PIER
PENN INSTITUTE for ECONOMIC RESEARCH
UNIVERSITY of PENNSYLVANIA

The Ronald O. Perelman Center for Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper
19-005

Dynamically Aggregating Diverse Information

ANNIE LIANG
University of Pennsylvania

XIAOSHENG MU
Cowles Foundation &
Columbia University

VASILIS SYRGKANIS
Microsoft Research

April 15, 2019

Dynamically Aggregating Diverse Information[*][†]

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

April 15, 2019

Abstract

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

[*]This paper grew out of and subsumes part of our working paper "Optimal and Myopic Information Acquisition." The current paper considers a continuous-time formulation and characterizes the optimal information acquisition strategy, while "Optimal and Myopic Information Acquisition" focuses on the myopic information acquisition strategy and studies its optimality properties in discrete time.

[†]We are grateful to Drew Fudenberg and Peter Norman Sørensen for helpful comments and suggestions.

[‡]Department of Economics, University of Pennsylvania

[§]Cowles Foundation and Columbia University

[¶]Microsoft Research

1 Introduction

Markets are increasingly saturated with numerous and diverse sources of information about consumers. Determining how to acquire information in these environments thus often involves a question of how to aggregate information across different sources. For example, suppose a hotel chain wants to forecast demand for a new location in Puerto Rico. Multiple data sets may be relevant to this forecasting problem—e.g. the hotel can acquire data on discretionary spending from US-based credit card companies to predict vacation travel from the United States, or acquire search trend data from Google to predict travel demand from Mexico, or acquire demographic data from data brokers to predict the volume of business travel to this location. The firm chooses which sources to acquire data from, and also how to allocate potentially limited time and resources across them. If additionally it can supplement current data with more data in the future, then the information acquisition decisions should also take this into account.

In this work, we consider a firm that has access to various *data sources*, each modeled as a Brownian motion whose drift is an unknown *attribute* that the data source provides information about. The firm can continuously allocate a budget of resources (e.g. employee hours) across these Brownian motions, where more resources allocated to any data source results in greater precision of information about the corresponding attribute value.[1] The firm acquires information until an endogenously chosen time, at which point it implements a decision based on the information acquired so far. Our key assumption is that the firm's final decision depends (only) on a weighted sum of the attribute values, which we call the *payoff-relevant state*.

What complicates the firm's information acquisition problem is the possibility for the different attribute values to be correlated. For example, one data source may provide information about vacation travel demand from American travelers, while another provides information about vacation travel demand from Canadian travelers. Although these two demand levels factor separately into the firm's total demand forecast, their correlation means that information about one affects the value of information about the other. Thus, the firm has to choose the optimal proportion of resources to allocate to each data source at any given moment, taking into account the potential complementarity or substitution among the different data sources.

[1] Often firms purchase raw data which must be re-processed, analyzed, or cross-referenced to existing data before it is useful. We can interpret allocation of attention/resources to the data source as allocation of employee hours towards extracting information from the given data set.

Under a condition on the prior that we provide, the optimal dynamic data acquisition strategy turns out to take a simple form. It consists of K stages, where K is the number of data sources. In the first stage, the firm exclusively observes the single most informative data source. In each subsequent stage k, the firm acquires information from one additional data source, and allocates resources in a constant fraction across these k sources. Once the firm reaches the final stage, it acquires information in a constant fraction over all data sources from that point on. Both the nested support sets of data sources observed at each stage, as well as the mixtures over them, are history-independent, and can be determined directly from the informational primitives.

This characterization reveals several properties about optimal information acquisition in our setting. First, despite the complexity of the information environment (i.e. multiple sources, flexible correlation), the firm's optimal rule is simple: Once it starts acquiring information from a given data source, it continues to acquire information from that source. At fixed times, the firm adds in a new source, re-weighting its resources over the sources observed in the past as well as the new source. Since the times at which the firm brings in new sources, and also the fraction of resources across these sources, do not depend on the history of signal realizations, the firm can map out and implement a deterministic plan for information acquisition from time 0.

Second, the optimal information acquisition strategy does not depend on the firm's payoff function or its choice of when to stop acquiring information, provided that the payoff-relevant state does not change. Returning to our previous example, so long as the (forecast of) demand for the Puerto Rico location remains a sufficient statistic for the firm's decision, the optimal strategy described above is robust to changes in the decision as well as changes in the firm's discount factor. That is, if the hotel's objective changed from pricing the Puerto Rico location to a decision about whether to open the location at all, its past resource allocations would remain optimal for this new objective. Likewise if the firm's data acquisition budget were unexpectedly cancelled, its choices up to that point would coincide with those it would have made had it foreseen this outcome.

The condition that we assume on the prior belief roughly requires that the covariances of the different attributes are small in magnitude relative to their variances. For the case of two attributes, it is sufficient for the covariances to be smaller than the variances; in general "how much smaller" will depend on the number of data sources. Intuitively, such a condition puts an upper bound on the possible complementarity or substitution effects between different data sources. As a result, the firm's short-run and long-run information acquisition incentives are aligned, so that it is optimal to focus on the most informative sources at this moment.[2] Although our condition on the prior belief is restrictive, we show that under optimal sampling from *any* prior belief, the firm's posterior beliefs will eventually satisfy the condition, at which point our characterization applies.

We conclude by showing how our characterization of the optimal information acquisition strategy can be applied to make advances in other problems. First, we consider an extension of our environment in which the data providers are themselves strategic, and can control the precision of the information that they provide.[3] We suppose that a mass of forward-looking firms optimally acquire information from the providers over time. Using our characterization of the agents' information acquisition strategy, we derive the equilibrium choices of precision. These precision levels turn out to be monotonically increasing in the data providers' discount rate and in the degree to which the unknown attributes are correlated. That is, the *more patient* the providers are, and the *less correlated* the attributes are, the *lower the precision* of the signals.

Second, we turn to a different but related setting: endogenous information acquisition for binary choice. In the classic binary choice problem, a consumer can choose between two goods with unknown payoffs, and learn about either payoff before making his decision. A result in Fudenberg et al. (2018) considers endogenous allocation of attention across learning about either payoff, and characterizes the optimal attention allocation strategy under the assumption that payoffs are independent and identically distributed.

The binary choice problem corresponds to the special case of our setting with two unknown attributes, where the agent wants to learn a difference of the attributes. The optimal attention allocation strategy we demonstrate holds in this case for all prior beliefs, thus generalizing the result in Fudenberg et al. (2018) to allow for correlation across the two payoffs. (We would expect payoffs to be correlated, for example, if the values of the goods depended on a common source of uncertainty—e.g., different portfolio choices, or consumer goods with shared features.) Using our characterization, we derive a new comparative static result with respect to prior uncertainty. We show that an increase in initial uncertainty about either payoff results in a *uniform* change in attention (that is, either weakly more attention paid to learning about that payoff at every instant, or weakly less), but the direction depends on the size of correlation between the unknown payoffs. Specifically: an increase in the initial uncertainty about one payoff results in *higher* attention to the corresponding data source

[2] Formally, these are the sources that maximize the marginal reduction in the posterior variance of the payoff-relevant state.

[3] For example, Google may decide whether to provide search trend data aggregated at the hour-level or the week-level.

when the two payoffs are weakly correlated in the prior, but results in *lower* attention when the payoffs are strongly correlated.

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

1.1 Related Literature

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

We primarily build on a large literature about optimal dynamic information acquisition. In contrast to an earlier focus in the literature on the choice of signal precisions (Moscarini and Smith, 2001), our framework characterizes the choice between different *kinds* of information, as in the work of Fudenberg et al. (2018) (where the information sources are two Brownian motions), and Che and Mierendorff (2019) and Mayskaya (2019) (where the sources are two Poisson signals). Compared to this work, our main contribution is to accommodate many sources that may be flexibly correlated.[4]

Another strand of the literature considers an agent who chooses from completely flexible information structures at entropic (or more generally, "posterior-separable") costs, such as in Steiner et al. (2017), Hébert and Woodford (2018) and Zhong (2018). Compared to these papers, our agent has access to a prescribed (physical) set of signals, and acquires information under a resource/attention capacity constraint. Thus the different signals in our setting are equally costly to acquire regardless of the current belief, which is the key distinction from measuring information acquisition costs by the reduction of uncertainty.

