

Dynamically Aggregating Diverse Information^{*,†}

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

July 15, 2019

Abstract

An agent has access to multiple data 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.

*This paper grew out of and subsumes parts of our earlier working paper “Optimal and Myopic Information Acquisition,” which studies a discrete-time setting. See Section 6 for a discussion of the relationship between the two papers.

†We are grateful to Yash Deshpande, Drew Fudenberg, George Mailath, Konrad Mierendorff, Lan Min, Peter Norman Sørensen, Jakub Steiner, and Philipp Strack for helpful comments and suggestions.

‡Department of Economics, University of Pennsylvania

§Cowles Foundation and Columbia University

¶Microsoft Research

1 Introduction

Markets are increasingly saturated with large quantities of information about consumers, which firms can use to learn about their preferences and behaviors. But since this information is usually not designed for the *exact* problem that the firm or decision-maker faces, learning often takes the form of acquiring and aggregating different kinds of information. A key question then is how to optimally acquire information for this goal.

To fix ideas, suppose a hotel chain is considering whether to open a new location in Puerto Rico, and wants to know the total amount of travel volume that the hotel would experience. There is not a direct source of information about this, but the firm can acquire data to learn about different aspects of its problem—for example, website traffic to the Puerto Rico tourism bureau helps the hotel to estimate tourism demand, while search data for conference venues helps it to estimate business demand. Data aggregation of this form is costly, not only because the firm needs to purchase the data, but also because the data needs to be processed—raw data is often unstructured, and employees (e.g., statisticians in a data science division) need to expend considerable effort to organize and analyze the data.

In this work, we present a simple model for dynamic aggregation of information given limited resources. Formally, a decision-maker (e.g., the management at the hotel chain) has access to various *data sources*, each modeled as a Brownian motion whose drift is an unknown *attribute* that the data source provides information about (e.g., tourism travel, business travel). The decision-maker can continuously allocate a budget of resources (e.g., employee hours) across these Brownian motions, where more resources allocated to any data source results in greater precision of information about the corresponding attribute value. The decision-maker acquires information until an endogenously chosen time, at which point he implements a decision based on the information acquired so far. We assume that the variable the decision-maker wants to learn (so called *payoff-relevant state*) is a linear combination of the attribute values. (In the example, total travel volume is a sum of demand from various consumer segments.)

The key difficulty in the decision-maker’s information acquisition problem comes from possible correlation across different attribute values. For example, one data source may provide information about vacation travel demand from the U.S., while another provides information about vacation travel demand from Canada, and we expect these values to be positively correlated. Although these two demand levels enter additively into the firm’s total demand forecast, their correlation means that information about one affects the value of information about the other. Thus, current information acquisitions have an immediate

impact (reducing present uncertainty), as well as an impact on the future value of different data sources. At any given moment, the decision-maker has to choose the optimal proportion of resources to allocate to each source, taking into account the potential complementarity or substitutability among the different sources.

Our main results demonstrate that under a condition on the prior belief, the optimal dynamic data acquisition strategy takes a simple form. Initially, the decision-maker exclusively observes the single most informative data 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 additionally from new data sources, dividing resources over these new sources and the ones he was learning from previously. Eventually, the decision-maker learns from all sources using a final and constant division of resources.

Crucially, the sources that are observed at each stage, and the way in which resources are divided across them, are history-independent—that is, they do not depend on the signal realizations. Thus, the decision-maker can completely determine its plan for information acquisition at time $t = 0$. We also show that the optimal information acquisition strategy is “robust” in the sense that it does not depend on the decision-maker’s discount factor or payoff function, so long as the variable that the decision-maker wants to learn remains the same.

The condition that we assume on the prior belief requires that the prior covariances of the different attributes are not too large in magnitude compared to their variances. For the case of two attributes, it is sufficient for the covariances to be smaller than the variances; in general, how much smaller depends on the number of data sources. Intuitively, such a condition puts an upper bound on the possible complementarity or substitution effects between different data sources. This helps to align the decision-maker’s short-run and long-run information acquisition incentives, so that it is optimal to focus on the most informative sources at any given moment.¹ Although in general our condition puts some restriction on the prior belief, we show that under optimal sampling from *any* prior belief, the decision-maker’s posterior beliefs will eventually satisfy this condition, after which point our characterization will hold.

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

¹Formally, these are the sources that maximize the marginal reduction of the posterior variance about the payoff-relevant state. See Section 3.3 for details.

comparative statics in primitives such as correlation across attributes). This makes the characterizations useful for deriving new substantive results in settings motivated by particular applications. We now illustrate this with two examples.

One such problem is endogenous information acquisition for binary choice. A large literature in economics and neuroscience (originating with [Ratcliff and McKoon \(2008\)](#)) has studied how a consumer chooses between two goods of unknown payoffs. A result in [Fudenberg et al. \(2018\)](#) characterizes the optimal way to learn about the two payoffs prior to making this decision, when these payoffs are Gaussian and i.i.d.

The model studied in [Fudenberg et al. \(2018\)](#) is nested in our framework as the case of *two* unknown attributes (the unknown payoffs), where 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 previously mentioned result in [Fudenberg et al. \(2018\)](#) to correlated payoffs.² In addition, we use our characterization to derive a new comparative static result with respect to prior uncertainty. An increase in the initial uncertainty about either payoff results in a uniform change in attention: either weakly more attention paid to learning about that payoff at every instant, or weakly less at every instant. But the direction of this change depends on the degree of prior correlation between the unknown payoffs. Specifically, we show that an increase in initial uncertainty results in uniformly *more* attention to that payoff when the two payoffs are weakly correlated in the prior, but results in uniformly *less* attention under strong prior correlation.

Our second application considers an extension of our environment in which the data providers (e.g., news sources) are themselves strategic, and can control the precision of the information they provide. We suppose that a mass of forward-looking decision-makers optimally acquire information from the sources over time. Using our characterization of the decision-makers' optimal information acquisition strategy, we derive the sources' equilibrium choices of precision.

These precision levels turn out to be monotonically increasing in the providers' discount rate and in the prior correlation between the unknown attributes. Specifically, the *more patient* the providers are, and the *less correlated* the attributes are, the *lower the precision* of the signals. We show this by studying the optimal information acquisition strategy, which starts with all resources directed towards one source (the more informative one), and eventually divides resources across the two sources in a ratio that depends on their relative levels of informativeness. When a source becomes more informative, it is more likely to

²We would expect payoffs to be correlated if the values of the goods depend on a common source of uncertainty—e.g., if the goods represent different portfolio choices, or consumer products with shared features.

be the source initially attended to, but will receive fewer resources in the long-run (since decision-makers need to spend less time to achieve the same level of information from that source). Thus, patient data providers choose less informative signals, while impatient data providers compete to be chosen initially. As far as we are aware, the effect of information precision on the *time path* of people’s information demand has not been noted in the previous literature, and our exact characterizations are what allow us to study this.

Our analyses in Section 4 (information acquisition for binary choice) and Section 5 (competing news sources) are only two problems whose solutions are facilitated by our main results, and we hope that the characterizations we provide can be used in future work on other applications.

