

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<https://www.econometricsociety.org/>

Econometrica, Vol. 90, No. 1 (January, 2022), 47–80

DYNAMICALLY AGGREGATING DIVERSE INFORMATION

ANNIE LIANG

Department of Economics and Department of Computer Science, Northwestern University

XIAOSHENG MU

Department of Economics, Princeton University

VASILIS SYRGKANIS

Microsoft Research

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

DYNAMICALLY AGGREGATING DIVERSE INFORMATION

ANNIE LIANG

Department of Economics and Department of Computer Science, Northwestern University

XIAOSHENG MU

Department of Economics, Princeton University

VASILIS SYRGKANIS

Microsoft Research

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

KEYWORDS: Information acquisition, dynamic Blackwell, binary choice.

1. INTRODUCTION

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

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

At every instant of continuous time, the decision-maker allocates a fixed budget of attention/resources across the information sources, where this allocation determines the precision of information extracted from each source. For example, the governor may have

Annie Liang: annie.liang@northwestern.edu

Xiaosheng Mu: xmu@princeton.edu

Vasilis Syrgkanis: vasy@microsoft.com

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

a limited number of tests to allocate across testing centers each day, where more tests lead to a more precise estimate of the incidence rate at that testing center. The decision-maker acquires information until an endogenously chosen stopping time, at which point he makes a decision whose payoff depends on the unknown state.

Our contribution is to demonstrate that the optimal dynamic acquisition strategy can be explicitly characterized under certain conditions on the prior belief, and to explain what those conditions are. Under the optimal strategy, the decision-maker initially exclusively acquires information from the single most informative source, where “more informative” is evaluated with respect to his prior belief over the unknown attribute values. At fixed times, the decision-maker begins learning from additional sources, and divides attention over these new sources as well as the ones he was learning from previously. Eventually, the decision-maker acquires information from all sources using a final and constant mixture. Notably, the optimal information acquisition strategy is not only history-independent but also robust across all decision problems. This implies, for example, that one does not need to solve for the optimal stopping time and information acquisition strategy jointly in this problem—optimal information acquisition is independent of when the decision-maker stops. We make use of this implication in Section 5.1 to derive new results about the optimal stopping behavior in a binary choice problem.

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

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

The question of whether a uniformly optimal strategy exists turns out to relate to a classic problem in consumer theory regarding the normality of demand—that is, whether a consumer’s demand for various goods is weakly increasing in income. In our setting, the decision-maker’s “utility function” is the negative of the posterior variance function, and his “income” is the budget constraint on attention. When a uniformly optimal strategy exists, this means that the decision-maker’s demand for information from each source is weakly increasing in his overall attention budget. One of our sufficient conditions for existence of a uniformly optimal strategy—“perpetual complementarity” of the different sources—directly connects to a known sufficient condition in the literature for normality of demand. We additionally utilize the specific structure of our problem to provide two new sufficient conditions, and show that all of these conditions can be simply stated in terms of the decision-maker’s prior belief. See Section 4 for further details.

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

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

Application 2: Attention Manipulation. Next, we use our framework to study the dynamic consequences of a one-time attention manipulation. Recently, [Gossner, Steiner, and Stewart \(2021\)](#) studied this question in a model where a decision-maker chooses between goods with independent payoffs. Under some assumptions, they showed that a one-time manipulation of attention towards a given good leads to persistently higher cumulative attention devoted to that good, and persistently lower cumulative attention to every other good. We derive a complementary result in our setting, focusing on how correlation across the unknown attributes affects the consequences of attention manipulation. We show that with two sources, [Gossner, Steiner, and Stewart's \(2021\)](#) insights hold under flexible patterns of correlation. On the other hand, with more than two sources, the nature of correlation matters. We identify a property of the prior belief under which all sources provide substitutable information, and show that under this property (but not in general), attention manipulation leads to persistently higher attention for that source and lower attention for others.

