigma_1)^2} = \frac{\rho \sigma_1^2 + 2\sigma_1 \sigma_2 + \rho \sigma_2^2}{(\sigma_2 + \rho \sigma_1)^2}.$$

The numerator is positive since $\rho \sigma_1^2 + 2\sigma_1 \sigma_2 + \rho \sigma_2^2 = \sigma_1(\sigma_2 + \rho \sigma_1) + \sigma_2(\sigma_1 + \rho \sigma_2)$. Thus Part (b) holds.

To prove Part (c), we need to show that

$$\frac{(\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2)(1 - \rho^2)}{(\sigma_1 + \rho \sigma_2)(\sigma_2 + \rho \sigma_1)}$$

is decreasing in ρ . The derivative with respect to ρ is

$$\frac{[-2\rho(\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2) + 2(1 - \rho^2)\sigma_1 \sigma_2] \cdot (\sigma_1 + \rho \sigma_2)(\sigma_2 + \rho \sigma_1) - (1 - \rho^2)(\sigma_1^2 + 2\rho \sigma_1 \sigma_2 + \sigma_2^2)^2}{(\sigma_1 + \rho \sigma_2)^2(\sigma_2 + \rho \sigma_1)^2}$$

We thus need to show

$$(1 - \rho^2)(\sigma_1^2 + 2\rho\sigma_1\sigma_2 + \sigma_2^2)^2 - 2(1 - \rho^2)\sigma_1\sigma_2(\sigma_1 + \rho\sigma_2)(\sigma_2 + \rho\sigma_1) + 2\rho(\sigma_1^2 + 2\rho\sigma_1\sigma_2 + \sigma_2^2)(\sigma_1 + \rho\sigma_2)(\sigma_2 + \rho\sigma_1) \geq 0.$$

Using $\sigma_1^2 + 2\rho\sigma_1\sigma_2 + \sigma_2^2 = \sigma_1(\sigma_1 + \rho\sigma_2) + \sigma_2(\sigma_2 + \rho\sigma_1)$, this inequality simplifies to

$$((1 + \rho^2)\sigma_1^2 + 2\rho\sigma_1\sigma_2) \cdot (\sigma_1 + \rho\sigma_2)^2 + ((1 + \rho^2)\sigma_2^2 + 2\rho\sigma_1\sigma_2) \cdot (\sigma_2 + \rho\sigma_1)^2 \geq 0.$$

We note that $1 + \rho^2 \geq 2|\rho|$. Thus it suffices to show

$$(2|\rho|\sigma_1^2 + 2\rho\sigma_1\sigma_2) \cdot (\sigma_1 + \rho\sigma_2)^2 + (2|\rho|\sigma_2^2 + 2\rho\sigma_1\sigma_2) \cdot (\sigma_2 + \rho\sigma_1)^2 \geq 0$$

If $\rho \geq 0$, then both terms on the LHS are positive and we are done. Suppose $\rho < 0$, then after taking out the common factor $2|\rho|$ it remains to show

$$(\sigma_1^2 - \sigma_1\sigma_2) \cdot (\sigma_1 + \rho\sigma_2)^2 + (\sigma_2^2 - \sigma_1\sigma_2) \cdot (\sigma_2 + \rho\sigma_1)^2 \geq 0.$$

With a little algebra, this inequality is equivalent to

$$(\sigma_1 - \sigma_2)^2 \cdot (\sigma_1^2 + \sigma_2^2 + (1 + 2\rho - \rho^2)\sigma_1\sigma_2) \geq 0,$$

which indeed holds because $\sigma_1^2 + \sigma_2^2 + (1 + 2\rho - \rho^2)\sigma_1\sigma_2 \geq \sigma_1^2 + \sigma_2^2 - 2\sigma_1\sigma_2 \geq 0$.