1.1 Related Literature

Our model resembles, but is not nested within, 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.³

We primarily 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 different *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). 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.⁴ Compared to these papers,

³This feature distinguishes our results relative to a classic literature on “learning by experimentation” (Easley and Kiefer, 1988; Aghion et al., 1991; Keller et al., 2005). In our paper, “myopic information acquisition” (one that maximizes the reduction of posterior variance) is *exactly* optimal because it achieves the fastest speed of learning. By contrast, Easley and Kiefer (1988) and Aghion et al. (1991) showed that if there is a unique “myopically optimal policy” (one that maximizes the flow payoff) at the limiting belief, then the optimal policies *eventually converge* to this policy. This argument does not apply to our setting, since every policy (i.e., resource allocation) is trivially myopic at the limiting belief, which is a point mass at the true parameters. Uniqueness fails.

⁴Gossner et al. (2019) study the consequences of attention manipulations, where the agent is forced to

we study a setting where the agent dynamically learns about many *correlated* attributes.⁵ Che 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 one of two states.

In the specific 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. Aside from having noisy signals, the main distinction of our model is that we allow for correlation across attributes. Correlation translates into complementarity or substitution effects between different data sources, and thus plays a key role in determining how the agent should aggregate information from these sources.

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).⁶ Compared to these papers, our agent has access to a prescribed (physical) set of signals, and acquires information under a resource/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.⁷

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.⁸ We did not obtain a characterization of the optimal strategy itself.

attend initially to one particular attribute. This interesting question bears certain high-level resemblances to our comparative statics in Section 4. 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.

⁵Relatedly, Callander (2011) considers sequential search from correlated signals. But the signals in Callander (2011) come from a single Brownian motion path, which yields a special correlation structure. Similar models are studied in Garfagnini and Strulovici (2016) and in Bardhi (2018).

⁶It 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.

⁷Formally, we consider a sequential sampling problem in which the flow cost of acquiring information only depends on the current time. Whereas in Morris and Strack (2019), for example, the flow cost is a function of the current belief.

⁸In 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

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, as we illustrate. The technical methods in this paper also differ from the prior work—see Section 6 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. [Bonatti and Cisternas \(2019\)](#) analyze a dynamic game in which firms use a consumer’s “score” to make inferences about his preferences and set prices. Different from these papers, we have a setting in which the firm has to dynamically aggregate multiple sources of information. Our characterizations trace out a time path of market demand for various kinds of information, which is absent from the literature.

2 Model

An agent (i.e., firm) has uncertainty about the values of K attributes $\theta = (\theta_1, \dots, \theta_K)'$,⁹ and his prior is that they are jointly normal with known mean vector $\mu \in \mathbb{R}^K$ and covariance matrix Σ , where Σ has full rank. The agent wants to learn an unknown payoff-relevant state $\omega = \sum_{i=1}^K \alpha_i \theta_i$, which is a linear combination of these attribute values. The weight vector $\alpha \in \mathbb{R}^K$ is known and fixed, and we assume for ease of exposition that each coordinate α_i is strictly positive, so that the state depends positively on all of the attribute values. Because any attribute value can be replaced with its negative, it is without loss to assume that the weights are non-negative. Moreover, any source with zero weight can be dropped from the model without affecting our results (see Appendix D.6).

Time is continuous. There is a data source that provides information about each attribute, and the agent divides his attention (i.e., resources) across these sources at every instant. Formally, we assume that the agent has one unit of attention in total at every point in time, and chooses attention allocations $\beta_1(t), \dots, \beta_K(t)$ subject to $\beta_i(t) \geq 0$ and

a myopic procedure will be asymptotically efficient. In another working paper [Liang and Mu \(2019\)](#), we provide a more thorough analysis of the conditions on the informational environment under which myopic acquisitions lead to long-run (in)efficient learning.

⁹Here and later, we use the apostrophe to denote vector or matrix transpose.

$\sum_i \beta_i(t) \leq 1$. See for example [Fudenberg et al. \(2018\)](#) and [Che and Mierendorff \(2019\)](#) for recent models with fixed budgets of attention.

These 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 For example, devoting T units of time with full attention on source i (that is, $\beta_i(t) = 1$ at every time t in this interval) is equivalent to a single observation of the signal $\theta_i + \mathcal{N}(0, 1/T)$, or T observations of $\theta_i + \mathcal{N}(0, 1)$.¹¹

Although we have assumed that the drift of each X_i is proportional to an individual attribute θ_i , the same analysis applies if this drift is instead some linear combination $a'_i \cdot \theta$ with $a_i \in \mathbb{R}^K$. This is because we can re-define the “primitive” attribute values $\tilde{\theta}_i = a'_i \cdot \theta$. Then, the vector of re-defined attributes $\tilde{\theta} = (\tilde{\theta}_1, \dots, \tilde{\theta}_K)'$ is again jointly normal, and the payoff-relevant state ω can be expressed as a (different) linear combination of $\tilde{\theta}_i$. This transformation is valid so long as the vectors a_1, \dots, a_K are linearly independent.

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 α_i is special. For example, with $K = 2$, Theorem 1 always holds 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 .

Throughout the paper, we use the observation that after devoting q_i units of attention to each data source i , the agent's posterior covariance matrix about $(\theta_1, \dots, \theta_K)$ is given by

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

where Σ is the prior covariance matrix and $\text{diag}(q)$ is the diagonal matrix with entries q_1, \dots, q_K . This formula reflects the general fact that in Gaussian environments, 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 ($\text{diag}(q)$ in this case). Since the agent is ultimately interested in learning about $\omega = \sum_{i=1}^K \alpha_i \theta_i$, what matters is his *posterior variance about ω* , as given by

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

See Appendix A for various properties of the V function.

3 Optimal Information Acquisition Strategy

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

3.1 $K = 2$

We begin by considering the case of two data sources and two attributes. The agent has a prior

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

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

and access to two Brownian motions. He seeks to learn $\omega = \alpha_1\theta_1 + \alpha_2\theta_2$, where each $\alpha_i > 0$.

We impose the following restriction on the agent's prior belief:

Assumption 1. *The prior covariance matrix satisfies $\alpha_1(\Sigma_{11} + \Sigma_{12}) + \alpha_2(\Sigma_{21} + \Sigma_{22}) \geq 0$.*

Since both variances Σ_{11}, Σ_{22} are positive, Assumption 1 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 1 holds for all priors.¹³ A different sufficient condition is for the attributes to be positively correlated ($\Sigma_{12} = \Sigma_{21} \geq 0$), in which case Assumption 1 holds for all weights α_1 and α_2 . We note that the set of beliefs satisfying Assumption 1 is absorbing: Once a belief satisfies Assumption 1, all subsequent posterior beliefs (following any strategy, not necessarily optimal) will as well.