Application 3: Biased News Sources. In our final application, we consider a stylized game between a liberal and a conservative news source that report on a common unknown (e.g., the fiscal cost of a policy proposal), where their reporting is biased in opposite directions. The sources choose the size of their bias, as well as the informativeness of their reporting, and compete over readers' attention. Using our characterization of information acquisition, we are able to derive the complete time path of readers' attention allocations given any precision and bias choices by the sources, which allows us to characterize equilibrium news provision in this model. One particular insight that emerges from

this analysis is that incentives for bias not only lead to greater polarization in equilibrium, but also lead to lower-quality news provision (i.e., larger noise choices). This analysis contributes to a literature about how competition across news sources affects the quality of news (Gentzkow and Shapiro (2008), Galperti and Trevino (2020), Chen and Suen (2019), Perego and Yuksel (2021)), where our work is distinguished in considering the role of the time path of information demand.

1.1. *Related Literature*

Our work builds on a large literature regarding dynamic acquisition of information. One part of this literature considers choice between unconstrained information structures at entropic (or more generally, “posterior-separable”) costs; see, for example, Yang (2015), Steiner, Stewart, and Matějka (2017), Hébert and Woodford (2021), Morris and Strack (2019), and Zhong (2019).¹ Under this modeling approach, the cost to acquiring information depends on the decision-maker’s current belief, and is often interpreted as a mental processing cost (Mackowiak, Matějka, and Wiederholt (2021)). In contrast, a second set of papers—to which our paper belongs—models agents as dynamically allocating a fixed budget of resources across a prescribed (and finite) set of experiments; see, for example, Che and Mierendorff (2019), Mayskaya (2020), Fudenberg, Strack, and Strzalecki (2018), Gossner, Steiner, and Stewart (2021), and Azevedo, Deng, Olea, Rao, and Weyl (2020). These papers, and ours, assume that the cost of information is independent of what the decision-maker currently knows. We view such information costs as a better match for applications in which the cost to information acquisition is physical, for example, a limit on the number of available COVID tests that can be administered in a given day.

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

Our work additionally connects to a large literature on sequential sampling in statistics and operations research. Since the information acquisition decisions in our model

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

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

are not directly linked to flow payoffs, our model does not fall under the classic multi-armed bandit framework (Gittins (1979), Bergemann and Välimäki (2008)). This feature also distinguishes our results relative to a classic literature on “learning by experimentation” (Easley and Kiefer (1988), Aghion, Bolton, Harris, and Jullien (1991), Keller, Rady, and Cripps (2005)). The “best-arm identification” problem (Bubeck, Munos, and Stoltz (2009), Russo (2016)) is more closely related to us, as it considers a decision-maker who samples for a number of periods before selecting an arm and receiving its payoff. Indeed, the special case of two arms with jointly normal payoffs is nested in our framework under the case of two attributes and equal payoff weights. Our Theorem 1 thus builds on a prior result of Frazier, Powell, and Dayanik (2008), which showed that myopic information acquisition—or the “knowledge gradient” policy in the language of that literature—is optimal when the two arms have independent normal payoffs. Our result generalizes Frazier, Powell, and Dayanik’s (2008) result by allowing for correlated payoffs and a broader class of decision problems.

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

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

2. MODEL

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

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

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

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

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

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

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

$$dX_i^t = \beta_i(t) \cdot \theta_i \cdot dt + \sqrt{\beta_i(t)} \cdot dB_i^t. \quad (1)$$

Above, each B_i is an independent Brownian motion, and the term $\sqrt{\beta_i(t)}$ is a standard normalizing factor to ensure constant informativeness per unit of attention devoted to each source. Thus, devoting T units of time to observing source i is equivalent to observing the normal signal $\theta_i + \mathcal{N}(0, 1/T)$, or receiving T independent observations of the standard normal signal $\theta_i + \mathcal{N}(0, 1)$. This formulation treats “attention” and “time” in the same way, in the sense that devoting $1/2$ attention to source i for a unit of time provides the same amount of information about θ_i as devoting full attention to source i for a $1/2$ unit of time. We also note that since all sources are assumed to be equally informative about their corresponding attributes, it is *with loss* to further normalize the payoff weights α_i to be equal to one another.⁵

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

Let $(\Omega, \mathbb{P}, \{\mathcal{F}_t\}_{t \in \mathbb{R}_+})$ describe the relevant probability space, where the information \mathcal{F}_t that the agent observes up to time t is the collection of paths $\{X_i^{\leq t}\}_{i=1}^K$ (with $X_i^{\leq t}$ representing the sample path of X_i from time 0 to time t). An *information acquisition strategy* S is a map from $\{X_i^{\leq t}\}_{i,t}$ into $\Delta(\{1, \dots, K\})$, representing how the agent divides attention at

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

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

each instant of time as a function of the observed diffusion processes.⁶ In addition to allocating his attention, the agent chooses how long to acquire information for; that is, at each instant of time, he determines (based on the history of observations) whether to continue acquiring information, or to stop and take an action. Formally, the agent chooses a *stopping time* τ , which is a map from Ω into $[0, +\infty]$ satisfying the measurability requirement $\{\tau \leq t\} \in \mathcal{F}_t$ for all t .