I Many Competing Information Providers

Here we demonstrate how the game in Section 7 generalizes to the case of $K > 2$ competing information sources. The setup is similar: the reader seeks to learn $\theta_1 + \dots + \theta_K$ where the noise level ζ_i about each θ_i is controlled by a separate information provider. We assume the reader's prior over these attributes is symmetric; specifically, each attribute has prior variance 1 and each pair of attributes has prior covariance ρ for some $\rho \in (-1, 1)$.⁴⁸

Using the transformation $\tilde{\theta}_i = \frac{\theta_i}{\zeta_i}$, we can reduce the reader's information acquisition problem to our main model with prior covariance matrix

$$\tilde{\Sigma} = \begin{pmatrix} \frac{1}{\zeta_1^2} & \frac{\rho}{\zeta_1\zeta_2} & \cdots & \frac{\rho}{\zeta_1\zeta_K} \\ \frac{\rho}{\zeta_1\zeta_2} & \frac{1}{\zeta_2^2} & \cdots & \frac{\rho}{\zeta_2\zeta_K} \\ \cdots & \cdots & \cdots & \cdots \\ \frac{\rho}{\zeta_1\zeta_K} & \frac{\rho}{\zeta_2\zeta_K} & \cdots & \frac{1}{\zeta_K^2} \end{pmatrix}.$$

⁴⁸By scaling the attributes, it is straightforward to generalize to the case where prior variances are σ^2 and prior covariances are $\rho\sigma^2$. The upshot is that sources simply scale their noise levels by σ .

and weight vector $\tilde{\alpha} = (\zeta_1, \dots, \zeta_K)'$.

Although $\tilde{\Sigma}$ does not in general satisfy Assumption 3, it turns out that the optimal attention allocations can still be characterized in the same way as Theorem 2, thanks to the symmetry in this problem. Specifically, we have:

Lemma 15. *Suppose $\zeta_1 \leq \zeta_2 \leq \dots \leq \zeta_K$. For $1 \leq k \leq K - 1$, define*

$$t_k = \frac{1}{1 - \rho} \sum_{i=1}^k \zeta_i (\zeta_{k+1} - \zeta_i)$$

and define $t_K = +\infty$. Then for any k , the optimal attention allocation is constant at all times $t \in [t_{k-1}, t_k)$ and supported on the first k sources, where each source $i \leq k$ receives attention proportional to its weight ζ_i .

Using this result, it is straightforward to solve for the *symmetric pure strategy equilibrium* of the game. Indeed, suppose the other sources all choose ζ^* ; then, source 1's payoff when choosing $\zeta_1 \leq \zeta^*$ is given by

$$\frac{1}{r} \left(1 - \frac{(K-1)\zeta^*}{\zeta_1 + (K-1)\zeta^*} \cdot e^{\frac{-r\zeta_1(\zeta^* - \zeta_1)}{(1-\rho)}} \right).$$

Differentiating this w.r.t. ζ_1 yields the first-order condition $r \cdot (\zeta_1 + (K-1)\zeta^*) \cdot (2\zeta_1 - \zeta^*) \leq 1 - \rho$ at $\zeta_1 = \zeta^*$, so that $\zeta^* \leq \sqrt{\frac{1-\rho}{Kr}}$.

On the other hand, by choosing $\zeta_1 > \zeta^*$, source 1 gets

$$\frac{\zeta_1}{\zeta_1 + (K-1)\zeta^*} \cdot e^{\frac{-r(K-1)\zeta^*(\zeta_1 - \zeta^*)}{1-\rho}}.$$

Differentiating w.r.t. ζ_1 yields another first-order condition $r \cdot \zeta_1 \cdot (\zeta_1 + (K-1)\zeta^*) \geq 1 - \rho$ at $\zeta_1 = \zeta^*$. Thus $\zeta^* \geq \sqrt{\frac{1-\rho}{Kr}}$, showing such an equilibrium is unique.