In previous work Liang et al. (2019), we studied a related setting in discrete time, introduced the notion of "myopic information acquisition" and studied its (approximate) optimality properties. Going beyond those results, the characterizations in the current paper

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

precisely (and more generally) describe the optimal path of attention allocations, which are useful in applications, as we illustrate. The technical methods in this paper also differ from the prior work: Without a clear definition of myopic information acquisition in continuous time, we instead analyze the Hessian matrix associated with the posterior variance function to derive the optimal allocations of attention.

Finally, this paper is related to recent work on data acquisition by firms. Azevedo et al. (2019) studies allocation of resources (i.e. test users) to learn about the quality of multiple innovation projects. These authors show that the tail distribution of innovation quality crucially affect the (static) optimal experimentation strategy. Immorlica et al. (2018) considers dynamic allocations of a budget of data samples for learning about an evolving state, and demonstrates (near) efficiency guarantees for certain classes of benchmark policies. Bonatti and Cisternas (2019) analyze a dynamic game in which firms use a consumer's "score" to infer about her preferences and set prices. Different from these papers, we have a setting in which the firm has to dynamically aggregate multiple sources of information. Our characterizations trace out a time path of market demand for various kinds of information, which is absent from the literature.

2 Model

An agent (i.e. firm) has uncertainty about the values of K attributes $\theta = (\theta_1, \ldots, \theta_K)'$,[5] and his prior is that they are jointly normal with known mean vector μ and covariance matrix Σ, where Σ has full rank. The agent wants to learn an unknown payoff-relevant state $\omega = \langle \alpha, \theta \rangle$, which is a linear combination of these attribute values. The weight vector $\alpha \in \mathbb{R}^K$ is known and fixed, and we assume without loss of generality that each coordinate α_i is strictly positive, so that the state depends positively on all of the attribute values.[6]

Time is continuous. There is a data source that provides information about each attribute, and the agent divides his attention (i.e. resources) across these sources at every instant. Formally, we assume that the agent has one unit of attention in total at every point in time, and chooses attention levels $\beta_1^t, \ldots, \beta_K^t$ subject to $\beta_i^t \geq 0$ and $\sum_i \beta_i^t \leq 1$.[7]

These choices influence the diffusion processes $X_1, \ldots, X_K$ (observed by the agent) in the

[5] Here and later, we use the apostrophe to denote vector or matrix transpose.

[6] It is without loss to assume that the weights are non-negative because any attribute value can be replaced with its negative. Those attribute values with zero weights can be dropped without affecting the results.

[7] See e.g. Fudenberg et al. (2018) and Che and Mierendorff (2019) for recent models with fixed budgets of attention.

following way:

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

Above, each B_i is an independent standard Brownian motion, and the term $\sqrt{\beta_i^t}$ is a normalizing factor to ensure constant informativeness per unit of attention devoted to each source.

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

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

At the endogenously chosen end time τ, the agent will choose from a set of actions A and receive the payoff $u(a, \omega)$, where u is a known payoff function that depends on the action taken a and the payoff-relevant state ω. The agent's posterior belief about ω at this time determines the action that maximizes his expected flow payoff $\mathbb{E}[u(a, \omega)]$. To ensure that optimal information acquisition is unique, we make a mild assumption about u such that the agent always *strictly* benefits from having more precise information.[8] Formally, we impose throughout:

Assumption. For any variance $\sigma^2 > 0$ and any action $a^* \in A$, there exists a positive measure of μ for which a^* does not maximize $\mathbb{E}[u(a, \theta) \mid \theta \sim \mathcal{N}(\mu, \sigma^2)]$.[9]

In words, the agent's expected value of θ affects the optimal action to take (holding fixed

[8] To be completely rigorous, the optimal information acquisition strategy is *uniquely determined up to the stopping time*, after which the attention choices do not matter. We will ignore this when stating our results.

[9] A sufficient condition is that the two limiting states $\theta \to +\infty$ and $\theta \to -\infty$ disagree about the optimal action. This is true for many applications of the model.

his belief variance). This rules out the existence of dominant actions, which would make information acquisition irrelevant. Besides this, the functional form of $u(\cdot)$ can be arbitrary.

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

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

where $c(\tau)$ is a non-negative and weakly increasing function that measures the cost of waiting until time τ.[10] Our focus throughout this paper is on the optimal information acquisition strategy S. In general, the strategies S and τ should be determined *jointly*, but our results will show that in fact these problems can be separated, with the optimal S characterized independently from the choice of τ.

3 Optimal Information Acquisition Strategy

3.1 $K = 2$

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

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

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

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

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

Since both variances Σ_{11}, Σ_{22} are positive, Assumption 1 can be understood as requiring that the covariance Σ_{12} is not too negative relative to the size of either variance. A sufficient condition is for the weights on the two attributes to be equal, i.e. $\alpha_1 = \alpha_2$, in which case Assumption 1 holds for all priors.[11] A different sufficient condition is for the attributes to be positively correlated ($\Sigma_{12} = \Sigma_{21} \geq 0$), in which case Assumption 1 holds for all weights α_1 and α_2.

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

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

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

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

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

where $x_1 = \alpha_1 \det(\Sigma)$, $y_1 = \alpha_1\Sigma_{11} + \alpha_2\Sigma_{12}$, $x_2 = \alpha_2 \det(\Sigma)$, *and* $y_2 = \alpha_1\Sigma_{21} + \alpha_2\Sigma_{22}$. *W.l.o.g. let* $y_i \geq y_j$. *Then the optimal information acquisition strategy* (β_1^t, β_2^t) *consists of two stages:*

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

Thus there are two stages of information acquisition. In the first stage, which ends at some t^* (that depends only on the prior covariance matrix Σ and the weight vector α), the agent allocates all of his attention to one of the attributes. After time t^*, he divides his attention across the attributes in a fraction that is constant across time. The long-run attention level is proportional to the weights α, which means that dividing attention according to this fraction achieves the most efficient aggregation of information about ω.

Observe additionally that the optimal information acquisition strategy does not depend on the agent's cost of waiting or the details of his payoff function, so long as α is unchanged and $\omega = \alpha' \cdot \theta$ remains the payoff-relevant state. Thus, when the prior belief satisfies Assumption 1, the optimal information acquisition strategy is constant across different objectives and also across different stopping rules. Relatedly, we can solve for the optimal stopping rule in this setting as if information acquisition were *exogenously* given by Theorem 1. We mention that Assumption 1 is also *necessary* for the optimal strategy to be constant across all cost functions $c(\cdot)$; see Lemma 8 and Lemma 9 for details.

3.2 General K

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

A key condition on the prior belief, parallel to the one used in Assumption 1, is the following:[12]

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

[12] In the appendix, we prove the result under the following weaker—but somewhat less interpretable—condition: The prior precision matrix Σ^{-1} satisfies $[\Sigma^{-1}]_{ii} \geq \sum_{j \neq i} |[\Sigma^{-1}]_{ij}| \quad \forall 1 \leq i \leq K$.

This condition requires that the covariance between every pair of attribute values is small relative to the variances. For the case of two attributes, this condition requires only that the covariance Σ_{12} is smaller in magnitude than both variances Σ_{11} and Σ_{22}, which would imply our previous Assumption 1. In general, the condition in Assumption 2 is harder to satisfy when the number of sources K is larger.

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

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

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

and nested sets

$$\emptyset = B_0 \subsetneq B_1 \subset \cdots B_{K-1} \subsetneq B_K = \{1, \ldots, K\},$$

such that for each $1 \leq k \leq K$, the optimal attention level 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 level at any time $t \geq t_{K-1}$ is proportional to α.

The times t_k as well as the attention level (including its support B_k) at each stage can be determined directly from the primitives α and Σ, and are independent of the signal realizations. Theorem 2 thus tells us that the agent can reduce the dynamic information acquisition problem to a sequence of K static problems, each of which involves finding the optimal constant attention for a fixed period of time (from t_{k-1} to t_k). Furthermore, as in the $K = 2$ case, the optimal information acquisition strategy does not depend on the agent's payoff function or stopping rule.

To interpret the use of Assumption 2, note that prior covariances measure the *complementarity or substitution effects* across the information provided by different data sources (i.e. whether information from one data source increases or decreases the learning benefits from other sources). Assumption 2 limits the magnitude of such complementarity/substitution, so that the agent's short-run and long-run information acquisition incentives are aligned.[13] It is this property that underlies the characterization in Theorem 2.