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

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

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

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

3. MAIN RESULTS

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

3.1. Two Attributes

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

ASSUMPTION 3: $\text{cov}_1 + \text{cov}_2 = \alpha_1(\Sigma_{11} + \Sigma_{12}) + \alpha_2(\Sigma_{21} + \Sigma_{22}) \geq 0$.

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

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

THEOREM 1: *Suppose $K = 2$ and Assumption 3 is satisfied. Without loss of generality, let $\text{cov}_i \geq \text{cov}_j$. Define*

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

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

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

There are two stages of information acquisition. In the first stage, which ends at some time t_i^* , the agent allocates all of his attention to the attribute i with higher prior covariance with the payoff-relevant state. After time t_i^* , he divides his attention across the attributes in a constant ratio, where the long-run instantaneous attention allocation is proportional to the weights α . Note that depending on the agent's stopping rule, Stage 2 of information acquisition may never be reached along some histories of the diffusion processes. But as long as the agent continues acquiring information, his attention allocations are as given above. In Appendix O.1 of the Supplemental Material (Liang, Mu, and Syrgkanis (2022)), we show that under mild technical assumptions, the optimal attention strategy is in fact unique up to the stopping time τ (after which attention allocations obviously do not matter).

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

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

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

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

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

3.2. K Attributes

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

ASSUMPTION 4—Perpetual Substitutes: Σ^{-1} has non-positive off-diagonal entries.

ASSUMPTION 5—Perpetual Complements: Σ has non-positive off-diagonal entries and $\Sigma \cdot \alpha$ has non-negative coordinates.

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

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

Assumption 5 imposes that all attributes are negatively correlated with one another, but that each attribute is initially positively correlated with the payoff-relevant state. Under this assumption, prior covariances are “mildly” negative. Proposition 8 in Appendix O.3

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

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

of the Supplemental Material shows that Assumption 5 is an if and only if condition for information from any pair of sources to be perpetually complementary, in the sense that the value of acquiring information from one source is increasing in the amount of information from the other source, again along the entire path of information acquisitions.

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

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

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

and nested sets

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

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

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

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

Finally, we mention that starting from any prior belief (including those that fail all three of the sufficient conditions we have provided), so long as the agent does not stop

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

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

learning about any attribute, his posterior belief must eventually satisfy Assumption 6.¹¹ Theorem 2 applies at these posterior beliefs, implying in particular that the agent's optimal attention allocation is eventually constant and proportional to the weight vector α .

4. EXPLANATION OF RESULTS

4.1. Fixed Stopping Time t

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

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

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

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

Define the *t-optimal* attention vector

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

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

4.2. Uniformly Optimal Strategy: Definition

In general, the family of solutions $\{n(t)\}_{t \geq 0}$ corresponding to optimal allocation of t units of attention does not determine the solution to the dynamic problem. To see this, suppose $n(1) = (0, 1)$, so that the optimal way to allocate attention for a decision at time $t = 1$ is to allocate it all to the second attribute, but $n(2) = (2, 0)$, so that the optimal way to allocate attention for a decision at time $t = 2$ is to allocate it all to the first attribute. Clearly, these *t-optimal* attention vectors are incompatible under a single dynamic strategy: Optimal attention allocation for a decision at time $t = 1$ precludes achieving the optimal attention allocation for time $t = 2$.¹² So when the agent faces the possibility of

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

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

stopping at either time $t = 1$ or $t = 2$ (depending on the realizations of the diffusion processes), he must trade off between achieving more precise information about the payoff-relevant state at either time. The optimal resolution of this trade-off depends on the specific decision problem.

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

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

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

Thus, whenever a uniformly optimal strategy exists, it must be an optimal strategy in our problem.¹³ It remains to show that under the assumptions we provided, a uniformly optimal strategy does exist, and has the structure described in Theorems 1 and 2.