Proof of Lemma 15. Fix any stage k and any time $t \in [t_{k-1}, t_k)$ with t_k defined in the lemma. Then, according to the lemma, the t -optimal attention vector $n(t)$ satisfies

$$n_i(t) = \frac{\zeta_i(\zeta_k - \zeta_i)}{1 - \rho} + \frac{\zeta_i}{\zeta_1 + \dots + \zeta_k} \cdot (t - t_{k-1}), \quad \forall 1 \leq i \leq k \quad (16)$$

and $n_i(t) = 0$ for $i > k$. Conversely, if we can show this vector $n(t)$ is indeed t -optimal, then the lemma would follow.

Let q denote this attention vector for ease of exposition. To prove q minimizes the posterior variance function, it is equivalent to check the first-order condition (noting that q is supported on the first k sources):

$$\partial_1 V(q) = \dots = \partial_k V(q) < \min_{i > k} \partial_i V(q).$$

Using Lemma 5, it suffices to show

$$\gamma_1 = \dots = \gamma_k \geq \gamma_{k+1} \geq \dots \geq \gamma_K > 0,$$

where as usual $\gamma = (\tilde{\Sigma} + \text{diag}(q))^{-1} \cdot \tilde{\alpha}$. Observe that the prior covariance $\tilde{\Sigma}$ in the transformed problem can be written as

$$\tilde{\Sigma} = \text{diag}(\zeta)^{-1} \cdot \Sigma \cdot \text{diag}(\zeta)^{-1},$$

with Σ being the matrix having “1”s on the diagonal and “ ρ ” everywhere off the diagonal, and ζ denoting the vector $(\zeta_1, \dots, \zeta_K)'$ (with a slight abuse of notation). From the above discussion, ζ is also the weight vector $\tilde{\alpha}$.

Thus, we can compute the key γ vector as follows:

$$\begin{aligned} \gamma &= (\tilde{\Sigma}^{-1} + \text{diag}(q))^{-1} \cdot \tilde{\alpha} \\ &= (\text{diag}(\zeta) \cdot \Sigma^{-1} \cdot \text{diag}(\zeta) + \text{diag}(q))^{-1} \cdot \zeta \\ &= (\Sigma^{-1} \cdot \text{diag}(\zeta) + \text{diag}(q/\zeta))^{-1} \cdot \text{diag}(\zeta)^{-1} \cdot \zeta \\ &= (\Sigma^{-1} \cdot \text{diag}(\zeta) + \text{diag}(q/\zeta))^{-1} \cdot \mathbf{1}, \end{aligned}$$

where we use $\text{diag}(q/\zeta)$ to denote the diagonal matrix with entries $q_1/\zeta_1, \dots, q_K/\zeta_K$.

We let M denote the matrix $\Sigma^{-1} \cdot \text{diag}(\zeta) + \text{diag}(q/\zeta)$. Then $M \cdot \gamma = \mathbf{1}$, so that

$$\sum_{j=1}^K M_{ij} \cdot \gamma_j = 1, \quad \forall i. \tag{17}$$

We will use these identities to show that each γ_j is positive and $\gamma_1 = \dots = \gamma_k$ are the largest coordinates of γ .

In fact, observe that Σ^{-1} is the matrix with diagonal entries equal to $a = \frac{1+(K-2)\rho}{(1-\rho)(1+(K-1)\rho)}$ and off-diagonal entries equal to $b = \frac{-\rho}{(1-\rho)(1+(K-1)\rho)}$. Thus from $M = \Sigma^{-1} \cdot \text{diag}(\zeta) + \text{diag}(q/\zeta)$ we deduce

$$M_{ij} = b\zeta_j + \left((a-b)\zeta_i + \frac{q_i}{\zeta_i}\right) \cdot \delta_{j=i},$$

with $\delta_{j=i}$ representing the indicator function for the event $j = i$. Plugging this into (17), we then obtain

$$\left((a-b)\zeta_i + \frac{q_i}{\zeta_i}\right) \cdot \gamma_i = 1 - \sum_{j=1}^K b\zeta_j \gamma_j, \quad \forall i.$$

Since the RHS is independent of i , we conclude that $\gamma_1, \dots, \gamma_K$ have the same sign and each γ_i is inversely proportional to $(a-b)\zeta_i + \frac{q_i}{\zeta_i}$.