We previously discussed the sense in which Assumption 1 is tight when there are just two data sources. With $K > 2$, Assumption 2 is sufficient but not in general necessary for the characterization in Theorem 2 to hold. In particular, optimal information acquisition may also consist of K stages as described in the theorem under alternative assumptions on Σ *and*

[13] Otherwise, the agent may (for example) prefer one data source for the short run but choose to learn from another pair of complementary data sources for the long run.

α; see Section 4 for an example.[14] However, the constant $\frac{1}{2K-3}$ in Assumption 2 cannot be improved upon in the statement of the theorem, as we formalize in Appendix C.6.

Moreover, it turns out that the agent's posterior beliefs under optimal sampling from *any* prior belief will eventually satisfy Assumption 2. In fact, optimal sampling is not required: Along any path in which each data source is infinitely sampled (which is necessary for complete learning of ω and thus satisfied under optimal sampling), the agent's beliefs will enter and stay within the set of beliefs defined by Assumption 2.

Formally, define *cumulated attention for source i* $q_i(t) = \int_0^t \beta_i^s \, ds$ under an arbitrary sampling strategy. We then have:

Lemma 1. *Starting from any prior belief, the optimal information acquisition strategy has the property that the induced cumulated attentions* $q_i(t) \to \infty$ *for each* $1 \leq i \leq K$ *as* $t \to \infty$.[15]

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

Combining Lemma 1, Lemma 2 and Theorem 2, we have the following result:

Proposition 1. *Starting from any prior belief, the optimal information acquisition strategy is eventually a constant attention level (across all data sources) proportional to the weight vector* α.[16]

3.3 Proof Sketch for Theorem 2

The plan of the proof is to first define a strategy which results in pointwise minimum variance along its path—we call this strategy *uniformly optimal*—and next to show that this strategy is the optimal information acquisition strategy. Then we characterize its structure.

Definition of a uniformly optimal strategy. As mentioned above, at every time t the agent's past attention levels integrate to a *cumulated attention vector*

[14] Generalizing Assumption 2 to a necessary and sufficient condition would be an interesting direction for future work.

[15] We note that starting from a general prior belief, $q_i(t)$ can be a random variable depending on past signal realizations. Thus the lemma asserts that each source is infinitely sampled regardless of signal realizations.

[16] More specifically, we show in the proof that there exists $\bar{t}$ depending only on α and Σ, such that the optimal attention level at any time $t \geq \bar{t}$ is proportional to α. This holds independently of the payoff function or past signal realizations.

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

describing how much attention has been paid to each source. These cumulated attention vectors $q(t)$ determine the agent's posterior variance about ω at time t via:

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

where $\mathrm{diag}(q)$ is the diagonal matrix with entries $q_1(t), \ldots, q_K(t)$. Recall that Σ is the agent's prior covariance matrix about the vector of attribute values, and α is the weight vector, both primitives of the model.

Define the *t-optimal* vector to be

$$n(t) = \operatorname*{argmin}_{q_1,\ldots,q_K \geq 0,\ \sum_i q_i = t} V(q_1, \ldots, q_K)$$

namely the allocation that minimizes posterior variance (among all attention vectors that allocate a budget of t).[17] We will say that a strategy is uniformly optimal if it achieves a t-optimal attention vector at every instant t.

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

That is, the strategy S deterministically leads to minimum posterior variance about ω at every time t. This is a strong property, and existence of such a strategy is in general not guaranteed.

When a uniformly optimal strategy exists, it is optimal. By definition, if a cumulated attention vector is t-optimal, it implies that the agent has learned as much about ω as possible in the interval $[0, t)$. Thus, if the agent stops acquiring information at instant t (and takes the optimal action), then his expected flow payoff is maximized among all strategies that stop at t.[18] Requiring that $q(t)$ is t-optimal at *every* time t then implies that the information acquisition strategy is most informative about ω at every history and maximizes expected payoffs given any *exogenous* stopping time. In our Gaussian environment, such a strategy also maximizes expected payoffs even when the stopping time can be endogenously chosen; this follows from a result of Greenshtein (1996). Given this discussion, whenever a uniformly optimal strategy exists, it must be the optimal strategy in our problem.

[17] We show in Lemma 5 that this minimizer is unique.

[18] Due to normal beliefs, achieving minimum posterior variance means that the agent's information up to time t is Blackwell more informative than under any other strategy (Hansen and Torgersen, 1974). Thus the form of the payoff function u does not matter.

Existence and structure of the uniformly optimal strategy. It remains to show that under Assumption 2, a uniformly optimal strategy exists, and has the structure described in Theorem 2.

Suppose without loss of generality that the t-optimal vector $n(t)$ at some time t is supported on the first k sources. The posterior variance function $V(q)$ is continuously differentiable, so the vector $n(t)$ is determined by a straightforward first-order condition. Specifically, the first k sources should maximize $|\partial_i V|$, the marginal reduction in the posterior variance. What is less obvious is that the same first-order condition continues to hold as the agent increases $n(t)$ in a direction that represents a mixture over the first k data sources, and moreover that this direction remains optimal for a period of time. The analysis here involves comparisons of *cross-partial derivatives of* V. Indeed, from a technical perspective, the marginal change of $n(t)$ as t increases can be found by differentiating the first-order condition, and is thus (inversely) related to the Hessian matrix of V.

Intuitively, the second derivatives of V capture how the latest information acquisition affects the marginal values of each data source. The condition in Assumption 2 ensures that as the agent optimally divides attention among the first k data sources, the specific mixture of new information reduces the marginal value of each of the k sources. This turns out to be sufficient for information acquisition in this mixture to remain optimal (i.e. maintain the first-order condition). Thus as the cumulated attention vector changes in this direction, the resulting posterior beliefs continue to minimize the posterior variance (for a while).

As the agent acquires information from the first k data sources, his beliefs about the first k attributes become more precise. The marginal values of learning about the remaining attributes (relative to learning about the first k) thus increase continuously during this time. Eventually, some new data source will have the same marginal value as the first k sources. At this point the agent expands his observation set to include that new source, and we can repeat the reasoning above. Demonstrating that the previous k data sources continue to receive positive attention (i.e. nested supported sets) is a key technical challenge, and we refer the reader to the appendix for details.

4 Application: Competing Information Providers

We now consider an application of our main results to a setting in which the data sources/providers are themselves strategic, and can control the precision of the information that they provide. For example, if the data source corresponds to search data or discretionary spending data, the data provider may control the resolution at which this information is provided (i.e. whether the data is aggregated at the week-level or month-level). Suppose a mass of forward-looking firms optimally acquire information from the sources over time, and each data provider seeks to maximize engagement with their data. What are the equilibrium choices of precision by the data providers?

In more detail, we suppose that firms seek to predict the sum of attributes $\theta_1 + \theta_2$, and their common prior over the parameters is

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

where $\rho \in (-1, 1)$ measures prior correlation between θ_1 and θ_2. Each of two data providers $i = 1, 2$ (freely) chooses a noise variance σ_i^2, so that the information it provides is

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

Note that there is no cost to choosing higher precision values.

Firms optimally allocate attention given these precision choices (which are fixed across time). Since we've assumed a common prior, all firms make the same information acquisition decisions, and it is without loss to consider a single firm whose allocation at time t is denoted (β_1^t, β_2^t). To map this setting into our main model, we normalize the noise terms to have *unit variances* as follows: Define $\tilde{\theta}_i = \frac{\theta_i}{\sigma_i}$, so that each data provider's signal is equivalent to $\tilde{\theta}_i$ plus standard Gaussian noise. Under this transformation, the firm seeks to learn $\sigma_1\tilde{\theta}_1 + \sigma_2\tilde{\theta}_2$, where its prior covariance matrix over $(\tilde{\theta}_1, \tilde{\theta}_2)$ is

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

Note that Assumption 1 is satisfied, since

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

Thus the optimal attention choices (β_1^t, β_2^t) are characterized by Theorem 1.

Each data provider i's payoff is the discounted average attention paid to that source $\int e^{-rt}\beta_i^t dt$, where r is a (common) discount rate. We can interpret this as reduced form for a profit margin, where each data provider receives a payoff proportional to the time the firm spends gathering its information.[19]

[19] Here, for the sake of illustrating the equilibrium, we are considering the case where firms sample forever. This corresponds to the limit as the waiting cost function $c(\cdot)$ decreases point-wise to zero.