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

Theorem 1. *Suppose Assumption 1 is satisfied. Define*

$$t_1^* := \frac{y_1 - y_2}{x_2}; \quad t_2^* := \frac{y_2 - y_1}{x_1}$$

where $x_1 = \alpha_1 \det(\Sigma)$, $y_1 = \alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{12}$, $x_2 = \alpha_2 \det(\Sigma)$, and $y_2 = \alpha_1 \Sigma_{21} + \alpha_2 \Sigma_{22}$. *W.l.o.g. let $y_i \geq y_j$. 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

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

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 so long as 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 functional form of $u(a, \omega)$. Thus, when the prior belief satisfies Assumption 1, 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.

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 $t_1 = 5/6$, his posterior covariance matrix is the identity matrix. After this time he optimally splits attention equally between the two attributes, which are now symmetric.

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 $t_1 = 5/2$, 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 $t_1 = 3/2$, 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 $y_1 = \alpha_1 \Sigma_{11} + \alpha_2 \Sigma_{12} \geq \alpha_1 \Sigma_{21} + \alpha_2 \Sigma_{22}$ 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 $y_1 \geq y_2$ equivalently says that the *marginal value* of learning about attribute θ_1 exceeds that of learning about θ_2 , according to the prior belief. The expression

$$y_i = \alpha_i \Sigma_{ii} + \alpha_j \Sigma_{ij}$$

takes into account the consequences of unequal payoff weights and possible correlation across the attributes. The agent re-weights uncertainty about θ_i in proportion to the weight α_i , since (all else equal) an increase in α_i means that the value of reducing uncertainty about θ_i is larger. Correlation further implies that learning about θ_i results in a reduction of variance not only about θ_i , but also about θ_j . The more the two attributes co-vary, the larger this spillover effect. Thus the marginal value of learning about θ_i includes the payoff consequences to indirectly learning about θ_j , as seen in the above formula for y_i .

Suppose without loss of generality that $y_1 \geq y_2$, then the agent initially learns exclusively about θ_1 , which has greater marginal value. As information about θ_1 accumulates, however, 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 data sources, so that acquiring information in this mixture remains optimal.

We provide a more involved proof outline in Section 3.3, but the intuition can already be seen through the examples above. 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 $t_1 = 3/2$ has the key property that the payoff-relevant state $\omega = \theta_1 + 2\theta_2$ is *independent* of $\theta_1 - \theta_2$ (since they are jointly normal and have zero covariance).¹⁴ As we show in Lemma 5, this independence property implies equal marginal values.¹⁵ 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.

3.2 General K

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

A key condition on the prior belief, parallel to the one stated in Assumption 1, is the following:

Assumption 2. *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 is bounded by an expression depending on the variances.¹⁶ 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 1.¹⁷ In general, the condition in Assumption 2 is more restrictive for larger numbers of sources K .

To interpret the use of Assumption 2, note that prior covariances measure the *complementarity or substitution effects* across the information provided by different data sources (i.e.,

¹⁴ $\text{Cov}(\theta_1 + 2\theta_2, \theta_1 - \theta_2) = \text{Var}(\theta_1) + \text{Cov}(\theta_1, \theta_2) - 2\text{Var}(\theta_2) = 3/5 + 1/5 - 2 \times 2/5 = 0$.

¹⁵Indeed, 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 $\theta_1 + \theta_2$.

¹⁶Note that this condition requires the covariances to be not too negative, and also not too positive, which differs from the previous Assumption 1. 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/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 $K = 2$.

¹⁷However, when $K = 2$ our previous Assumption 1 is strictly weaker. So Theorem 1 does not follow as a corollary from Theorem 2 below.

whether information from one data source increases or decreases the learning benefits from other sources). Assumption 2 limits the magnitude of such complementarity/substitution, so that the agent’s short-run and long-run information acquisition incentives are aligned. In Section 3.4, we provide a counterexample to illustrate that misalignment can occur when Assumption 2 is violated.

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

Theorem 2. *Suppose Assumption 2 is satisfied. 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 involves constant instantaneous attention 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 .

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’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 the prior satisfies Assumption 2. The agent wants to learn $\omega = \theta_1 + \theta_2 + \theta_3$.

The optimal information acquisition strategy consists of three stages:

Stage 1. The agent begins by putting 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 $t_1 = 1/12$, 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 $t_2 = 13/44$, 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 $(1/3, 1/3, 1/3)$.

We provide next an outline for the proof of this result, which also formalizes some of the intuitions hinted at in the examples above—e.g., that the marginal values of learning about different attributes in the set B_k are held constant during each stage k .

3.3 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 in fact exist, and have the structure that we characterize.

Definition of a uniformly optimal strategy. 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. These cumulated attention vectors $q(t)$ determine the agent’s posterior variance about ω , via the function $V(q(t))$ introduced in (2).

Define the *t-optimal* cumulated attention vector to be

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

namely the allocation of t units of attention that minimizes posterior variance (among all attention vectors that allocate a budget of t).¹⁸ We will say that an attention allocation strategy is uniformly optimal if it integrates to the *t-optimal* vector at every time 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.

That is, the strategy S deterministically leads to minimum posterior variance about ω at every possible stopping time t . 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), then 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).

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

¹⁸We show in Lemma 6 that this minimizer is unique.

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). Given this discussion, whenever a uniformly optimal strategy exists, it must be the optimal strategy in our problem.¹⁹ It remains to show that under Assumption 2, 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 1. *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 to have higher amount of attention allocated to each source 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.

Suppose the agent has achieved the t -optimal vector $n(t)$ at some time t , and is trying to reach $n(t')$ at some future time t' slightly larger than t . When the marginal value of learning about some attribute is strictly largest at time t , it is optimal to focus on that attribute for a while. More often, however, there will be multiple sources that have the same maximal marginal value. In these cases, the agent turns from the “first-order” comparison of marginal values to a “second-order” comparison of mixtures over this set of sources, since all these mixtures have the same first-order effect. Formally, different mixtures have the same marginal value at time t , but due to second-order effects that we describe below, they would have different marginal values at future instants.²⁰

¹⁹While 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.

²⁰We 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 4. 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 optimal mixture depends (roughly) on whether the sources of information (that maximize the marginal value at t) are substitutes or complements. If information about different attributes are *substitutes*—so that information about attribute 1 has a negative impact on the marginal value of information about attribute 2—then the agent prefers *not* to observe both data sources in positive amounts from t to t' .²¹ Instead, he would ideally like to take away some attention given to attribute 1 before time t and re-distribute it to attribute 2 between t and t' . This would create a failure of monotonicity, since $n_1(t')$ is smaller than $n_1(t)$. In these situations there does not exist a uniformly optimal strategy, as we illustrate by example in the next subsection.

In contrast, if information about different attributes are complements, then the agent optimally chooses a *positive* mixture to take advantage of the complementarity. In this case the t -optimal vectors are weakly increasing in all coordinates, and uniform optimality is attainable. What we show in the proof is that Assumptions 1 and 2 are sufficient to guarantee that different data sources are complements whenever their marginal values equalize.²² Thus a uniformly optimal strategy exists under the stated assumptions.

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 $n(t)$. Since $n(t)$ itself is history-independent, so is $\beta(t)$. This delivers the first part of our theorems: existence of an optimal strategy that is deterministic.