Now recall that $\gamma = (\tilde{\Sigma}^{-1} + \text{diag}(q))^{-1} \cdot \tilde{\alpha}$. So $\tilde{\alpha}' \cdot \gamma = \tilde{\alpha}' \cdot (\tilde{\Sigma}^{-1} + \text{diag}(q))^{-1} \cdot \tilde{\alpha}$, which is positive since $(\tilde{\Sigma}^{-1} + \text{diag}(q))^{-1}$ is a positive-definite matrix. It follows that the coordinates

of γ cannot all be less than or equal to zero. By the preceding analysis, they must all be positive. Finally, to show $\gamma_1, \dots, \gamma_k$ are equal and larger than the remaining coordinates, it suffices to consider their inverses, which are proportional to $(a - b)\zeta_i + \frac{q_i}{\zeta_i}$. From (16) and $a - b = \frac{1}{1-\rho}$ we indeed have

$$(a - b)\zeta_i + \frac{q_i}{\zeta_i} = \frac{1}{1 - \rho} \cdot \zeta_k + \frac{t - t_{k-1}}{\zeta_1 + \dots + \zeta_k}, \quad \forall 1 \leq i \leq k.$$

The RHS is the same for $i \leq k$ and smaller than $(a - b)\zeta_{k+1}$ when $t < t_k$. This completes the proof that $\gamma_1 = \dots = \gamma_k \geq \gamma_{k+1} \geq \dots \geq \gamma_K$. Lemma 15 follows. \square

References

- AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.
- AZEVEDO, E., A. DENG, J. L. M. OLEA, J. RAO, AND G. WEYL (2019): “A/B Testing with Fat Tails,” Working Paper.
- BARDHI, A. (2019): “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.
- BLACKWELL, D. (1951): *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley and Los Angeles: University of California Press, chap. Comparison of Experiments, 93–102.
- 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.
- CHAN, J. AND W. SUEN (2008): “A Spatial Theory of News Consumption and Electoral Competition,” *Review of Economic Studies*, 75, 699–728.
- CHE, Y.-K. AND K. MIERENDORFF (2019): “Optimal Dynamic Allocation of Attention,” *American Economic Review*, forthcoming.
- CHEN, H. AND W. SUEN (2020): “Competition for Attention in the News Media Market,” Working Paper.
- 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,” *Informatics 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, 3651–84.
- GALPERTI, S. AND I. TREVINO (2020): “Competition for Attention in the News Media Market,” Working Paper.
- GARFAGNINI, U. AND B. STRULOVICI (2016): “Social Experimentation with Interdepen-

- dent and Expanding Technologies,” *Review of Economic Studies*, 83, 1579–1613.
- GENTZKOW, M. AND J. SHAPIRO (2006): “Media Bias and Reputation,” *Journal of Political Economy*.
- (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 (2019): “Attention Please!” Working Paper.
- 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 (2018): “Information Costs and Sequential Information Sampling,” Working Paper.
- IMMORLICA, N., I. KASH, AND B. LUCIER (2018): “The Evolving Value of Data,” Working Paper.
- 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 Economic Behavior*, 88, 190–206.
- LIANG, A. AND X. MU (2020): “Complementary Information and Learning Traps,” *Quarterly Journal of Economics*, 135, 389–448.
- LIANG, A., X. MU, AND V. SYRGKANIS (2017): “Optimal and Myopic Information Acquisition,” Working Paper.
- MAYSKAYA, T. (2019): “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.
- MOSCARINI, G. AND L. SMITH (2001): “The Optimal Level of Experimentation,” *Econometrica*, 69, 1629–1644.
- MULLAINATHAN, S. AND A. SHLEIFER (2005): “The Market for News,” *American Economic Review*.
- PEREGO, J. AND S. YUKSEL (2020): “Media Competition and Social Disagreement,” Working Paper.
- 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.
- 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.
- ZHONG, W. (2019): “Optimal Dynamic Information Acquisition,” Working Paper.