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

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

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

We note that the less patient the information providers are, the more precise the signals are in equilibrium. Intuitively, when data provider i increases the precision of the information it provides, there are two opposing effects: On the one hand, weakly more attention is attracted to i early on, since i's information becomes more valuable relative to source j, and it is more likely to be the source chosen in phase 1. On the other hand, increasing precision lowers the long-run frequency $\frac{\sigma_i}{\sigma_i+\sigma_j}$ with which i is viewed, since firms need fewer observations of i's information to achieve the same level of precision about θ_i. Thus, less patient data providers compete over short-run profits (i.e. being chosen in phase 1) and provide precise signals, while patient data providers compete for long-run profits (i.e. long-run frequency) and provide imprecise signals. As far as we are aware, the effect of information precision on the *time path* of people's information demand has not been noted in the previous literature.

Additionally, the more positively correlated the unknown parameters are (higher covariance ρ), the higher the precision of the signals provided in equilibrium. This is because (as we derive in the proof of the proposition) the threshold $t_i^* = \frac{(\sigma_j - \sigma_i)\sigma_i}{1-\rho}$ increases in ρ, which increases the value to being the information source chosen in phase 1. The competition for short-run profits thus drives the signals to be more precise. We provide a rough intuition for why t_i^* increases monotonically in ρ. Note that when the unknowns θ_1, θ_2 are negatively correlated ($\rho < 0$), the two signals about these unknowns are *complements* for estimating their sum: This can be formalized by computing the cross-partial derivative $\frac{\partial^2 V(q_1,q_2)}{\partial q_1 \partial q_2}$ and checking that it has the same sign as ρ.[20] Complementarity between the signals implies a stronger incentive to observe them together, and so phase 1 is shorter when ρ is more negative.

From a social welfare perspective, these comparative statics tell us that more information is released into society (and hence society learns faster) when information providers are *less forward-looking*, and when the information they provide is *more similar*. In the appendix, we generalize these insights to a game where $K > 2$ information providers compete, and where society seeks to predict $\theta_1 + \cdots + \theta_K$. Observe that the transformed prior covariance matrix $\tilde{\Sigma}$

[20] Since the firm seeks to *minimize* V, having a negative cross-partial means the marginal reduction of uncertainty by the signal about θ_2 becomes *bigger* after having observed the signal about θ_1.

does not in general satisfy Assumption 2; this can occur if ρ is large or if the providers choose very different signal precision levels. Nonetheless, we directly compute the uniformly optimal strategy (as defined in the proof sketch) and show that the characterization in Theorem 1 extends to this setting. We find $\sigma^* = \sqrt{\frac{1-\rho}{Kr}}$ to be the equilibrium precision in the unique symmetric pure strategy equilibrium. Thus the findings in this section are robust to the presence of many data sources.

5 Application: Binary Choice

The framework we study relates to a large body of work regarding "binary choice tasks," in which an agent has a choice between two goods with payoffs θ_1 and $-\theta_2$ (we introduce the negative here for expositional simplicity), and can devote effort towards learning about these payoffs before making his decision. The leading model in this domain, the drift-diffusion model (Ratcliff and McKoon, 2008), supposes that the agent observes a Brownian motion whose drift depends on which good yields the higher payoff. In our framework, this model corresponds to a case in which the agent's prior belief is supported on two points—either $(\theta_1, -\theta_2) = (\theta', \theta'')$ or $(\theta_1, -\theta_2) = (\theta'', \theta')$ where $\theta' > \theta''$ are known quantities.[21] Thus the agent has uncertainty over which good is better, but not over how much better it is. Fudenberg et al. (2018) recently generalized this model to allow additionally for the latter kind of uncertainty. In their *uncertain drift-diffusion* model, the agent has a jointly normal prior over $(\theta_1, -\theta_2)$, and has access to two Brownian motions with drifts corresponding to these unknown payoffs.

Both the classic DDM model and also Fudenberg et al. (2018) focus on the optimal stopping rule given *exogenous* information. But Fudenberg et al. (2018) additionally proposes a version of their model in which the agent endogenously acquires information by choosing attention levels (subject to an budget constraint) that scale the drifts of the Brownian motions. Indeed, this corresponds exactly to our framework with $K = 2$ and equal weights (since the payoff difference $\theta_1 + \theta_2$ is a sufficient statistic for the agent's decision). They impose further that the agent's prior is *independent*—that is, $\Sigma = I$—and find that the agent optimally devotes equal attention to both information sources at all times.

Applying Theorem 1 with $\alpha_1 = \alpha_2 = 1$, we obtain the following immediate generalization

[21] Although in our setting the agent observes two Brownian motions with drifts θ_1 and θ_2, his decision is only a function of their sum, which is a single Brownian motion with drift $\theta_1 + \theta_2$. The classic DDM thus assumes that the magnitude of $\theta_1 + \theta_2$ is known to the agent.

of this result on optimal information acquisition.[22]

Corollary 1. *Suppose $K = 2$, $\alpha_1 = \alpha_2 = 1$ and $\Sigma_{ii} \geq \Sigma_{jj}$. The agent's optimal information acquisition strategy (β_1^t, β_2^t) consists of two stages:*

- ***Phase 1:*** *At all times*
$$t \leq t_i^* = \frac{\Sigma_{ii} - \Sigma_{jj}}{\det(\Sigma)},$$
the agent optimally allocates all attention to source i.

- ***Phase 2:*** *At times $t > t_i^*$, the agent allocates attention in the constant fraction $(\frac{1}{2}, \frac{1}{2})$.*

This result improves upon Theorem 5 in Fudenberg et al. (2018) by allowing for possible correlation as well as asymmetry between the two unknown payoffs. Fudenberg et al. (2018) point out that their result does not characterize "off-equilibrium" attention allocation, since it no longer applies if the agent has paid unequal attention to the two sources in the past. In contrast, our result applies to all prior beliefs and thus allows for characterization of the optimal information acquisition strategy at any history, including those in which the agent has previously behaved sub-optimally.

When $\Sigma = I$, the thresholds are $t_1^* = t_2^* = 0$, so that the agent splits his attention evenly from the beginning. This returns the solution in Fudenberg et al. (2018). For general prior covariance matrices Σ, the agent also eventually acquires information according to the stationary fraction of $(\frac{1}{2}, \frac{1}{2})$. However, the agent begins by learning about the good over which he has greater initial uncertainty, until some time t^* at which his posterior variances about the two unknown payoffs equalize, and his attention choice jumps discontinuously to $(\frac{1}{2}, \frac{1}{2})$.

From Corollary 1 we see additionally that the prior belief impacts the agent's attention strategy only by determining which source is observed in phase 1, and for how long that phase lasts. Thus, changes in the prior result in the agent paying *uniformly* more or less attention to either source. In particular if we consider the impact of changes in initial uncertainty about the attribute values, we have the following comparative static:

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

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

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

Appendix

A Preliminaries

We first introduce a definition that captures the objective of minimizing the posterior variance about ω at some given time t.

Definition 2. *Say a cumulated attention vector $n(t)$ is t-*optimal *if*

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

In this way, we can rephrase the requirement of uniform optimality as follows: A strategy is uniformly optimal if and only if its cumulated attention vector $q(t)$ is t-optimal for each t.

Lemma 3. *The posterior variance about ω can be written in two ways:*

$$V(q_1, \ldots, q_K) = \alpha' \left[(\Sigma^{-1} + \operatorname{diag}(q))^{-1}\right] \alpha = \alpha' \left[\Sigma - \Sigma(\Sigma + \operatorname{diag}(1/q))^{-1}\Sigma\right] \alpha$$

where $\operatorname{diag}(q)$ denotes the diagonal matrix with entries $q_1, \ldots, q_K$ and $\operatorname{diag}(1/q)$ has entries $1/q_1, \ldots, 1/q_k$. This function V extends to a rational function (quotient of polynomials) over all of $\mathbb{R}^K$ (even if some q_i are negative).