It remains to characterize the structure of this strategy. By our discussion above, at each time t the agent divides attention across learning those attributes that maximize the marginal reduction of posterior variance. Suppose that $n(t)$ is supported on some set B_k , then these sources have the highest marginal value at time t . We demonstrate a specific positive mixture over these sources, such that the marginal values of these sources remain the same when the agent divides attention according to this mixture. This allocation is thus optimal for a while after time t .²³

Nonetheless, as beliefs about the attributes in B_k become precise, the marginal values of

²¹This intuition is rough because it does not take into account the substitution effect a source of information has with past sampling of the *same* source. More precisely, it is not optimal to observe both data sources in positive amounts if and only if the substitution between them is sufficiently strong.

²²The formal version of this claim is Lemma 10 in the appendix. Note that complementarity or substitution of two data sources is captured by the relevant cross-partial derivative of the posterior variance function V , given in Lemma 5.

²³Since $V(q)$ is a convex function, the first-order condition of having equal marginal values is not only necessary but also sufficient for a cumulated attention vector to be t -optimal.

learning about the remaining attributes increase continuously relative to the value of learning about those in B_k . Eventually some new data source(s) will have the same marginal value as those in B_k . 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.

The final difficulty is to argue that given any observation set B_k , the specific mixture over these sources that maintains equal marginal values is constant over time. For this we directly compute the time-derivative of $n(t)$: Since $n(t)$ is characterized by equal partial derivatives of V , its time-derivative is inversely related to the Hessian matrix (i.e., second derivatives) of V . We are able to derive this Hessian matrix (Lemma 5) and use its properties to characterize the evolution of $n(t)$ (Lemmata 11 and 12). This completes the proof.

3.4 Arbitrary Priors

We next comment on optimal information acquisition for prior beliefs that do not satisfy the assumptions given above.

It turns out that the agent’s posterior beliefs under optimal sampling from *any* prior belief will eventually satisfy Assumption 2. In fact, optimal sampling is not required: Along any path in which each data 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 2.

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

Lemma 2. *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 $t \rightarrow \infty$.*²⁴

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

Once Assumption 2 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 2, Lemma 3 and Theorem 2 that:

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

²⁴We 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.

weight vector α .²⁵

Thus, in general, the optimal information acquisition strategy will eventually have the properties described in the previous subsections: independence of signal realizations, of the payoff function $u(a, \omega)$ and of the waiting cost $c(\tau)$.

These robustness properties need not hold from $t = 0$ for all priors. Specifically, the condition provided in Assumption 1 for $K = 2$ is not only sufficient but also necessary for our characterization to hold independently of the agent’s payoff criterion—see the example below and Proposition 4 for details. The condition we provide in Assumption 2 for general K is sufficient but not necessary for the characterization to hold.²⁶ Nonetheless, the constant $\frac{1}{2K-3}$ in Assumption 2 is tight, in a sense that we formalize in Appendix D.7.

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 3.3) 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

²⁵More 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 . This holds independently of the payoff function or past signal realizations.

²⁶That is, under alternative assumptions on Σ and α , optimal information acquisition may also consist of K stages as described in Theorem 2. In fact, in the appendix we prove Theorem 2 under a weaker condition than Assumption 2.

at $t = 1/4$, since $\omega = \theta_1 + 4\theta_2$ is independent of $\theta_1 + \theta_2$ (see Footnote 15).²⁷ However, to 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 $n(t) = \left(\frac{-t+1}{3}, \frac{4t-1}{3}\right)$, and the optimal cumulated attention toward attribute 1 is decreasing in this interval.²⁸

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 1). Hence the optimal information acquisition strategy varies according to when the agent expects to stop, and Theorem 1 cannot hold independently of the payoff criterion (Lemma 8).

4 Application: 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 θ_1 and $-\theta_2$ (we introduce the negative here for expositional simplicity), 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 $(\theta_1, -\theta_2) = (\theta', \theta'')$ or $(\theta_1, -\theta_2) = (\theta'', \theta')$ where $\theta' > \theta''$ are known quantities. Thus the agent has uncertainty over which good is better, but not over how much better it is.²⁹ 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 $(\theta_1, -\theta_2)$, and has access to two Brownian motions with drifts corresponding to these unknown payoffs.

Both the classic DDM model and also Fudenberg et al. (2018) focus on the optimal stop-

²⁷The 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 1 and 2 essentially rule out the other possibility that ω is independent of $\theta_1 + \theta_2$.

²⁸Subsequently, at times $t \in (1, 3]$, the t -optimal vector is $n(t) = (0, t)$, allocating all attention to attribute 2. Finally, at times $t \geq 3$, $n(t) = \left(\frac{t-3}{5}, \frac{4t+3}{5}\right)$, allocating attention proportional to α .

²⁹That is, the classic DDM assumes that the magnitude of $\theta_1 + \theta_2$ is known to the agent.

ping rule given *exogenous* information. But Section E of [Fudenberg et al. \(2018\)](#) additionally considers 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. Indeed, this corresponds exactly to our framework with $K = 2$ and equal payoff weights (since the payoff difference $\theta_1 + \theta_2$ is a sufficient statistic for the agent’s decision). These authors impose further that the agent’s prior is *independent and symmetric*—that is, $\Sigma = I$ —and find that the agent optimally devotes equal attention to both information sources at all times.

Applying Theorem 1 with $\alpha_1 = \alpha_2 = 1$, we obtain the following immediate generalization of this result on optimal information acquisition.³⁰

Corollary 1. *Suppose $K = 2$, $\alpha_1 = \alpha_2 = 1$ and $\Sigma_{ii} \geq \Sigma_{jj}$. The agent’s optimal information acquisition strategy $(\beta_1(t), \beta_2(t))$ consists of two stages:*

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

$$t \leq t_i^* = \frac{\Sigma_{ii} - \Sigma_{jj}}{\det(\Sigma)},$$

the agent optimally allocates all attention to source i .

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

When $\Sigma = I$, the thresholds are $t_1^* = t_2^* = 0$, 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 whenever the prior belief is ex-ante “asymmetric,” the agent initially devotes all attention to learning about the payoff he deems more uncertain.

We note additionally that, as [Fudenberg et al. \(2018\)](#) point out, their result does not characterize “off-equilibrium” attention allocation (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.

³⁰Note that [Fudenberg et al. \(2018\)](#) additionally provide results about the optimal stopping time, which we do not pursue here.