4.3. Uniformly Optimal Strategy: Existence

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

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

Whether $n(t)$ is monotone turns out to be closely related to a classic problem in consumer theory: Suppose a consumer has a utility function $U(q_1, \dots, q_K)$ over consumption of q_k units of each of K goods, and let $D(\mathbf{p}, w)$ denote his Marshallian demand subject to the budget constraint $\mathbf{p} \cdot \mathbf{q} \leq w$. Then, the consumer’s demand is *normal* if each coordinate of $D(\mathbf{p}, w)$ increases with income w . In our setting, we can set $U = -V(q_1, \dots, q_K)$

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

to be the negative of the posterior variance, so that minimizing V is the same as maximizing U . Our t -optimal attention vector $n(t)$ is then precisely the Marshallian demand, when prices are identically equal to 1 and income is equal to t . Thus, normality of the consumer's demand under utility function $U = -V$ is equivalent to existence of a uniformly optimal strategy.

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

The key difference between these conditions and ours is that the known conditions apply to the utility function U . In our environment, $U = -V = -\alpha'(\Sigma^{-1} + \text{diag}(q))^{-1}\alpha$ is a composite function of the primitive objects Σ (prior belief) and α (payoff weights). Our Assumptions 3–6 apply to these primitives, which have no natural analogue in the consumer demand problem. Moreover, the conditions in the literature are stated for the utility function U evaluated at *all* demand levels q . In contrast, we show it is possible to reduce such conditions to conditions on Σ and α only, which are much easier to verify.

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

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

4.4. Uniformly Optimal Strategy: Structure

When $n(t)$ is indeed monotone in t , the uniformly optimal strategy that achieves these vectors is simply the one that sets each attention allocation $\beta(t)$ to be the time-derivative of $n(t)$. Under this strategy, the agent divides attention at every moment across those

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

attributes that maximize the *instantaneous* marginal reduction of posterior variance V . While the set of attributes is pinned down by first-order conditions, the precise mixture over those attributes is pinned down by second-order conditions, which ensure that this set of attributes continue to have equal and marginal values at future instants. Specifically, for each set of attributes, there is a unique linear combination that corresponds to the “learnable component” of ω given those attributes (formally, it is a projection of ω). It turns out that at each stage, it is optimal to acquire information in a mixture proportional to the weights of this linear combination, thus producing an unbiased signal about the learnable component of ω . In the final stage, when the agent pays attention to every attribute—so that the learnable component of ω is ω itself—the optimal mixture is proportional to the payoff weights α .

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

5. APPLICATIONS

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

In Section 5.1, we consider optimal attention allocation for choice between two goods and generalize recent results from Fudenberg, Strack, and Strzalecki (2018). In Section 5.2, we consider the dynamic implications of a one-time attention manipulation, complementing a recent exercise in Gossner, Steiner, and Stewart (2021). In Section 5.3, we develop a game between biased news sources, and characterize the degree of polarization and the quality of information in equilibrium. The applications in Sections 5.1 and 5.2 show that we can use our main results to tractably introduce correlation in settings that have been previously studied under strong assumptions of independence. Our applications in Sections 5.1 and 5.3 show how our results about information acquisition can be used as an intermediate step to derive results about other economic behaviors.

5.1. Application 1: Binary Choice

Building on a large literature regarding “binary choice” problems, Fudenberg, Strack, and Strzalecki (2018) (henceforth FSS) recently proposed the *uncertain drift-diffusion* model: An agent has a choice between two goods with unknown payoffs v_1 and v_2 , and can learn about those payoffs before making a choice. The two payoffs are jointly normal and i.i.d.:

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

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

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

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

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

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

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

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

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

- Stage 1: *At all times*

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

the agent optimally allocates all attention to θ_1 .