Lemma 4. *Given a cumulated attention vector $q \geq \mathbf{0}$, define*

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

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

$$\partial_i V = -\gamma_i^2 \qquad \partial_{ij} V = 2\gamma_i\gamma_j \cdot \left[(\Sigma^{-1} + \operatorname{diag}(q))^{-1}\right]_{ij}.$$

As a corollary, the Hessian matrix is $2 \cdot \operatorname{diag}(\gamma)(\Sigma^{-1} + \operatorname{diag}(q))^{-1} \cdot \operatorname{diag}(\gamma)$ is positive semi-definite, and so V is decreasing and convex in $q_1, \ldots, q_K$ whenever $q_i \geq 0$.

Proof. From Lemma 3 and the formula for matrix derivatives, we have

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

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

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

as we desire to show. □

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

Lemma 5. *For each $t \geq 0$, there is a unique t-optimal vector $n(t)$.*

Proof. Suppose for contradiction that two vectors $(r_1, \ldots, r_K)$ and $(s_1, \ldots, s_K)$ 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 $\lambda \to +\infty$. 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 $\lambda \to +\infty$, $\operatorname{diag}(1/|q^\lambda|)$ has the same limit as $\operatorname{diag}(1/q^\lambda)$. Thus by the second expression for V (see Lemma 3), $\lim_{\lambda\to\infty} V(|q^\lambda|) = \lim_{\lambda\to\infty} V(q^\lambda) = V(r)$. 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 $V(q^*) = V(r)$ implies $\partial_i V(r) = 0$ for $1 \leq i \leq l$.

Consider the vector $\gamma = (\Sigma^{-1} + \operatorname{diag}(r))^{-1}\alpha$. By Lemma 4, $\partial_i V(r) = -\gamma_i^2$ for $1 \leq i \leq K$. Thus $\gamma_1 = \cdots = \gamma_l = 0$. Since γ is not the zero vector,[23] 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 as $r_1 > s_1$) and increased the j-th coordinate by the same amount. This contradicts the assumption that r is t-optimal. Hence the lemma holds. $\square$

Given Lemma 5, a uniformly optimal strategy must have cumulated attention vector $n(t)$ at each time t. Thus a necessary condition for such a strategy to exist is that $n(t)$ weakly increases (in each coordinate) over time. Conversely, when $n(t)$ is monotonic, we can define instantaneous attention levels β^t to be the time-derivative of $n(t)$. Then this strategy indeed achieves uniform optimality. To summarize, we have

Lemma 6. *A uniformly optimal strategy exists if and only if the t-optimal attention vector $n(t)$ weakly increases over time.*

[23] This follows because α is not the zero vector, by assumption.

The following result ensures that a strategy that minimizes the posterior variance uniformly at all times is a dynamically optimal strategy in any decision problem. It is a continuous-time version of Greenshtein (1996).

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

We also have a simple converse result:

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

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

B Proof of Theorem 1

By Lemma 4 and direct computation, we have

$$\begin{aligned} \partial_1 V(q_1, q_2) &= \frac{-(x_1 q_2 + y_1)^2}{\det^2(\Sigma Q + I)} \\ \partial_2 V(q_1, q_2) &= \frac{-(x_2 q_1 + y_2)^2}{\det^2(\Sigma Q + I)} \end{aligned} \tag{1}$$

where x_1, x_2, y_1, y_2 are as defined in Theorem 1:

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

Q is the diagonal matrix with entries q_1, q_2 and I is the identity matrix.

Note that Assumption 1 translates into $y_1 + y_2 \geq 0$. Under this assumption, we will characterize the t-optimal attention vector $(n_1(t), n_2(t))$ and show it is increasing. Without loss assume $y_1 \geq y_2$, then y_1 is non-negative. Let $t_1^* = \frac{y_1 - y_2}{x_2}$. Then when $q_1 + q_2 \leq t_1^*$ we always have $x_1 q_2 + y_1 \geq x_2 q_1 + y_2$ as well as $x_1 q_2 + y_1 \geq -(x_2 q_1 + y_2)$. Thus (1) implies that $\partial_1 V(q_1, q_2) \leq \partial_2 V(q_1, q_2)$ at such attention vectors q. So for any budget of attention $t \leq t_1^*$, putting all attention to source 1 minimizes the posterior variance function V. That is, $n(t) = (t, 0)$ for $t \leq t_1^*$.

For $t > t_1^*$, observe that $\partial_1 V(0, t) < \partial_2 V(0, t)$ as well as $\partial_1 V(t, 0) > \partial_2 V(t, 0)$. These imply that $n(t)$ is interior, and the first-order condition $\partial_1 V = \partial_2 V$ yields the solution

$n(t) = (\frac{x_1 t + y_1 - y_2}{x_1 + x_2}, \frac{x_2 t - y_1 + y_2}{x_1 + x_2})$. Thus $n(t)$ is indeed increasing in t, and the instantaneous attention levels are as described by Theorem 1. This completes the proof.

We mention that the assumption $y_1 + y_2 \geq 0$ is also necessary for a uniformly optimal strategy to exist.

Lemma 9. *Suppose the prior covariance matrix fails Assumption 1. Then a uniformly optimal strategy does not exist.*

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

1. If $t \leq \frac{-(y_1+y_2)}{x_1}$, then $x_2 q_1 + y_2$ is negative and has larger magnitude than $x_1 q_2 + y_1$ whenever $q_1 + q_2 = t$. By (1), this means $\partial_1 V \geq \partial_2 V$ and so $n(t) = (0, t)$ devotes all attention to source 2.

2. If $\frac{-(y_1+y_2)}{x_1} < t < \frac{-(y_1+y_2)}{x_2}$, then $\partial_1 V(0, t) < \partial_2(0, t)$ and $\partial_1 V(t, 0) > \partial_2(t, 0)$. These imply that $n(t)$ is interior, and the first-order condition yields $x_1 n_2(t) + y_1 = -(x_2 n_1(t) + y_2)$ (for t in this range, $x_2 q_1 + y_2$ is always negative). Together with $n_1(t) + n_2(t) = t$, we can solve that $n(t) = (\frac{x_1 t + y_1 + y_2}{x_1 - x_2}, \frac{-x_2 t - y_1 - y_2}{x_1 - x_2})$.

3. If $\frac{-(y_1+y_2)}{x_2} \leq t \leq \frac{y_1 - y_2}{x_2}$, then $(x_1 q_2 + y_1)^2 - (x_2 q_1 + y_2)^2 = (y_1 - y_2 - x_2 q_1 + x_1 q_2) \cdot (y_1 + y_2 + x_1 q_2 + x_2 q_1) \geq 0$ whenever $q_1 + q_2 = t$. Thus $\partial_1 V(q_1, q_2) \leq \partial_2 V(q_1, q_2)$, implying that the t-optimal attention vector should be $n(t) = (t, 0)$.

4. Finally, if $t > \frac{y_1 - y_2}{x_2}$, then as in the second case $\partial_1 V(0, t) < \partial_2(0, t)$ and $\partial_1 V(t, 0) > \partial_2(t, 0)$. So $n(t)$ is interior and satisfies the first-order condition $x_1 n_2(t) + y_1 = x_2 n_1(t) + y_2$. This yields the solution $n(t) = (\frac{x_1 t + y_1 - y_2}{x_1 + x_2}, \frac{x_2 t - y_1 + y_2}{x_1 + x_2})$ and completes the analysis.

Note that in Case 2 above, as t increases in the range, $n_2(t)$ *actually decreases*. This proves that a uniformly optimal strategy does not exist. □

C Proof of Theorem 2

C.1 Weaker Assumption

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

Assumption 3. *The inverse of the prior covariance matrix Σ^{-1} is diagonally-dominant. That is,*

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

This is implied by Assumption 2 via the following lemma.

Lemma 10. *Suppose the prior covariance matrix Σ satisfies Assumption 2, then its inverse matrix satisfies $[\Sigma^{-1}]_{ii} \geq (K-1) \cdot |[\Sigma^{-1}]_{ij}| \quad \forall i \neq j$ and is thus diagonally-dominant.*

Proof. We focus on the first row of Σ^{-1}, as the other rows can be symmetrically handled. Let $x_i = [\Sigma^{-1}]_{1i}$ for $1 \leq i \leq K$, and without loss assume x_2 has the greatest absolute value among $x_2, \ldots, x_K$. It suffices to show $x_1 \geq (K-1)|x_2|$.

From $\Sigma^{-1} \cdot \Sigma = I$ we have $\sum_{i=1}^{K} [\Sigma^{-1}]_{1i} \cdot \Sigma_{i2} = 0$. Thus by symmetry, $\sum_{i=1}^{K} x_i \cdot \Sigma_{2i} = 0$. Rearranging yields

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

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

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

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

C.2 Technical Property of γ

The following technical lemma will be repeatedly used.