From Corollary 1 we see that the prior belief affects the agent’s attention strategy only by determining which source is observed in Stage 1, and for how long that stage lasts. Thus, changes in the prior belief result in the agent paying *uniformly* more or less attention to either source. If we consider in particular the impact of changes in the initial uncertainty about one of the payoffs, we have the following comparative static:

Corollary 2. *Suppose $K = 2$, $\alpha_1 = \alpha_2 = 1$ and $\Sigma_{ii} \geq \Sigma_{jj}$. Then, if $\Sigma_{jj} \geq |\Sigma_{ij}|$, an increase in Σ_{ii} results in uniformly higher attention towards source i (i.e., $\beta_i(t)$ is weakly larger at every t). Otherwise if $\Sigma_{jj} < |\Sigma_{ij}|$, an increase in Σ_{ii} results in uniformly lower attention towards source i .*

The case in which larger Σ_{ii} results in uniformly higher attention towards source i is intuitive, since the agent wants to make up for greater initial uncertainty about θ_i . But the comparative static is reversed when the covariance Σ_{ij} is larger in magnitude than Σ_{jj} . To interpret this finding, note that whenever the two payoffs are correlated, any information acquired about θ_i also provides information about θ_j . So in general, an increase in Σ_{ii} has two opposing effects on t_i^* . On the one hand, greater asymmetry in the prior belief means it should take longer time to “balance out” the beliefs (the intuition given above). On the other hand, holding fixed Σ_{ij} and Σ_{jj} , larger Σ_{ii} *decreases the correlation* between the unknowns, so that each unit of attention devoted to θ_i now reveals less about the other payoff θ_j . It should then be faster for the posterior variance about θ_i to “catch up” with the posterior variance about θ_j . Therefore, whether attention is uniformly increased or decreased depends on which of these two effects dominates. As stated in the corollary, the effect of (decreased) correlation is dominant when Σ_{ij} is large in magnitude; that is, when correlation is high to begin with.

5 Application: Competing News Sources

Next, we apply our results to a setting in which the data sources are themselves strategic, and can control the precision of the information that they provide.

For example, consider two news sources, each of which has expertise on a particular topic—for example, one source may specialize in a politician’s financial dealings (i.e., how corrupt he is), while another specializes in that politician’s handling of international relations (i.e., how competent he is). The news sources primarily earn revenue by running ads on an online site, so what they aim to maximize is time spent on their site. They do this by strategically controlling how informative their articles are; for example, they may either

reveal everything they know at once, or reveal it slowly across many articles. On the other hand, readers care to learn the overall quality of the politician, which depends both on how corrupt and how competent the politician is. We ask how informative the news articles will be in equilibrium.

In more detail, we suppose that a mass of readers seek to learn the sum of attributes θ_1 and θ_2 , and their common prior over these parameters is

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

where $\rho \in (-1, 1)$ measures prior correlation between θ_1 and θ_2 . Each of two news sources $i = 1, 2$ (freely) chooses a noise variance σ_i^2 , so that a unit of time spent on their site generates the signal

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

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 variances σ_i .

Readers optimally allocate attention given these noise variances (which are fixed across time). Since we have assumed a common prior, all readers make the same information acquisition decisions, and it is without loss to consider a single reader whose allocation at time t is denoted $(\beta_1(t), \beta_2(t))$. To map this setting into our main model, we normalize the noise terms to have *unit variances* as follows: Define $\tilde{\theta}_i = \frac{\theta_i}{\sigma_i}$, so that each unit of time spent on source i generates a signal about $\tilde{\theta}_i$ with standard Gaussian noise. Under this transformation, the reader seeks to learn $\sigma_1 \tilde{\theta}_1 + \sigma_2 \tilde{\theta}_2$, and his prior covariance matrix over $(\tilde{\theta}_1, \tilde{\theta}_2)$ is

$$\tilde{\Sigma} = \begin{pmatrix} \frac{1}{\sigma_1^2} & \frac{\rho}{\sigma_1 \sigma_2} \\ \frac{\rho}{\sigma_1 \sigma_2} & \frac{1}{\sigma_2^2} \end{pmatrix}.$$

Note that Assumption 1 is satisfied in this transformed problem, since

$$\sigma_1(\tilde{\Sigma}_{11} + \tilde{\Sigma}_{12}) + \sigma_2(\tilde{\Sigma}_{21} + \tilde{\Sigma}_{22}) = (1 + \rho) \left(\frac{1}{\sigma_1} + \frac{1}{\sigma_2} \right) \geq 0.$$

Thus the optimal attention choices $(\beta_1(t), \beta_2(t))$ are characterized by Theorem 1.

Each news source i 's payoff is the discounted average attention paid to that source $\int e^{-rt} \beta_i(t) dt$, where r is a (common) discount rate. We can interpret this as reduced form for advertising revenue, where each news source receives a payoff proportional to the time

the reader spends on its site.³¹

Proposition 2. *The unique equilibrium is a pure strategy equilibrium (σ^*, σ^*) with*

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

Equilibrium precision $(1/(\sigma^)^2 = \frac{2r}{1-\rho})$ is monotonically increasing in the discount rate r and also in the prior correlation ρ .*

Thus, the less patient the information providers are, the more precise the signals are in equilibrium. Intuitively, when news source i increases the informativeness of the articles it provides, there are two opposing effects: On the one hand, weakly more attention is attracted to i early on and it is more likely to be the source chosen in Stage 1, since i 's information becomes more valuable initially. On the other hand, increasing precision lowers the long-run frequency $\frac{\sigma_i}{\sigma_i + \sigma_j}$ with which i is viewed, since readers need to spend less time on site i to achieve the same level of information about θ_i . Thus, less patient data providers compete over short-run profits (i.e., being chosen in Stage 1) and provide precise signals, while patient data providers compete for long-run profits (i.e., long-run frequency) and provide imprecise signals. We are not aware of prior literature that studies this effect of information precision on the *time path* of people's information demand.

Additionally, the more positively correlated the unknown attributes are (i.e., higher covariance ρ), the higher the precision of signals provided in equilibrium. This is because (as we derive in the proof of the proposition) the threshold $t_i^* = \frac{(\sigma_j - \sigma_i)\sigma_i}{1-\rho}$ increases in ρ , which increases the value of being chosen in Stage 1. The competition for short-run profits thus drives the news sources to be more informative.

The parameter ρ measures the degree to which what the news sources know overlap. The case of $\rho = 1$ corresponds to full competition, in which case the sources choose perfectly precise signals in equilibrium ($\sigma^* = 0$). In general, the less substitutable the sources are (smaller ρ), the less competition there is and the more information is withheld.

From the perspective of social welfare, these comparative statics tell us that more information is released into society (and hence society learns faster) when information providers are *less forward-looking*, and when the information they provide is *more similar*. Note that a crucial aspect of the game we have analyzed is that each news source has a ‘‘monopoly’’ on

³¹Here, for the sake of illustrating the equilibrium, we are considering the case where readers sample forever. This corresponds to the limit as the waiting cost $c(\cdot)$ decreases to zero.

some kind of information, so that agents eventually need to observe both sources. An interesting direction for future work would be to consider the case in which there are redundancies in the kinds of information provided by the different sources.

Finally, in Appendix G we generalize these insights to a game where $K > 2$ information providers compete, and where agents seek to learn $\theta_1 + \dots + \theta_K$. Observe that the transformed prior covariance matrix $\tilde{\Sigma}$ does *not* in general satisfy Assumption 2.³² Nonetheless, we directly compute the uniformly optimal strategy (defined in Section 3.3) and show that our K -stage characterization of optimal information acquisition extends to this setting. Generalizing the result in Proposition 2, we find $\sigma^* = \sqrt{\frac{1-\rho}{Kr}}$ to be the equilibrium precision in the unique symmetric pure strategy equilibrium among competing news sources.