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

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

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

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

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

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

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

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

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

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

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

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

5.2. Application 2: Attention Manipulation

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

Gossner, Steiner, and Stewart (2021) (henceforth GSS) recently studied this question in a model in which an agent sequentially learns about the quality of a number of goods by allocating attention to one good each period. One of their main results (Theorem 1) resolves the two questions posed above in the following way:

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

A key assumption in GSS is that the attention strategy used by the agent satisfies a version of Independence of Irrelevant Alternatives (IIA): Conditional on *not* focusing on the good to which attention is diverted, the agent’s belief about that good does not affect the relative probabilities of focusing on the remaining goods. Proposition 5 in Gossner, Steiner, and Stewart (2021) shows that when the agent adopts a class of “satisficing” stopping rules, the optimal attention strategy satisfies IIA for goods with independent values.¹⁶

We can use our framework and main characterization to study a related but different problem, where the agent learns about multiple attributes of an unknown (one-dimensional) payoff-relevant state, and—importantly—those attribute values can be correlated. Additionally, we differ from GSS by focusing on the optimal attention allocation strategy and how it is affected by attention manipulation. Outside of the special case of independent attributes, the optimal strategy in our setting fails IIA when there are more than two attributes. Nevertheless, we show that GSS’s finding in (1) holds for flexible patterns of correlation, and the finding in (2) holds under an additional condition that we provide.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5.3. Application 3: Biased News Sources

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

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

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

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

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

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

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

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

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

The first part, $\int_0^\infty r e^{-rt} \beta_i^{(\phi, \zeta)}(t) dt$, is the discounted average attention paid to source i given the common discount rate r ,²⁰ which we interpret as a reduced form for advertising revenue. The second part, $-\lambda(1 - \phi_i)^2$, penalizes the news source for distance from the ideal bias for its party. This part of the payoff function is maximized by choosing $\phi_i = 1$, in which case source 1's signal is centered at $\omega + b$ and source 2's signal is centered at $\omega - b$. The parameter $\lambda \in \mathbb{R}_+$ determines the strength of incentive for bias, relative to the incentive for viewership.

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

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

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

and

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

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

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

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

²¹This assumption is imposed to guarantee the existence of a pure strategy equilibrium. Our analysis shows that $\lambda \geq 1.6$ is sufficient for existence, whereas a weaker condition $\lambda \geq \frac{9}{10}$ is necessary.

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

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

COROLLARY 2—Informativeness of News: *The equilibrium noise level, $\zeta^* = \frac{\sigma_b}{2\sqrt{r}}(1 + \sqrt{1 - \frac{1}{2\lambda}})$,*

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

Part (a) says that incentives for greater bias not only increase polarization, which is expected, but also lead to a reduction in the quality of news. To understand this result, consider the incentives for source i 's choice of precision. Applying our characterization in Theorem 1, there are up to two stages of information acquisition: In Stage 1, if there is a strictly more informative source, then that source receives all viewership; in Stage 2, both sources receive a constant proportion of viewership. We show that, for any equal bias intensity choices $\phi_1 = \phi_2$, source i 's long-run share is $\frac{\zeta_i}{\zeta_1 + \zeta_2}$, while source i is chosen in Stage 1 if and only if its noise term is smaller ($\zeta_i < \zeta_{3-i}$). Thus, sources face a trade-off between optimizing for greater long-run viewership—where a larger noise choice ζ_i increases the long-run share—versus competing to be chosen in the short run—which encourages smaller ζ_i . Intuitively, more precise information improves the competitive value of a source at the beginning of time, but reduces the value of continual engagement with that source. In equilibrium, sources choose the same ζ_i , thus washing out the first stage of information acquisition.

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

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

Part (c) of Corollary 2 holds by very similar reasoning. Less patient news sources compete over short-run profits (i.e., being chosen in Stage 1), and thus prefer precise signals, while patient sources compete for long-run profits (i.e., long-run proportion), and thus

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

prefer imprecise signals. So the less patient the sources are (larger r), the more precise their signals will be in equilibrium (smaller ζ^*).

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

6. DISCUSSION

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

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

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

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

Intertemporal Decision Problems. Our main model assumes that the agent takes only one action, which simplifies the exposition. But since our analysis based on the notion of uniform optimality is independent of details of the payoff function, it can be easily generalized to a setting where the agent takes N actions a_1, \dots, a_N at times $\tau_1 \leq \dots \leq \tau_N$. Our characterization of the optimal attention strategy extends for any (intertemporal) payoff function $u(\tau_1, \dots, \tau_N, a_1, \dots, a_N, \omega)$ that is decreasing in the decision times τ_1, \dots, τ_N .

APPENDIX A: PRELIMINARIES

A.1. Posterior Variance Function

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

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

LEMMA 2: *It holds that*

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

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

This function V extends to a rational function (quotient of polynomials) over all of \mathbb{R}^K ; that is, even if some q_i values 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 . Q.E.D.