Lemma 11. *Suppose Σ^{-1} is diagonally-dominant. Given an arbitrary attention vector q, define γ as in Lemma 4 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 $\mathrm{diag}(q)$. Since $(\Sigma^{-1} + Q)^{-1}\alpha = \gamma$, we equivalently have

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

Suppose for contradiction that $\gamma_i \le 0$ 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

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

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

C.3 The Last Stage

To prove Theorem 2 under Assumption 3, we first consider those times t when each of the K sources has been sampled. The following lemma shows that it is optimal to maintain a constant attention level proportional to α ever after.

Lemma 12. *Suppose Σ^{-1} is diagonally-dominant. If at some time $\underline{t}$, the $\underline{t}$-optimal vector satisfies $\partial_1 V(n(\underline{t})) = \cdots = \partial_K V(n(\underline{t}))$, then the t-optimal vector at each time $t \ge \underline{t}$ is given by*

$$n(t) = n(\underline{t}) + \frac{t}{\alpha_1 + \cdots + \alpha_K} \cdot \alpha.$$ [24]

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

We check the equalities $\partial_1 V = \cdots = \partial_K V$ 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 $\mathrm{diag}(q)$ by Q to save notation, this marginal change equals

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

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

by Lemma 4. If we can show that $\gamma_1 = \cdots = \gamma_K$, then the above simplifies to

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

where the second equality follows from $(\Sigma^{-1} + Q)^{-1} \cdot \alpha = \gamma$. Thus, assuming that $\gamma_1 = \cdots = \gamma_K$, the marginal change of $\partial_i V$ is the same for all i. Hence $\partial_1 V = \cdots = \partial_K V$ continues to hold.

Because $\partial_1 V = \cdots = \partial_K V$ holds at the attention vector q, Lemma 4 implies that all γ_i have the same absolute value. They are thus equal by Lemma 11. □

C.4 Earlier Stages

In general, we need to show that even when the agent is choosing from a subset of the sources, the t-optimal vector $n(t)$ is still increasing over time. This is guaranteed by the following lemma, which says that the agent optimally attends to those sources that maximize the marginal reduction of V, until a new source becomes another maximizer. For ease of exposition we state the lemma under a slightly stronger assumption that Σ^{-1} is *strictly* diagonally-dominant. Later we will discuss how the lemma should be modified without this strictness.

Lemma 13. *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 point $\operatorname{argmin}_i \partial_i V(n(\bar{t}))$ is a strict superset of B. When $|B| = K$, it holds that $\bar{t} = \infty$ and β is proportional to α, as given by Lemma 12.

Proof. Without loss we assume $B = \{1, \ldots, k\}$ with $1 \leq k < K$. First we let $q = n(\underline{t})$ and define γ as before. By Lemma 11, γ_i is the same positive number for $i \leq k$. Next, the first-order condition for t-optimality implies that $q_j = 0$ whenever $j > k$. Otherwise V could be reduced by decreasing q_j and increasing q_1.

We now use a trick to deduce the current lemma from the previous Lemma 12. Specifically, given the prior covariance matrix Σ, we can choose another basis of the attributes $\theta_1, \ldots, \theta_k, \hat{\theta}_{k+1}, \ldots, \hat{\theta}_K$ with two properties:

1. each $\hat{\theta}_j$ $(j > k)$ is a linear combination of the original attributes $\theta_1, \theta_2, \ldots, \theta_K$;

2. $\mathrm{Cov}[\theta_i, \hat{\theta}_j] = 0$ for all $i \le k < j$, where the covariance is computed according to the prior belief.

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

$$\tilde{\alpha} = (\Sigma_{TL})^{-1} \cdot (\Sigma_{TL},\ \Sigma_{TR}) \cdot \alpha \tag{2}$$

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

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

Thus, in this smaller problem, the prior covariance matrix is Σ_{TL} and the payoff weights are $\tilde{\alpha}$. Assuming that $\tilde{\alpha}$ has positive coordinates, we can then apply Lemma 12: As long as the agent attends to the first k sources proportional to $\tilde{\alpha}$, $\partial_1 V = \cdots = \partial_k V$ continues to hold.[25] Moreover, at $q = n(\underline{t})$, the definition of the set B implies that these k partial derivatives are smaller (more negative) than the rest. By continuity, the same comparison holds until some time $\bar{t} > \underline{t}$. Thus, when $t \in [\underline{t}, \bar{t}]$, the cumulated attention vector (under this strategy) still satisfies the first-order condition $B = \mathrm{argmin}_{1 \le i \le K}\, \partial_i V$ and $q_j = 0$ for $j \notin A$. Since V is convex, this must be the t-optimal vector as desired.

It remains to prove that $\beta_i = \tilde{\alpha}_i$ is positive for $1 \le i \le k$. To this end, define $\tilde{Q} = \mathrm{diag}(q_1, \ldots, q_k)$ to be the $k \times k$ top-left submatrix of Q, and

$$\tilde{\gamma} = ((\Sigma_{TL})^{-1} + \tilde{Q})^{-1} \tilde{\alpha}. \tag{3}$$

We will show that $\tilde{\gamma}$ is just the first k coordinates of γ. Indeed, observe that $(\Sigma_{TL}^{-1} + \tilde{Q})^{-1}$ is

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

also the $k \times k$ top-left submatrix of $(\Sigma^{-1} + Q)^{-1}$.[26] Using (2) and (3), we have

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

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

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

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

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

which indeed holds.[27]

Hence $\tilde{\gamma}_i = \gamma_i$ for $1 \le i \le k$, and it is the same positive number by Lemma 11. Finally, we rewrite (3) as $\tilde{\alpha} = ((\Sigma_{TL})^{-1} + \tilde{Q})\tilde{\gamma}$. Thus $\tilde{\alpha}_i$ is (proportional to) the i-th row sum of the matrix $(\Sigma_{TL})^{-1} + \tilde{Q}$, which is just the row sum of $(\Sigma_{TL})^{-1}$ plus q_i. A theorem of Carlson and Markham (1979) says that if Σ^{-1} is (strictly) diagonally-dominant, then so is $(\Sigma_{TL})^{-1}$ for any principal submatrix Σ_{TL}. Consequently the row sums of $(\Sigma_{TL})^{-1}$ are all positive, implying that $\tilde{\alpha}_i > 0$. □

C.5 Completing the Proof

We now apply Lemma 13 repeatedly to prove Theorem 2. Continuing to assume strict diagonal dominance, we can apply Lemma 13 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

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

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

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

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

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

These two equalities together yield the desired result.

on $B^1 = \text{argmin}_{1\le i\le K} \partial_i V(\mathbf{0})$. Applying Lemma 13 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 $B^2 = \text{argmin}_{1\le i\le K} \partial_i V(n(t^1))$. So on and so forth. Since the sets $\emptyset = B^0, B^1, B^2, \dots$ are nested by construction, we eventually have $B^m = \{1, \dots, K\}$ for some m, and consequently $t^m = \infty$.

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

Finally, suppose Σ^{-1} is only weakly diagonally-dominant. The proof of Lemma 13 is still applicable, except that we can no longer conclude $\beta_i = \tilde{\alpha}_i$ is strictly positive. Thus the attention vector β is non-negative and potentially supported on a subset of B. Nonetheless, recall that $\tilde{\alpha}_i$ is proportional to the i-th row sum of $(\Sigma_{TL})^{-1}$ plus q_i. So we have $\beta_i > 0$ whenever $q_i > 0$. This implies that any source that has received attention in the past will receive positive attention at every future moment.

Using this property, we can redo the above proof of Theorem 2. First, the agent can achieve t-optimality up to time t^1 with a constant attention strategy supported on some set B^1, which is now a potential subset of $\text{argmin}_{1\le i\le K} \partial_i V(\mathbf{0})$. He can then maintain t-optimality from t^1 to t^2 with another constant attention level supported on some set $B^2 \subset \text{argmin}_{1\le i\le K} \partial_i V(n(t^1))$. So on and so forth, until $\text{argmin}_{1\le i\le K} \partial_i V(n(t^{m-1})) = \{1, \dots, K\}$, at which point $B^m = \{1, \dots, K\}$ as well by Lemma 12. This process must end, because by construction the set of minimizers $\text{argmin}_{1\le i\le K} \partial_i V(n(t^0))$, $\text{argmin}_{1\le i\le K} \partial_i V(n(t^1)), \dots$ strictly expands.