6 Comparison with Discrete Setting

Our continuous-time setting can be seen as the limit of a sequence of discrete-time models, which we studied in earlier work Liang et al. (2017). We expand on the connection and differences in this section.

Consider the following formulation: There are K unknown attribute values $\theta_1, \dots, \theta_K$ that are jointly normal, and a payoff-relevant state $\omega = \alpha' \cdot \theta$ with a known and positive weight vector α . Time is discrete with period length Δ , and at each time $t = 0, \Delta, 2\Delta, \dots$, the agent chooses from among K information sources. Choice of source i produces an independent observation of the signal

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

The agent also chooses when to stop acquiring information and face the decision problem.

This discrete-time model is a reformulation of Liang et al. (2017), and relates to our continuous-time setting here in the following way. Imagine that in our continuous-time model the agent is constrained to put all attention to one of the sources over each of the time intervals $[0, \Delta), [\Delta, 2\Delta)$, etc. Then at time t , choosing source i means he will observe the path of a diffusion process with drift θ_i and unit volatility, from time t to $t + \Delta$. As is well known, the difference between the values of this process at t and $t + \Delta$ is a sufficient statistic for learning about θ_i . Thus the agent’s information from time t to $t + \Delta$ is equivalent to a normal signal with mean $\theta_i \cdot \Delta$ and variance Δ , which is just $\Delta \cdot Y_i$ with Y_i given above.

A key difference is that the discrete model has non-divisible signals, and thus faces an “integer problem.” As $\Delta \rightarrow 0$, the long-run frequencies over signals in the discrete-time

³²This can happen if ρ is large or if the chosen signal noise levels $\sigma_1, \dots, \sigma_K$ are very different.

model converge to the optimal (instantaneous) attention allocation in continuous time, but the necessity of an integer approximation complicates characterization of the full sequence of information acquisitions. Studying the problem in continuous-time, as we do here, permits a sharper characterization both of the optimal information acquisition strategy itself, and of the conditions needed for this characterization to hold.

Another point of comparison is the notion of *myopic information acquisition*. In the discrete setting of [Liang et al. \(2017\)](#), we defined a signal choice strategy to be myopic if it maximizes the immediate reduction of posterior variance given any history. Our results there demonstrated conditions under which the myopic strategy coincides with the optimal strategy from the beginning, or eventually. The optimal attention allocation strategy for the continuous-time model, characterized in the present paper, has a similar flavor of being “myopic,” since at each time the agent divides attention across those sources that have the greatest marginal reduction of posterior variance. As we explained, however, the optimal allocation of attention is *not* pinned down by having the greatest marginal value at this moment. Rather, given a set of sources that maximize the marginal value at time t , the optimal mixture over them is chosen so that their marginal values remain equal at future instants.

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 3. *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,³³ 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

³³This 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 as³⁴

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

³⁴This 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 1 and 2 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 3. *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 3 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 3 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 x_1, x_2, y_1, y_2 as in in Theorem 1:

$$x_i = \alpha_i \det(\Sigma), \quad y_i = \alpha_1 \Sigma_{i1} + \alpha_2 \Sigma_{i2}.$$

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

$$\begin{aligned}
\gamma &:= (\Sigma^{-1} + Q)^{-1} \cdot \alpha \\
&= (\Sigma^{-1} \cdot (I + \Sigma Q))^{-1} \cdot \alpha \\
&= (I + \Sigma Q)^{-1} \cdot \Sigma \cdot \alpha \\
&= (I + \Sigma Q)^{-1} \cdot \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \\
&= \frac{1}{\det(I + \Sigma Q)} \begin{pmatrix} 1 + q_2 \Sigma_{22} & -q_2 \Sigma_{12} \\ -q_1 \Sigma_{21} & 1 + q_1 \Sigma_{11} \end{pmatrix} \cdot \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \\
&= \frac{1}{\det(I + \Sigma Q)} \begin{pmatrix} x_1 q_2 + y_1 \\ x_2 q_1 + y_2 \end{pmatrix}.
\end{aligned}$$

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

$$\begin{aligned}
\partial_1 V(q_1, q_2) &= \frac{-(x_1 q_2 + y_1)^2}{\det^2(I + \Sigma Q)}, \\
\partial_2 V(q_1, q_2) &= \frac{-(x_2 q_1 + y_2)^2}{\det^2(I + \Sigma Q)}.
\end{aligned} \tag{5}$$

Note that Assumption 1 translates into $y_1 + y_2 \geq 0$. 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 $y_1 \geq y_2$, then y_1 is non-negative. Let $t_1^* = \frac{y_1 - y_2}{x_2}$. Then when $q_1 + q_2 \leq t_1^*$ we always have

$$x_1 q_2 + y_1 \geq y_1 \geq x_2 q_1 + y_2,$$

since $x_1 q_2 \geq 0$ and $x_2 q_1 \leq x_2(q_1 + q_2) \leq x_2 t_1^* = y_1 - y_2$. We also have

$$x_1 q_2 + y_1 \geq -(x_2 q_1 + y_2),$$

since $x_1 q_2, x_2 q_1 \geq 0$ and by assumption $y_1 + y_2 \geq 0$. 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

$$n(t) = \left(\frac{x_1 t + y_1 - y_2}{x_1 + x_2}, \frac{x_2 t - y_1 + y_2}{x_1 + x_2} \right).$$

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 Necessity of Assumption 1

We show here that the assumption $y_1 + y_2 \geq 0$ is also *necessary* for the existence of a uniformly optimal strategy. The result generalizes Example 5 in the main text.

Proposition 4. *Suppose Assumption 1 is violated. Then a uniformly optimal strategy does not exist.*

Proof. Suppose that $y_1 + y_2 < 0$. First note that one of y_1, y_2 is positive, because $\alpha_1 y_1 + \alpha_2 y_2 = \alpha' \Sigma \alpha > 0$. So without loss we can assume $y_2 > 0 > -y_2 > y_1$. Moreover, from $\alpha_1 y_1 + \alpha_2 y_2 > 0$ we obtain $\alpha_2 > \alpha_1$ and hence $x_2 > x_1$. Below we characterize the t -optimal attention vector $n(t)$:

1. If $t \leq \frac{-(y_1+y_2)}{x_2}$, then $x_1 q_2 + y_1$ is negative and has larger absolute value than $x_2 q_1 + y_2$ (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 $\frac{-(y_1+y_2)}{x_2} < t < \frac{-(y_1+y_2)}{x_1}$, 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

$$x_1 n_2(t) + y_1 = -(x_2 n_1(t) + y_2),$$

where we use the fact that for t in this range, $x_1 q_2 + y_1$ is always negative. Together with $n_1(t) + n_2(t) = t$, we can solve that $n(t) = \left(\frac{-x_1 t - y_1 - y_2}{x_2 - x_1}, \frac{x_2 t + y_1 + y_2}{x_2 - x_1} \right)$.