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

LEMMA 3: *Given a cumulative attention vector $q \geq 0$, define*

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

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

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

PROOF: From Lemma 2 and the formula for matrix derivatives, we have

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

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

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

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

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

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

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

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

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

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

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

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

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

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

A.2. Optimality and Uniform Optimality

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

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

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

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

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

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

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

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

To formulate the resulting optimization problem, for each $\nu \in [0, v_0]$ we denote by $T^*(\nu)$ the time t such that $v_t^* = v_0 - \nu$; T^* is a deterministic and increasing function of ν . Then, under the attention strategy S^* , the agent's optimal payoff can be rewritten as

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

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

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

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

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

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

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

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

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

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

We also have a simple converse result:

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

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

A.3. Sufficient Condition for Assumption 6

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

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

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

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

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

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

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

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

APPENDIX B: PROOF OF THEOREM 1

Define $\text{cov}_1, \text{cov}_2$ as in the statement of Theorem 1, and define $x_i = \alpha_i \det(\Sigma)$ to ease notation.

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

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

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

$$\partial_1 V(q_1, q_2) = \frac{-(x_1 q_2 + \text{cov}_1)^2}{2 \det(I + \Sigma Q)}; \quad \partial_2 V(q_1, q_2) = \frac{-(x_2 q_1 + \text{cov}_2)^2}{2 \det(I + \Sigma Q)}. \quad (10)$$

Note that Assumption 3 translates into $\text{cov}_1 + \text{cov}_2 \geq 0$. Under this assumption, we will characterize the t -optimal vector $(n_1(t), n_2(t))$ and show it is increasing over time. Without loss assume $\text{cov}_1 \geq \text{cov}_2$; then cov_1 is non-negative. Let $t_1^* = \frac{\text{cov}_1 - \text{cov}_2}{x_2}$. Then, when

$q_1 + q_2 \leq t_1^*$, we always have

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

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

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

since $x_1 q_2, x_2 q_1 \geq 0$ and by assumption $\text{cov}_1 + \text{cov}_2 \geq 0$. Thus, (10) implies that $\partial_1 V(q_1, q_2) \leq \partial_2 V(q_1, q_2)$ at such attention vectors q . So for any budget of attention $t \leq t_1^*$, putting all attention to source 1 minimizes the posterior variance V . That is, $n(t) = (t, 0)$ for $t \leq t_1^*$.

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

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

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

APPENDIX C: PROOF OF THEOREM 2

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

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

C.1. Technical Property of γ

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

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

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

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

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

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

Rearranging, we then have

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

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

C.2. The Last Stage

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

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

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

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

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

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

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

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

Hence $\partial_1 V = \dots = \partial_K V$ continues to hold, completing the proof. *Q.E.D.*

C.3. Earlier Stages

In general, we need to show that even when the agent is choosing from a subset of the sources, the t -optimal vector $n(t)$ is still increasing over time. This is guaranteed by the

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

following lemma, which says that the agent optimally attends to those sources that maximize the marginal reduction of V , until a new source becomes another maximizer. For ease of exposition, we work under the stronger assumption that Σ^{-1} is *strictly* diagonally-dominant, in the sense that $[\Sigma^{-1}]_{ii} > \sum_{j \neq i} |[\Sigma^{-1}]_{ij}|$ for all $1 \leq i \leq K$. Later, we discuss how the lemma should be modified without this strictness.

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

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

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 . Moreover, when $|B| < K$, it holds that $\bar{t} < \infty$, whereas when $|B| = K$, it holds that $\bar{t} = \infty$ and β is proportional to α .

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

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

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

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

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

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

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

Thus, in this smaller problem, the prior covariance matrix is Σ_{TL} and the payoff weights are $\tilde{\alpha}$. Assuming that $\tilde{\alpha}$ has strictly positive coordinates, we can then apply Lemma 9: As

long as the agent attends to the first k sources proportional to $\tilde{\alpha}$, $\partial_1 V = \dots = \partial_k V$ continues to hold.²⁵ Moreover, at $q = n(\underline{t})$, the definition of the set B implies that these k partial derivatives have greater magnitude (i.e., more negative) than the rest. By continuity, the same comparison holds until some time $\bar{t} > \underline{t}$. Thus, when $t \in [\underline{t}, \bar{t})$, the cumulative attention vector (under this strategy) still satisfies the first-order condition $B = \operatorname{argmin}_{1 \leq i \leq K} \partial_i V$ and $q_j = 0$ for $j \notin B$. Since V is convex, this must be the t -optimal vector as we desire to show. It also follows that $\bar{t} < \infty$, because at $t = \infty$ the minimum possible posterior variance is zero, which cannot be achieved by attending only to a subset B of sources.