By the earlier analysis, we have the crucial observation that the sets $B^1, B^2, \dots$ are weakly increasing. From the perspective of information acquisition, we can in fact assume these sets are strictly nested because the attention level is unchanged at time t^l whenever $B^{l+1} = B^l$.[28] Theorem 2 then follows from the same argument as before, after including

[28] To see this, define $B^* = \text{argmin}_i \partial_i V(n(t^{l-1}))$ and $B^{**} = \text{argmin}_i \partial_i V(n(t^l))$. Then the constant attention level between times t^l and t^{l+1} is proportional to the vector α^{**}, which has the property that the payoff-relevant state ω can be written as α^{**} times (inner-product) the states in B^{**} plus a residual term orthogonal to these states. If $B^{l+1} = B^l$, then α^{**} which is supported on B^{l+1} is also supported on $B^l \subset B^*$. Since the aforementioned residual term is orthogonal to the smaller set of states in B^*, we deduce that α^{**}

some redundant times.

C.6 Tightness of $\frac{1}{2K-3}$

Here we provide an example to show that the constant $\frac{1}{2K-3}$ in Assumption 2 is *tight for the existence of a uniformly optimal strategy*. In other words, for any $\rho > \frac{1}{2K-3}$ we demonstrate a prior covariance matrix Σ satisfying $|\Sigma_{ij}| \leq \rho \cdot \Sigma_{ii}$ for all $i \neq j$, as well as *some* weight vector α, such that uniform optimality cannot be achieved given the primitives Σ and α.

Let Σ have diagonal entries 1 and off-diagonal entries $-\rho$, with $\rho > \frac{1}{2K-3}$. This means the agent's prior belief over the attributes is symmetric. We also choose $\alpha_2 = \cdots = \alpha_K = 1$ and α_1 equal to a small positive number.

For this problem, we will show that the t-optimal attention vector $n(t)$ is *not* monotonic over time. Note that the last $K-1$ sources have symmetric prior and symmetric payoff weights. Thus, the posterior variance function $V(q_1, q_2, \ldots, q_K)$ is symmetric in its last $K-1$ arguments. This implies that the t-optimal vector $n(t)$ must satisfy $n_2(t) = \cdots = n_K(t)$; otherwise it would not be unique.

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

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

That is, the agent optimally divides attention between signal 1 and the remaining signals (which always receive equal attention).

The posterior belief of such a agent can be derived by Bayesian updating on the following K normal signals: $\theta_1 + \mathcal{N}\left(0, \frac{1}{q_1}\right)$ and $\theta_i + \mathcal{N}\left(0, \frac{1}{q_2}\right)$ for $2 \leq i \leq K$. We now show that in terms of predicting the payoff-relevant state $\alpha_1\theta_1 + \sum_{i>1}\theta_i$, the agent's belief is the same as if he had observed just *two* signals: $\theta_1 + \mathcal{N}\left(0, \frac{1}{q_1}\right)$ and $\frac{1}{K-1}\sum_{i>1}\theta_i + \mathcal{N}\left(0, \frac{1}{(K-1)q_2}\right)$, which is the average of the last $K-1$ signals previously. Clearly, those $K-1$ signals provide at least as much information as their average, so we focus on the converse. Indeed, by symmetry we know that the agent's posterior belief about $\sum_{i>1}\theta_i$ is unchanged whether he observes the $K-1$ signals or their average. Moreover, conditional on $\sum_{i>1}\theta_i$, the $K-1$ signals only provide information about the differences $\theta_i - \theta_j$ (with $i, j > 1$). Since $\theta_i - \theta_j$ is independent from θ_1 conditional on $\sum_{i>1}\theta_i$ (it is in fact independent from both), the extra information does not change the conditional belief about θ_1. As such, the $K-1$ signals $\theta_i + \mathcal{N}\left(0, \frac{1}{q_2}\right)$ for $i > 1$ are equally informative about the payoff-relevant state as their average.[29]

coincides with the constant attention level α^* between times t^{l-1} and t^l.

[29] This can also be proved by directly computing the posterior covariance matrix. We omit the details.

Given this equivalence, we can relate t-optimality in the original information environment with K sources to a smaller problem with just two sources. Specifically, define $\theta_1^* = \theta_1$, $\theta_2^* = \frac{1}{K-1}\sum_{i>1}\theta_i$, $\alpha_1^* = \alpha_1$, $\alpha_2^* = K-1$. Then the payoff-relevant state ω is rewritten as $\alpha_1^* \cdot \theta_1^* + \alpha_2^* \cdot \theta_2^*$. The discussion in the preceding paragraph shows that the posterior variance function V^* in this two-by-two problem satisfies

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

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

We compute the prior covariance matrix Σ^* to be

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

In particular, since $\rho > \frac{1}{2K-3}$, $\Sigma_{21}^* + \Sigma_{22}^*$ is negative. Thus if $\alpha_1^* = \alpha_1$ is sufficiently small, this matrix violates Assumption 1. By Lemma 9, we conclude that the t-optimal attention vectors in this smaller problem are not monotonic. The same holds for the original problem, completing the proof.

D Proof of Proposition 1

D.1 Proof Outline

As discussed in the main text, we only need to prove that each source receives infinite attention (Lemma 1) and that Theorem 2 applies at any posterior belief after each source is sufficiently sampled. The latter is easy: Observe that the agent's posterior *precision* matrix is given by $\Sigma^{-1} + Q$, where Q is the diagonal matrix with entries $q_1, \ldots, q_K$. As $q_i \to \infty$ to each i, clearly the matrix $\Sigma^{-1} + Q$ is diagonally-dominant. So the conclusion of Theorem 2 holds.[30]

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

Lemma 14. *Fix Σ and α. Given any $\underline{q} \in \mathbb{R}_+$, there exists $\overline{q} \in \mathbb{R}_+$ such that the cumulated attention vectors $q(t)$ under the optimal strategy have the following property: Whenever $q_i(t) < \underline{q}$ for some source i, it holds that $q_j(t) \leq \overline{q}$ for every source i.*

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

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

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

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

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

Our proof method will be to construct a profitable deviation strategy (of how to allocation attention) following this history, so that optimality is violated. Thanks to the main theorem of Greenshtein (1996), any deviation strategy is profitable so long as it *decreases the posterior variance of* ω *at all future times.* Given a deviation strategy, let $\tilde{q}(t)$ denote the induced cumulated attention vector, which is distinguished from $q(t)$. Then the deviation is profitable whenever the following inequality holds:[31]

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

D.2 The Deviation

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

(a) Suppose that the original strategy S devotes positive attention to source K at time T. Then under the deviation strategy, the agent *diverts this attention (evenly) toward those sources* i *with* $\tilde{q}_i(T) < \underline{q}$.[32] If no such source exists, the deviation strategy devotes the same amount of attention to source K.

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

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

(b) Suppose that the original strategy devotes attention to some source in $k+1,\dots,K-1$. Then the deviation strategy devotes the same attention to this source.

(c) Suppose that the original strategy devotes attention to source $i \le k$. If $\tilde{q}_i(T) < \underline{q}$ or $\tilde{q}_i(T) = q_i(t)$, then the deviation strategy also observes source i. Otherwise we have $\tilde{q}_i(T) = \underline{q} > q_i(T)$, and in this case the deviation strategy *diverts this amount of attention to source K instead.*

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

D.3 Four Kinds of Sources

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

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

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

2. For sources $i \in I_2 \subset \{1,\dots,k\}$, we have $q_i < \tilde{q}_i = \underline{q}$. These are the sources that have reached the target level $\underline{q}$ under the deviation strategy, but not under the original strategy. Let x_i denote the difference $\tilde{q}_i - q_i$, then by construction we have $x_i \le x$, which is defined above.

3. For sources $i \in I_3$, we have $q_i = \tilde{q}_i \ge \underline{q}$. These include the sources $k+1,\dots,K-1$, which the deviation strategy does not affect. Also included are those sources in $1,\dots,k$ that have reached cumulated attention $\underline{q}$ under both the original and deviation strategies.