3. If $\frac{-(y_1+y_2)}{x_1} \leq t \leq \frac{y_2 - y_1}{x_1}$, then $(x_2 q_1 + y_2)^2 - (x_1 q_2 + y_1)^2 = (y_2 - y_1 - x_1 q_2 + x_2 q_1) \cdot (y_1 + y_2 + x_1 q_2 + x_2 q_1) \geq 0$ 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 $t > \frac{y_2 - y_1}{x_1}$, 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

$$x_1 n_2(t) + y_1 = x_2 n_1(t) + y_2,$$

since both terms are now positive. This together with $n_1(t) + n_2(t) = t$ yields the solution $n(t) = \left(\frac{x_1 t + y_1 - y_2}{x_1 + x_2}, \frac{x_2 t - y_1 + y_2}{x_1 + x_2} \right)$ 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 .³⁵ 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 4. *The inverse of the prior covariance matrix Σ^{-1} is diagonally-dominant. That is,*

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

This is implied by Assumption 2 via the following lemma.

Lemma 9. *Suppose the prior covariance matrix Σ satisfies Assumption 2, 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 $x_j = [\Sigma^{-1}]_{1j}$ for $1 \leq j \leq K$, and without loss assume x_2 has the greatest absolute value among x_2, \dots, x_K . It suffices to show $x_1 \geq (K - 1)|x_2|$.

³⁵This 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.

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

$$|x_1 \cdot \Sigma_{21}| = |x_2 \cdot \Sigma_{22} + \sum_{j>2} x_j \cdot \Sigma_{2j}| \geq |x_2 \cdot \Sigma_{22}| - \sum_{j>2} |x_j \cdot \Sigma_{2j}| \geq |x_2 \cdot \Sigma_{22}| - \sum_{i>2} \frac{|x_2 \cdot \Sigma_{22}|}{2K-3},$$

where the last inequality uses $|x_j| \leq |x_2|$ and $|\Sigma_{2j}| \leq \frac{1}{2K-3} |\Sigma_{22}|$ for $j > 2$. The above inequality simplifies to

$$|x_1 \cdot \Sigma_{21}| \geq \frac{K-1}{2K-3} \cdot |x_2 \cdot \Sigma_{22}|.$$

And since $\Sigma_{21} \leq \frac{1}{2K-3} |\Sigma_{22}|$, we conclude that $|x_1| \geq (K-1)|x_2|$ as desired. Note that $x_1 = [\Sigma^{-1}]_{11}$ is necessarily positive, thus $x_1 \geq (K-1)|x_2|$. \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 4, 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*

$$n(t) = n(\underline{t}) + \frac{t}{\alpha_1 + \dots + \alpha_K} \cdot \alpha.^{36}$$

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 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.

³⁶That is, $n_i(t) = n_i(\underline{t}) + \frac{t}{\alpha_1 + \dots + \alpha_K} \cdot \alpha_i$ for each i .

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.³⁷ 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 $(\Sigma^{-1} + Q)^{-1}$.³⁸ 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

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

which indeed holds.³⁹

³⁷To 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.

³⁸This 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 $\tilde{Q}^{-1} - \tilde{Q}^{-1}(\tilde{Q}^{-1} + \Sigma_{TL})^{-1}\tilde{Q}^{-1} = ((\Sigma_{TL})^{-1} + \tilde{Q})^{-1}$.

³⁹Consider the identity $(\Sigma^{-1} + Q)^{-1} \cdot (\Sigma^{-1} + Q) = I_K$. The top-right block of the product is zeros, so by

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.

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.⁴⁰ 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 2 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 1. 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

⁴⁰It 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 $\frac{1}{K-1} \sum_{i>1} \theta_i + \mathcal{N}\left(0, \frac{1}{(K-1)q_2}\right)$.⁴¹

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 1. 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 2) and that Theorem 2 applies at any posterior belief after each source is

⁴¹This 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.⁴²

It remains to prove Lemma 2. 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 i .*

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 2.

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:⁴³

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

⁴²This argument shows that Assumption 4 is satisfied when each q_i is large. It can be shown that in fact, the stronger Assumption 2 is also satisfied if we take q_i even larger (i.e. Lemma 3 holds).

⁴³Such 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 $\tilde{q}_i(T) < \underline{q}$.*⁴⁴ 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

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

⁴⁴Formally, 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

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

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 $\partial_i V(q) > \frac{-C_L}{\underline{q}^2}$ 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 Proof of Proposition 2

We first consider pure strategy equilibria, and then use the constant-sum feature of the game to argue there are no mixed equilibria. Fix arbitrary $\sigma_1, \sigma_2 > 0$. From the agent's perspective, the informational environment is equivalent to one in which he seeks to predict $\sigma_1 \tilde{\theta}_1 + \sigma_2 \tilde{\theta}_2$ and holds the prior belief

$$\begin{pmatrix} \tilde{\theta}_1 \\ \tilde{\theta}_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \frac{\mu_1}{\sigma_1} \\ \frac{\mu_2}{\sigma_2} \end{pmatrix}, \begin{pmatrix} \frac{1}{\sigma_1^2} & \frac{\rho}{\sigma_1 \sigma_2} \\ \frac{\rho}{\sigma_1 \sigma_2} & \frac{1}{\sigma_2^2} \end{pmatrix} \right).$$

Since Assumption 1 is satisfied, we can apply Theorem 1 to this transformed environment.

Without loss of generality we assume $\sigma_1 \leq \sigma_2$ in equilibrium. Then the agent puts all attention on source 1 until time $t_1^* = \frac{(\sigma_2 - \sigma_1)\sigma_1}{1 - \rho}$. At all times after t_1^* , he allocates attention in the constant proportion $\left(\frac{\sigma_1}{\sigma_1 + \sigma_2}, \frac{\sigma_2}{\sigma_1 + \sigma_2} \right)$. Source 1's payoff function is thus

$$U_1(\sigma_1, \sigma_2) = \int_0^{t_1^*} e^{-rt} dt + \int_{t_1^*}^{\infty} e^{-rt} \frac{\sigma_1}{\sigma_1 + \sigma_2} dt = \frac{1}{r} \left(1 - \frac{\sigma_2}{\sigma_1 + \sigma_2} e^{-rt_1^*} \right).$$

The derivative with respect to the source's action σ_1 is

$$\left. \frac{\partial U_1}{\partial \sigma_1} \right|_{(\sigma_1, \sigma_2)} = \frac{\sigma_2}{r(\sigma_1 + \sigma_2)^2} e^{-rt_1^*} \left(1 - \frac{r(\sigma_1 + \sigma_2)(2\sigma_1 - \sigma_2)}{1 - \rho} \right). \quad (10)$$

Equilibrium requires

$$r(\sigma_1 + \sigma_2)(2\sigma_1 - \sigma_2) \leq 1 - \rho \quad \text{with equality if } \sigma_1 < \sigma_2. \quad (11)$$

On the other hand, since $\beta_1(t) + \beta_2(t) = 1$ at every t , the game has constant sum $\frac{1}{r}$. So source 2's payoff is simply

$$U_2(\sigma_1, \sigma_2) = \frac{1}{r} - U_1(\sigma_1, \sigma_2) = \frac{\sigma_2}{r(\sigma_1 + \sigma_2)} e^{-rt_1^*}.$$