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

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

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

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

C.4. Piecing Together Different Stages

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

C.5. The Case of Weak Diagonal Dominance

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

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

sampled, which would violate the requirement of constant attention allocation for a given observation set.

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

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

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

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

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

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

C.6. The Case of Perpetual Substitutes or Perpetual Complements

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

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

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

If instead Assumption 5 is satisfied, then Σ itself is an M -matrix, and so is the principal sub-matrix Σ_{TL} . Thus, $(\Sigma_{\text{TL}})^{-1}$ has non-negative entries off the diagonal and strictly

positive entries on the diagonal. From (12), we have $\tilde{\alpha} = ((\Sigma_{\text{TL}})^{-1} + \tilde{Q}) \cdot \tilde{\gamma}$. Since $\tilde{\gamma}$ is a positive vector (with equal coordinates), we deduce $\tilde{\alpha} \gg 0$. This completes the proof of Theorem 2.

APPENDIX D: ALGORITHM FOR COMPUTING THE OPTIMAL STRATEGY

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

REFERENCES

- AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN (1991): "Optimal Learning by Experimentation," *Review of Economic Studies*, 58, 621–654. [51]
- ALARIE, Y., AND C. BRONSARD (1990): "Preferences and Normal Goods: A Necessary and Sufficient Condition," *Journal of Economic Theory*, 51, 423–430. [59]
- AZEVEDO, E., A. DENG, J. L. M. OLEA, J. RAO, AND G. WEYL (2020): "A/B Testing With Fat Tails," *Journal of Political Economy*, 128, 4614–4672. [50]
- BARDHI, A. (2020): "Attributes: Selective Learning and Influence," Working Paper. [50]
- 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. [51]
- BILANCINI, E., AND L. BONCINELLI (2010): "Preferences and Normal Goods: An Easy-to-Check Necessary and Sufficient Condition," *Economics Letters*, 108, 13–15. [59]
- BLACKWELL, D. (1951): "Comparison of Experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley and Los Angeles: University of California Press. [48,51]
- BUBECK, S., R. MUNOS, AND G. STOLTZ (2009): "Pure Exploration in Multi-Armed Bandits Problems," in *Algorithmic Learning Theory. ALT 2009. Lecture Notes in Computer Science*, Vol. 5809. Berlin, Heidelberg: Springer. [51]
- CALLANDER, S. (2011): "Searching and Learning by Trial and Error," *American Economic Review*, 101, 2277–2308. [50]
- CARLSON, D., AND T. L. MARKHAM (1979): "Schur Complements of Diagonally Dominant Matrices," *Czechoslovak Mathematical Journal*, 29, 246–251. [76]
- CHE, Y.-K., AND K. MIERENDORFF (2019): "Optimal Dynamic Allocation of Attention," *American Economic Review*, 108, 2993–3029. [50]
- CHEN, H., AND W. SUEN (2019): "Competition for Attention and News Quality," Working Paper. [50]
- CHICK, S. E., AND P. FRAZIER (2012): "Sequential Sampling With Economics of Selection Procedures," *Management Science*, 58, 550–569. [51]
- CHIPMAN, J. S. (1977): "An Empirical Implication of Auspitz–Lieben–Edgeworth–Pareto Complementarity," *Journal of Economic Theory*, 14, 228–231. [59]
- EASLEY, D., AND N. M. KIEFER (1988): "Controlling a Stochastic Process With Unknown Parameters," *Econometrica*, 56, 1045–1064. [51]
- 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. [51]
- (2009): "The Knowledge-Gradient Policy for Correlated Normal Beliefs," *Informs Journal on Computing*, 21, 599–613. [51]
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2018): "Speed, Accuracy, and the Optimal Timing of Choices," *American Economic Review*, 108, 3651–3684. [49,50,60]
- GALPERTI, S., AND I. TREVINO (2020): "Coordination Motives and Competition for Attention in Information Markets," *Journal of Economic Theory*, 188. [50]
- GARFAGNINI, U., AND B. STRULOVICI (2016): "Social Experimentation With Interdependent and Expanding Technologies," *Review of Economic Studies*, 83, 1579–1613. [50]
- GENTZKOW, M., AND J. SHAPIRO (2008): "Competition and Truth in the Market for News," *Journal of Economic Perspectives*, 22, 133–154. [50]
- GITTINS, J. C. (1979): "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society, Series B*, 148–177. [51]
- GOSSNER, O., J. STEINER, AND C. STEWART (2021): "Attention Please!" *Econometrica*, 89, 1717–1751. [49,50,60,63]
- GREENSHTEIN, E. (1996): "Comparison of Sequential Experiments," *The Annals of Statistics*, 24, 436–448. [48, 51,58]
- HANSEN, O. H., AND E. N. TORGERSEN (1974): "Comparison of Linear Normal Experiments," *The Annals of Statistics*, 2, 367–373. [48,51,57]
- HÉBERT, B., AND M. WOODFORD (2021): "Rational Inattention When Decisions Take Time," Working Paper. [50]
- KE, T. T., AND J. M. VILLAS-BOAS (2019): "Optimal Learning Before Choice," *Journal of Economic Theory*, 180, 383–437. [51]
- KELLER, G., S. RADY, AND M. CRIPPS (2005): "Strategic Experimentation With Exponential Bandits," *Econometrica*, 73, 2787–2802. [51]
- KLABJAN, D., W. OLSZWSKI, AND A. WOLINSKY (2014): "Attributes," *Games and Economic Behavior*, 88, 190–206. [50]