4. Finally source K is the only source with $q_i > \tilde{q}_i$. In fact we have

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

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

D.4 Comparison of Posterior Variances

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

Lemma 15. *There exists a positive constant C_H depending only on Σ and α, such that for all $q_1, \dots, q_K \geq 0$,*

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

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

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

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

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

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

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

$$V(\tilde{q}) = V(\tilde{q}_j, \tilde{q}_{-j}) \leq V(q_j, \tilde{q}_{-j}) + (\tilde{q}_j - q_j) \cdot \partial_j V(\tilde{q}) \leq V(q_j, \tilde{q}_{-j}) - \frac{x_j \cdot C_L}{\underline{q}^2} \leq V(q_1, \dots, q_{K-1}, \tilde{q}_K) - \frac{x_j \cdot C_L}{\underline{q}^2}. \tag{4}$$

The first inequality uses the convexity of V. The second inequality uses the second part of Lemma 15 (which applies because $\tilde{q}_i \geq \underline{q}$ for all i when I_1 is empty), as well as $\tilde{q}_j = \underline{q}$ (since $j \in I_2$). The last inequality uses the monotonicity of V and $\tilde{q}_i \geq q_i$ for all but the last source.

On the other hand, we also have

$$V(q) \geq V(q_1, \dots, q_{K-1}, \tilde{q}_K) + (q_K - \tilde{q}_K) \cdot \partial_K V(q_1, \dots, q_{K-1}, \tilde{q}_K) \geq V(q_1, \dots, q_{K-1}, \tilde{q}_K) - \frac{(K-1)x_j \cdot C_H}{(\tilde{q}_K)^2}, \tag{5}$$

where the first inequality is by convexity, and the second uses the first part of Lemma 15 and $q_K - \tilde{q}_K = \sum_{i \in I_2} x_i \leq (K-1)x_j$ by our choice of j.

Recall that $\tilde{q}_K \geq \overline{q}$. Thus whenever $\overline{q}$ is much larger compared to $\underline{q}$, the above inequalities (4) and (5) imply that $V(\tilde{q}) < V(q)$, as we desire to show.

Next we consider the case where I_1 is non-empty. By the third part of Lemma 15, we can choose $j \in I_1$ such that $\partial_j V(\tilde{q}) \leq \frac{-C_L}{\underline{q}^2}$. Then, similar to (4) we have

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

with x replacing the role of x_j. Likewise, we have the following analogue of (5):

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

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

Hence we are once again able to deduce $V(\tilde{q}) < V(q)$ so long as $\tilde{q}_K \geq \overline{q}$ is much larger than $\underline{q}$. This completes the proof of Proposition 1 modulo Lemma 15.

D.5 Proof of Lemma 15

In light of Lemma 4, the key will be to estimate the size of the different coordinates of $\gamma = (\Sigma^{-1} + Q)^{-1} \cdot \alpha$.

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

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

Comparing the i-th coordinate on both sides, we have $\alpha_i = e_i' \cdot \Sigma^{-1} \cdot \gamma + q_i \gamma_i$. This then implies that the product $q_i \gamma_i$ is bounded across different possible q. Since $\partial_i V(q) = -\gamma_i^2$, the first part of Lemma 15 is proved.

For the second part, we use the matrix identity

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

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

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

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

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

E Proof of Proposition 2

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

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

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

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

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

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

$$\left.\frac{\partial U_1}{\partial \sigma_1}\right|_{(\sigma_1,\sigma_2)} = \frac{\sigma_2}{r(\sigma_1+\sigma_2)^2}e^{-rt_1^*}\left(1 - \frac{r(\sigma_1+\sigma_2)(2\sigma_1-\sigma_2)}{1-\rho}\right). \tag{6}$$

Equilibrium requires

$$r(\sigma_1+\sigma_2)(2\sigma_1-\sigma_2) \le 1-\rho \quad \text{with equality if } \sigma_1 < \sigma_2. \tag{7}$$

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

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

The derivative with respect to its action σ_2 is

$$\left.\frac{\partial U_2}{\partial \sigma_2}\right|_{(\sigma_1,\sigma_2)} = \frac{\sigma_1}{r(\sigma_1+\sigma_2)^2}e^{-rt_1^*}\left(1 - \frac{r(\sigma_1+\sigma_2)\sigma_2}{1-\rho}\right). \tag{8}$$

Equilibrium requires

$$r(\sigma_1 + \sigma_2)\sigma_2 \geq 1 - \rho \quad \text{with equality if } \sigma_1 < \sigma_2 \tag{9}$$

Combining (7) and (9), it is immediate that any pure strategy equilibrium must have $\sigma_1 = \sigma_2$.[33] Then the two inequalities (7) and (9) together give $\sigma_1 = \sigma_2 = \sqrt{\frac{1-\rho}{2r}} = \sigma^*$ as desired. Moreover, this is an equilibrium because (6) and (8) show that any deviation (not just local deviations) is not profitable. In fact, given $\sigma_j = \sigma^*$, the unique best response of source i is to choose the same σ_i. Since the game has a constant sum, this proves that the pure strategy equilibrium we have found is the unique equilibrium, pure or mixed.

F Many Competing Providers

Here we demonstrate how the game in Section 4 generalizes to the case of $K > 2$ competing data sources. We maintain essentially the same setup, except that the agent seeks to predict $\theta_1 + \cdots + \theta_K$ where the precision of information about each θ_i is controlled by a separate data provider. Using the transformation $\tilde{\theta}_i = \frac{\theta_i}{\sigma_i}$, we can reduce the agent's information acquisition problem to our main model with prior covariance matrix

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

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

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

Lemma 16. *Suppose $\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_K$. For $1 \leq k \leq K-1$, define*

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

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

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

Using this result, it is straightforward to solve for the symmetric pure strategy equilibrium of the game. Indeed, suppose the other sources all choose σ_2; then, source 1's payoff when choosing $\sigma_1 \leq \sigma$ is given by

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

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

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

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

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

References

Azevedo, E., A. Deng, J. L. M. Olea, J. Rao, and G. Weyl (2019): "A/B Testing with Fat Tails," Working Paper.

Bardhi, A. (2018): "Optimal Discovery and Influence Through Selective Sampling," Working Paper.

Bergemann, D. and J. Välimäki (2008): "Bandit Problems," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Basingstoke: Palgrave Macmillan.

Bonatti, A. and G. Cisternas (2019): "Consumer Scores and Price Discrimination," Working Paper.

Callander, S. (2011): "Searching and Learning by Trial and Error," *American Economic Review*, 101, 2277–2308.

Carlson, D. and T. L. Markham (1979): "Schur Complements of Diagonally Dominant Matrices," *Czechoslovak Mathematical Journal*, 29, 246–251.

Che, Y.-K. and K. Mierendorff (2019): "Optimal Dynamic Allocation of Attention," *American Economic Review*, forthcoming.

Fudenberg, D., P. Strack, and T. Strzalecki (2018): "Stochastic Choice and Optimal Sequential Sampling," *American Economic Review.*

Garfagnini, U. and B. Strulovici (2016): "Social Experimentation with Interdependent and Expanding Technologies," *Review of Economic Studies*, 83, 1579–1613.

Gittins, J. C. (1979): "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society, Series B*, 148–177.

Greenshtein, E. (1996): "Comparison of Sequential Experiments," *The Annals of Statistics*, 24, 436–448.

Hansen, O. H. and E. N. Torgersen (1974): "Comparison of Linear Normal Experiments," *The Annals of Statistics*, 2, 367–373.

Hébert, B. and M. Woodford (2018): "Rational Inattention with Sequential Information Sampling," Working Paper.

Immorlica, N., I. Kash, and B. Lucier (2018): "The Evolving Value of Data," Working Paper.

Liang, A., X. Mu, and V. Syrgkanis (2019): "Optimal and Myopic Information Acquisition," Working Paper.

Mayskaya, T. (2019): "Dynamic Choice of Information Sources," Working Paper.

Moscarini, G. and L. Smith (2001): "The Optimal Level of Experimentation," *Econo-*

metrica, 69, 1629–1644.

Ratcliff, R. and G. McKoon (2008): "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks," *Neural Computation*, 20, 873–922.

Steiner, J., C. Stewart, and F. Matějka (2017): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *Econometrica*, 85, 521–553.

Zhong, W. (2018): "Optimal Dynamic Information Acquisition," Working Paper.