The derivative with respect to its action σ_2 is

$$\left. \frac{\partial U_2}{\partial \sigma_2} \right|_{(\sigma_1, \sigma_2)} = \frac{\sigma_1}{r(\sigma_1 + \sigma_2)^2} e^{-rt_1^*} \left(1 - \frac{r(\sigma_1 + \sigma_2)\sigma_2}{1 - \rho} \right). \quad (12)$$

Equilibrium requires

$$r(\sigma_1 + \sigma_2)\sigma_2 \geq 1 - \rho \quad \text{with equality if } \sigma_1 < \sigma_2 \quad (13)$$

Combining (11) and (13), it is immediate that any pure strategy equilibrium must have $\sigma_1 = \sigma_2$.⁴⁵ Then the two inequalities (11) and (13) together give $\sigma_1 = \sigma_2 = \sqrt{\frac{1-\rho}{2r}} = \sigma^*$ as desired. Moreover, this is an equilibrium because (10) and (12) show that any deviation (not just local deviations) is not profitable. In fact, given $\sigma_j = \sigma^*$, the unique best response of source i is to choose the same σ_i . Since the game has a constant sum, this proves that the pure strategy equilibrium we have found is the unique equilibrium, pure or mixed.

G Many Competing Providers

Here we demonstrate how the game in Section 5 generalizes to the case of $K > 2$ competing data sources. We maintain essentially the same setup, except that the agent seeks to predict $\theta_1 + \dots + \theta_K$ where the precision of information about each θ_i is controlled by a separate data provider. Using the transformation $\tilde{\theta}_i = \frac{\theta_i}{\sigma_i}$, we can reduce the agent's information acquisition problem to our main model with prior covariance matrix

$$\tilde{\Sigma} = \begin{pmatrix} \frac{1}{\sigma_1^2} & \frac{\rho}{\sigma_1\sigma_2} & \cdots & \frac{\rho}{\sigma_1\sigma_K} \\ \frac{\rho}{\sigma_1\sigma_2} & \frac{1}{\sigma_2^2} & \cdots & \frac{\rho}{\sigma_2\sigma_K} \\ \cdots & \cdots & \cdots & \cdots \\ \frac{\rho}{\sigma_1\sigma_K} & \frac{\rho}{\sigma_2\sigma_K} & \cdots & \frac{1}{\sigma_K^2} \end{pmatrix}.$$

⁴⁵Otherwise both inequalities hold equal, which yields $2\sigma_1 - \sigma_2 = \sigma_2$ and again $\sigma_1 = \sigma_2$.

and weight vector $\tilde{\alpha} = (\sigma_1, \dots, \sigma_K)'$.

Although $\tilde{\Sigma}$ does not in general satisfy Assumption 2, 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 $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_K$. For $1 \leq k \leq K - 1$, define*

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

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 σ_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 σ_2 ; then, source 1's payoff when choosing $\sigma_1 \leq \sigma$ is given by

$$\frac{1}{r} \left(1 - \frac{(K-1)\sigma}{\sigma_1 + (K-1)\sigma} \cdot e^{\frac{-r\sigma_1(\sigma - \sigma_1)}{(1-\rho)}} \right).$$

Differentiating this w.r.t. σ_1 yields the first-order condition $r \cdot (\sigma_1 + (K-1)\sigma) \cdot (2\sigma_1 - \sigma) \leq 1 - \rho$ at $\sigma_1 = \sigma$, so that $\sigma \leq \sqrt{\frac{1-\rho}{Kr}}$.

On the other hand, by choosing $\sigma_1 > \sigma$, source 1 gets

$$\frac{\sigma_1}{\sigma_1 + (K-1)\sigma} \cdot e^{\frac{-r(K-1)\sigma(\sigma_1 - \sigma)}{1-\rho}}.$$

Differentiating w.r.t. σ_1 yields another first-order condition $r \cdot \sigma_1 \cdot (\sigma_1 + (K-1)\sigma) \geq 1 - \rho$ at $\sigma_1 = \sigma$. Thus $\sigma \geq \sqrt{\frac{1-\rho}{Kr}}$, 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

$$n_i(t) = \frac{\sigma_i(\sigma_k - \sigma_i)}{1 - \rho} + \frac{\sigma_i}{\sigma_1 + \dots + \sigma_k} \cdot (t - t_{k-1}), \quad \forall 1 \leq i \leq k \quad (14)$$

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

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

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

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

with Σ being the matrix having “1”s on the diagonal and “ ρ ” everywhere off the diagonal, and σ denoting the vector $(\sigma_1, \dots, \sigma_K)'$ (with a slight abuse of notation). From the above discussion, σ is also the weight vector $\tilde{\alpha}$.

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

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

where we use $\text{diag}(q/\sigma)$ to denote the diagonal matrix with entries $q_1/\sigma_1, \dots, q_K/\sigma_K$.

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

$$\sum_{j=1}^K M_{ij} \cdot \gamma_j = 1, \quad \forall i. \tag{15}$$

We will use these identities to show that each γ_j is positive and $\gamma_1 = \cdots = \gamma_k$ 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 $M = \Sigma^{-1} \cdot \text{diag}(\sigma) + \text{diag}(q/\sigma)$ we deduce

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

with $\delta_{j=i}$ representing the indicator function for the event $j = i$. Plugging this into (15), we then obtain

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

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 $(a-b)\sigma_i + \frac{q_i}{\sigma_i}$.

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 $(a - b)\sigma_i + \frac{q_i}{\sigma_i}$. From (14) and $a - b = \frac{1}{1-\rho}$ we indeed have

$$(a - b)\sigma_i + \frac{q_i}{\sigma_i} = \frac{1}{1 - \rho} \cdot \sigma_k + \frac{t - t_{k-1}}{\sigma_1 + \dots + \sigma_k}, \quad \forall 1 \leq i \leq k.$$

The RHS is the same for $i \leq k$ and smaller than $(a - b)\sigma_{k+1}$ 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. (2018): “Optimal Discovery and Influence Through Selective Sampling,” 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.
- BONATTI, A. AND G. CISTERNAS (2019): “Consumer Scores and Price Discrimination,” Working Paper.
- CALLANDER, S. (2011): “Searching and Learning by Trial and Error,” *American Economic Review*, 101, 2277–2308.
- CARLSON, D. AND T. L. MARKHAM (1979): “Schur Complements of Diagonally Dominant Matrices,” *Czechoslovak Mathematical Journal*, 29, 246–251.
- CHE, Y.-K. AND K. MIERENDORFF (2019): “Optimal Dynamic Allocation of Attention,” *American Economic Review*, forthcoming.
- 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,” *Inform 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*.
- GARFAGNINI, U. AND B. STRULOVICI (2016): “Social Experimentation with Interdependent and Expanding Technologies,” *Review of Economic Studies*, 83, 1579–1613.
- 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 (2019): “Complementary Information and Learning Traps,” Working Paper.
- 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.
- 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.