- LIANG, A., X. MU, AND V. SYRGKANIS (2019): “Optimal and Myopic Information Acquisition,” Working Paper. [68]
- (2022): “Supplement to ‘Dynamically Aggregating Diverse Information’,” *Econometrica Supplemental Material*, 90, <https://doi.org/10.3982/ECTA18324>. [54]
- MACKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2021): “Rational Inattention: A Review,” *Journal of Economic Literature*. (forthcoming). [50]
- MAYSKAYA, T. (2020): “Dynamic Choice of Information Sources,” Working Paper. [50]
- MORRIS, S., AND P. STRACK (2019): “The Wald Problem and the Relation of Sequential Sampling and Ex-ante Information Costs,” Working Paper. [50]
- NIEUWERBURGH, S. V., AND L. VELDKAMP (2010): “Information Acquisition and Under-Diversification,” *Review of Economic Studies*, 77, 779–805. [57]
- PEREGO, J., AND S. YUKSEL (2021): “Media Competition and Social Disagreement,” *Econometrica*. (forthcoming). [50]
- PLEMMONS, R. (1977): “M-Matrix Characterizations. I—Nonsingular M-Matrices,” *Linear Algebra and Its Applications*, 18, 175–188. [55]
- QUAH, J. K.-H. (2007): “The Comparative Statics of Constrained Optimization Problems,” *Econometrica*, 75, 401–431. [59]
- RATCLIFF, R., AND G. MCKOON (2008): “The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks,” *Neural Computation*, 20, 873–922. [49]
- REVUZ, D., AND M. YOR (1999): *Continuous Martingales and Brownian Motion*. New York: Springer. [71]
- RUSSO, D. (2016): “Simple Bayesian Algorithms for Best-Arm Identification,” *Operations Research*, 68, 1625–1647. [51]
- SANJURJO, A. (2017): “Search With Multiple Attributes: Theory and Empirics,” *Games and Economic Behavior*, 104, 535–562. [50]
- STEINER, J., C. STEWART, AND F. MATĚJKA (2017): “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” *Econometrica*, 85, 521–553. [50]
- YANG, M. (2015): “Coordination With Flexible Information Acquisition,” *Journal of Economic Theory*, 158, 721–738. [50]
- YONG, J., AND X. Y. ZHOU (1999): *Stochastic Controls: Hamiltonian Systems and HJB Equations*. New York: Springer. [53]
- ZHONG, W. (2019): “Optimal Dynamic Information Acquisition,” Working Paper. [50]

Guest Co-editor Marzena J. Rostek handled this manuscript.

Manuscript received 7 April, 2020; final version accepted 3 June, 2021; available online 4 June, 2021.